Electronic Design®

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

MARCH/APRIL 2019
electronicdesign.com

BODY CONTROL MODULES

The Automotive Networking Solution p12

$10.00

Broadcom Optocouplers

Certified to the IEC 60747-5-5 Safety Standard with Proven High Voltage Life Time Performance

Are the High Voltage Circuits in Your EV/HEV using IEC Certified Components for Safe Isolation and Long Working Life Time?

Broadcom offers a wide range of digital, IPM interface, high voltage Photo-MOSFET, gate drive optocouplers and isolation amplifiers for the automotive industry. The first supplier to introduce 125°C optocouplers for automotive applications, Broadcom R²Coupler® products are AEC-Q100 qualified with applications in CANBus interface isolation, motor inverter drives, high voltage leakage detection, AC-DC and DC-DC converters for battery chargers, battery/inverter voltage monitoring, fault detection and battery management systems.

- Certified to Regulatory Standards IEC 60747-5-5, CSA Notice 5A, UL 1577
- High voltage isolation leadership and life-time proven in field for more than 40 years

BROADCOM®
connecting everything®

Altech® PUSH-IN PLUGS

Printed Circuit Board Terminal Blocks

Time Saving
Push-In Technology

Easy wire insertion
No tool required !

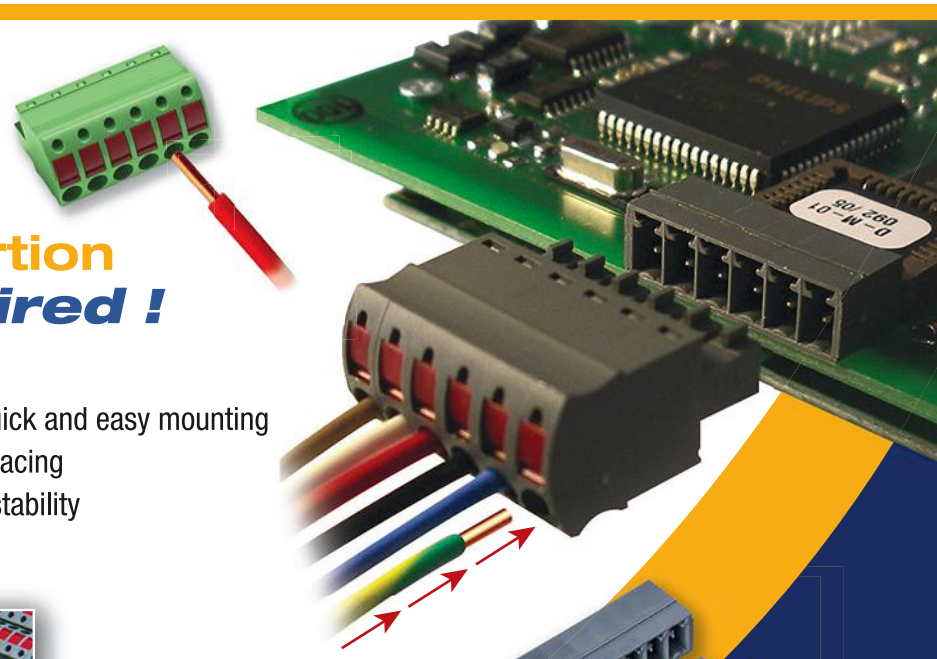

FEATURES

- Unique ergonomic grip design for quick and easy mounting
- 3.50, 3.81, 5.00 and 5.08mm pin spacing
- Screw flange version for additional stability
- Imprinting and coding available

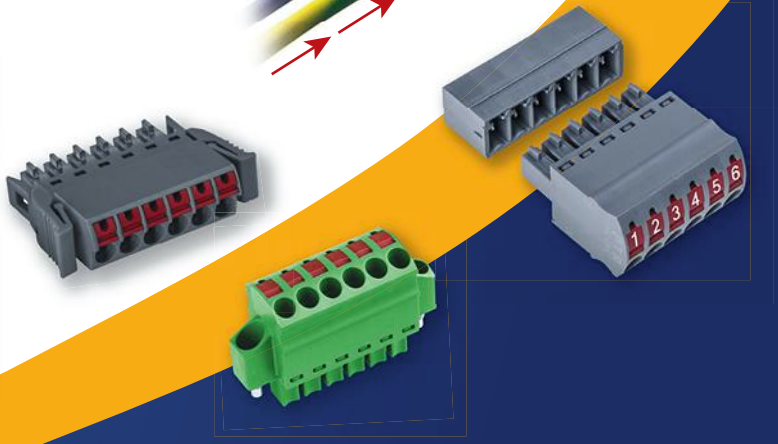

c UL us

VDE IEC

Complete line of Plug-In Connectors

- All pin spacings
- Proven screw clamp design
- Compatible headers

Altech Corp.®

Your Source for Automation & Control Components

AltechCorp.com/PCBPushin
908.806.9400 for all the details & catalog

IN THIS ISSUE

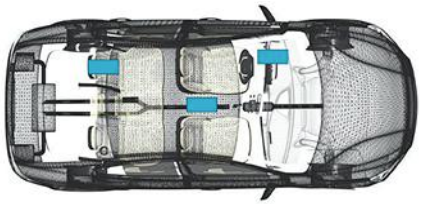

12

FEATURES

28

41

45

COLUMNS & DEPARTMENTS

EDITORIAL MISSION:
To provide the most current, accurate, and in-depth technical coverage of the key emerging technologies that engineers need to design tomorrow's products today.

ELECTRONIC DESIGN (ISSN 0013-4872) is published bi-monthly in Jan/Feb, Mar/Apr, May/June, July/Aug, Sept/Oct and Nov/Dec by Informa Media Inc., 9800 Metcalf Ave., Overland Park, KS 66212-2216. Paid rates for a one-year subscription are as follows: $60 U.S., $90 Canada, $120 International. Periodicals postage paid at Kansas City, MO, and additional mailing offices. Editorial and advertising addresses: ELECTRONIC DESIGN, 605 Third Avenue, New York, NY 10158. Telephone (212) 204-4200. Printed in U.S.A. Title registered in U.S. Patent Office. For subscriber services or to order single copies, write to Electronic Design, PO Box 2100, Skokie, IL 60076. POSTMASTER: Send change of address to Electronic Design, PO Box 2100, Skokie, IL 60076. Canadian Post Publications Mail agreement No. 40612608. Canada return address: IMEX Global Solutions, P.O. Box 25542, London, ON N6C 6B2.

ON ELECTRONICDESIGN.COM

The Biggest Utility-Scale Solar Farms on Earth

Superpowers on every continent are in an all-out sprint for the title of the largest utility-scale solar project on the planet. Who will win the race for clean energy, and what will we lose?

https://www.electronicdesign.com/power/biggest-utility-scale-solar-farms-earth

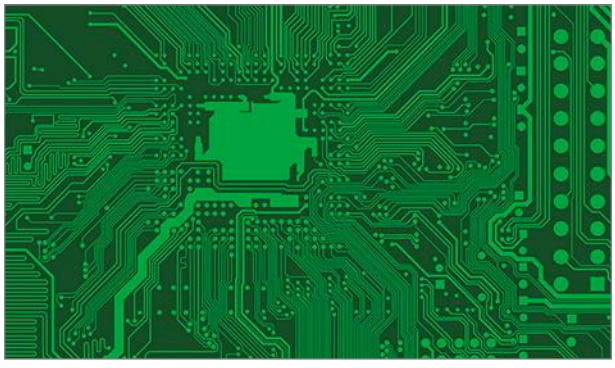

Via-In-Pad Guidelines for PCBs

Though generally frowned upon, here's a quick rundown on why "via in pad" may be an option to consider in certain designs, and how to place them properly.

https://www.electronicdesign.com/analog/pad-guidelines-pcbs

VHS or Betamax…CCIX or CXL…So Many Choices

Intel provided the technology behind Compute Express Link, a high-speed interconnect for hardware accelerators. How will it compete with Cache Coherent Interconnect for Accelerators?

https://www.electronicdesign.com/industrial-automation/vhs-or-betamax-ccix-or-cxl-so-many-choices

What's All This Subthreshold Stuff, Anyhow?

Operating transistors and ICs in the "leakage" zone lowers power consumption by an order of magnitude. Contributing Tech Expert Paul Rako explains.

https://www.electronicdesign.com/analog/what-s-all-subthreshold-stuff-anyhow

ElectronicDesign®

FOLLOW US ON FACEBOOK AND TWITTER

(ElectronicDesign) (@ElectronicDesgn)

PUSH PERFORMANCE TO THE TOP

PCB Terminal Blocks and Connectors for Power Electronics

- Compact dimensions and high current carrying capacity – up to 85 A
- Fast, easy connection with Push-in CAGE CLAMP®
- Optimal handling for any application
- Comprehensive range covering 24 - 4 AWG

www.wago.us/powerelectronics

WAGO

ElectronicDesign

MARCH/APRIL 2019

INFORMA MEDIA INC.

605 THIRD AVENUE
NEW YORK, NY 10158 USA **T** | 212.204.4200

informa

Electronic Design | Machine Design | Microwaves & RF | Source ESB | Hydraulics & Pneumatics | Global Purchasing | Distribution Resource | Power Electronics | Defense Electronics

GPS disciplined *10 MHz reference* ...

... and so much more !

- GPS/GNSS disciplined 10 MHz
- TCXO, OCXO or Rb timebase
- Time tagging to GPS and UTC
- Frequency counter with 12 digits/s
- Source out: sine, square, triangle & IRIG-B
- Built-in distribution amplifiers
- Ethernet and RS-232 interfaces

The FS740 GPS disciplined10 MHz reference delivers cesium equivalent stability and phase noise at a fraction of the cost.

It's host of features includes a 12-digit/s frequency counter, a DDS synthesized source with adjustable frequency and amplitude, built-in distribution amplifiers, and event time-tagging with respect to UTC or GPS.

The optional OCXO or rubidium clock (PRS10) provide better than –130 dBc/Hz phase noise.

FS740 ... $2495 (U.S. list)

SRS Stanford Research Systems
Tel: (408) 744-9040 • www.thinkSRS.com
www.thinkSRS.com/products/FS740.htm

Editorial

WILLIAM WONG | Senior Technology Editor
bill.wong@informa.com

AI Continues to Trend Smaller and Faster

My daughter competed in the Intel International Science and Engineering Fair for three years using custom-built robots that incorporated a Parallax Basic Stamp and Java Stamp along with a VGA camera that did basic object recognition with an 8 by 8 resolution. The reduced resolution was needed to process at only a few frames per second. The last project earned a second-place award. This was less than 20 years ago.

Contrast that platform with the JetBot *(Fig. 1)*, an open-source project that's built on NVIDIA's new Jetson Nano, a DIMM with the functionality of the Jetson TX1. The Jetson Nano costs only $129 in single quantities, but it can process up to eight 1080p video streams while running multiple neural networks doing object recognition in real time. Performance in this space has increased by more than a factor of 1000 in a little over a decade.

The thing is, the Jetson Nano just continues the artificial-intelligence/machine-learning (ML) trend that's essentially based in deep neural networks (DNNs). Platforms like the new BeagleBone AI *(Fig. 2)* take advantage of Texas Instruments' AM5729 with C66x DSP cores and an embedded vision engine (EVE) to handle ML applications. Though targeting a different class of applications, it can still process video streams in real time.

Image processing is a computationally demanding ML application, but far from the only AI application these days. Renesas' RX66x microcontroller is supported by the company's Failure Detection e-AI Solution. This provides ML motor-control support for up to four motors, allowing developers to add preventative maintenance support while reducing the amount of information that would be sent to the cloud by massaging data locally.

Microcontrollers are gaining ML support as well. For example, STMicroelectronics' STM32Cube.AI helps developers create ML software for a range of STM32 platforms based on ARM's 32-bit Cortex-M family.

These examples are merely the tip of the iceberg when it comes to ML support. AI certainly isn't needed for all applications. However, a greater number can benefit from the technology even for applications where space and power are limited. The key to success doesn't simply surround hardware acceleration—it also involves the software support that's extensive when it comes to AI tools. It's the software support that makes these solutions and others like them stand out. ed

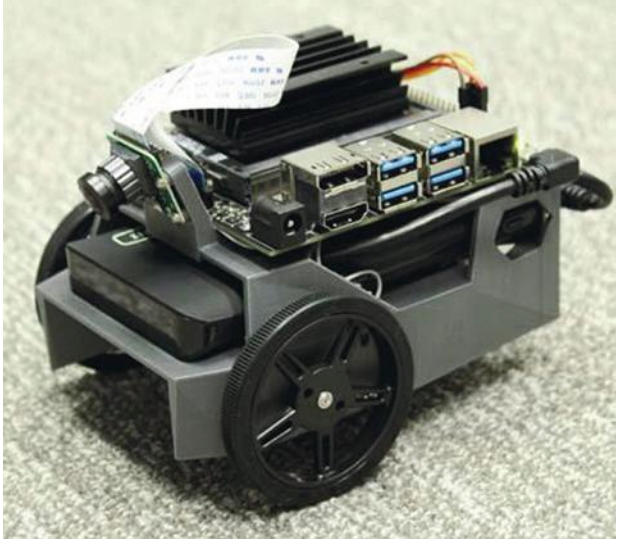

1. The JetBot open-source project takes advantage of NVIDIA's Jetson Nano.

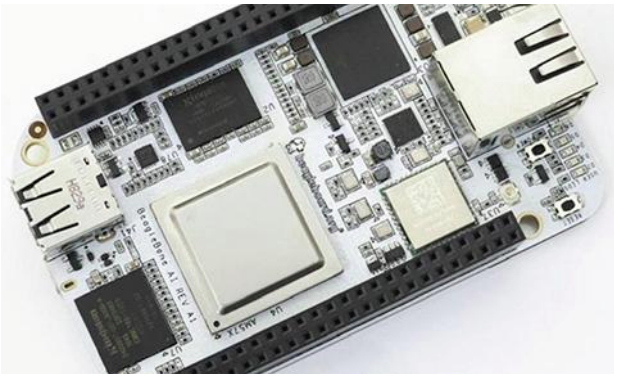

2. Texas Instruments' AM5729 with C66x DSP cores, plus an embedded vision engine, help the BeagleBone AI platform tackle machine-learning chores.

THINK™
HIGH PERFORMANCE

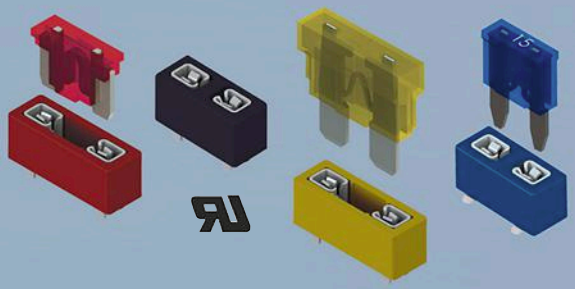

Color Keyed Automotive Blade Fuse Holders

• Available in colors to match fuse amperage colors for simplified fuse location • Available for Mini, LP Mini and Standard automotive blade fuses • Fully insulated contacts • Available in Thru Hole Mount (THM) rated for 30 Amps and Surface Mount (SMT) rated for 20 Amps • Available in Bulk or on Tape and Reel

IT'S WHAT'S ON THE INSIDE THAT COUNTS

KEYSTONE®
ELECTRONICS CORP.

www.keyelco.com • (516) 328-7500 • (800) 221-5510

REQUEST CATALOG!

VISIT US AT BOOTH 215 EDS 2019

News

INSIDE OUT—It's All in the Positioning

Head-mounted displays (HMDs) for virtual and augmented reality (VR/AR) need sensors to determine head movement and positioning, otherwise the effects provided by the HMD can range from annoying to sickening. Even latency can cause a problem even if the tracking is accurate.

What's needed is fast and accurate tracking. This is accomplished by using sensors on the HMD, called inside-out tracking, versus external sensors, or outside-in. The former has the advantage of simplifying the overall system, at least from a user's perspective.

1. Intel's T265 RealSense Tracking Camera combines a 3D RealSense camera system with the Movidius Myriad 2 machine-learning, video-processing system.

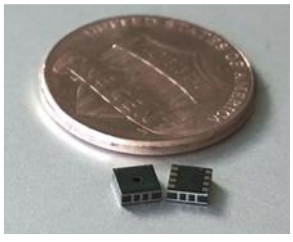

2. TDK's Chirp CH-101 ultrasonic transceiver is tiny and power-efficient. ▶

INSIDE-OUT 3D CAMERA DELIVERS V-SLAM

Two very different solutions have cropped up recently. One, from Intel, is based on its RealSense 3D camera. The RealSense family recently acquired a built-in IMU, although this makes the version of the camera useful for other reasons. The newest addition targets inside-out positioning. The RealSense T265 *(Fig. 1)* employs Intel's Movidius Myriad 2 machine-learning (ML) video-processing system.

The T265 implements a visual inertial odometry simultaneous localization and mapping (V-SLAM) system that tracks objects in the real world. It delivers six degrees of freedom (6DOF) of inside-out tracking support. Essentially, it tracks objects and their movement and position with respect to the camera. The RealSense sensors provide the 3D depth information that's then processed by the Movidius SoC. Fish-eye lenses provide a wide, 170-degree field of view (FOV). Pricing starts at $199.

The T265 can be mounted on an HMD to provide V-SLAM support. It also will deliver video and depth information for other aspects of an AR/VR application. On top of that, the sensor will be useful for robotics and drones, especially in GPS-restricted environments.

CHIRPING ULTRASONICALLY

Chirp is TDK's latest acquisition. Instead of infrared, 3D cameras, Chirp's SonicTrack solution employs tiny ultrasonic transceivers, which can often outperform infrared sensors in a range of applications. They're able to provide range information by themselves, although not to the degree of a RealSense 3D camera. On the other hand, the CH-101 chip is tiny *(Fig. 2)*.

Ultrasonic sensors have been used for range sensing applications in a number of areas like proximity sensors on cars. Things get a bit more interesting when multiple sensors are combined with the kind of intelligence Intel added to its RealSense cameras.

One application that TDK is working on delivers inside-out positioning information about the 3D hand controllers for AR/VR applications; in particular, HTC's Vive HMD *(Fig. 3)*. Three Chirp sensors are mounted on each controller and additional sensors are on the HMD. These are used cooperatively to provide accurate relative positioning information via triangulation. Thus, the controllers can be used without initial calibration and set up.

The CH-101 is an omnidirectional sensor. A single sensor is only able to provide range information, not positional information. Likewise, the sensor can be used as a transmitter, receiver, or both albeit one mode at a time.

TDK's InvenSense Fusion software can combine fuse sensor data from a number of devices, including the new ICM42688 and ICM42686 6-axis motion sensors. The software is designed to run on standard microcontrollers. ■

3. TDK is working with HTC and its Vive VR HMD to provide inside-out tracking of the hand controllers with the VR HMD.

From Concept to Reality

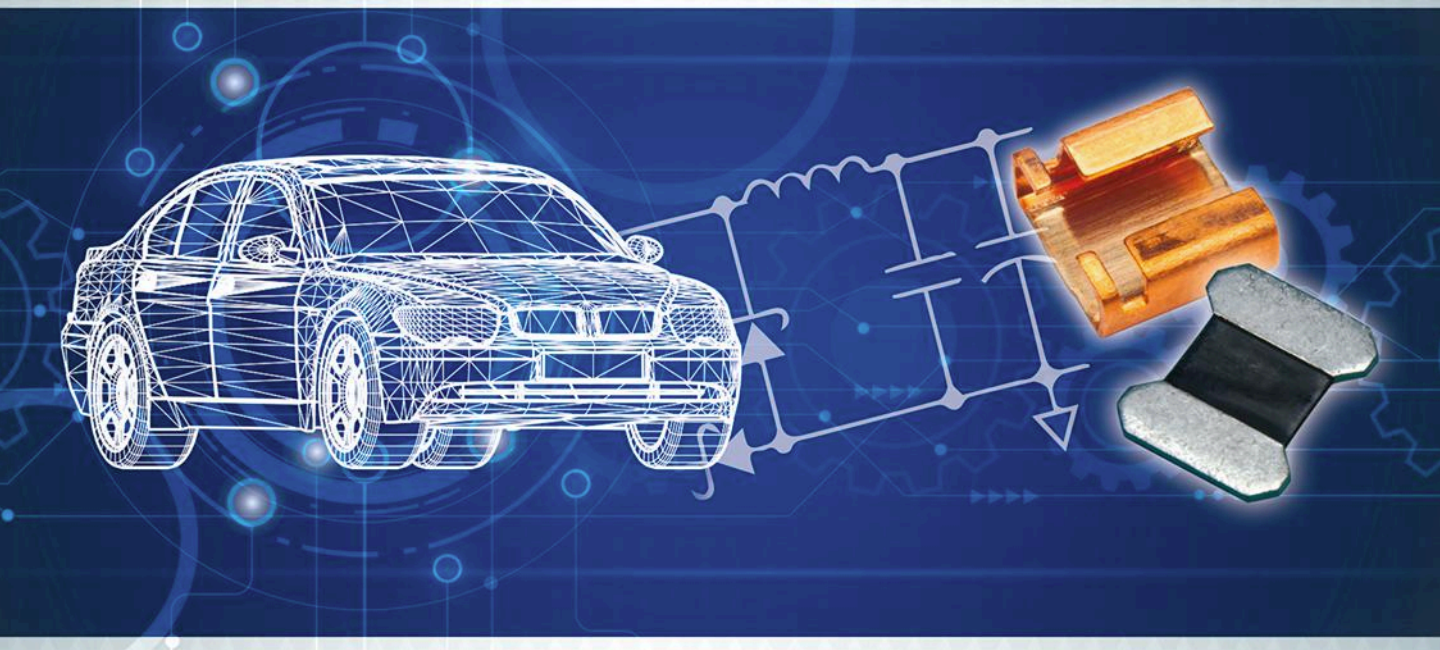

Our new passive component solutions will help you make the leap!

Whether you're designing the car of tomorrow or connecting the Internet of Things, you need cutting-edge product solutions to bring your design to fruition. At KOA Speer, we're the ideal partner to help you do just that. Our constantly expanding line of passive components will give you the solution to make your concept a reality!

More than just Resistors - yes we're the world leader in thick film resistors, but our product line is broad and growing every day, including:

- Shunt/Current Sense Resistors
- Wide Terminal Resistors
- High Voltage Resistors
- Thin Film Chip Resistors
- Surge Resistors

Passive Component Solutions - our engineering group will work with your design team to develop the passive component solution you need. Many of the hundreds of new products we've introduced in recent years are the result of such collaboration.

Award Winning Passive Supplier - we're the industry's most recognized and awarded supplier for achieving the highest product quality, on-time delivery and responsive customer service.

NEW! From KOA Speer

Precision Thick Film Resistors (RK73G)

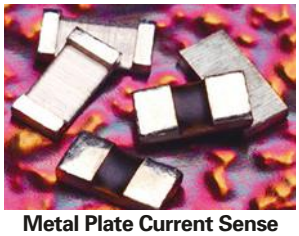

Metal Plate Current Sense Resistors (TLR, TLRZ & TLRH)

Wide Terminal Current Sense Resistors (WK73S & WU73)

KOA SPEER ELECTRONICS, INC.

Make the Leap from Concept to Reality, visit KOASpeer.com

Design Feature

JOHN GRIFFITH | Engineer, Automotive Body Electronics and Lighting, Texas Instruments

NETWORKING in Automotive Body Control Modules

The role of BCMs in automotive design continues to expand to handle the increase in features and networking bandwidth requirements. Here's a look at the architectures and interfaces in play.

Automotive body control modules (BCMs) are present in all modern vehicles to handle comfort, security, and lighting functions *(Fig. 1)*. They have quickly transformed over the last few years, driven by the growing number of features, the desire to replace fuses and relays, the increased bandwidth requirements on in-vehicle networks, and the shrinking size requirements of original equipment manufacturers (OEMs).

This article reviews the evolution of networking in BCMs, how different vehicle architectures affect in-vehicle networks, and common networking interfaces and requirements.

VEHICLE ARCHITECTURES

The increasing automotive body loads distributed around the vehicle have created control, load, and networking issues. Some solutions control the loads from one central module; others control each load or a subset of loads locally. This creates a range of possible vehicle architectures.

The body-domain architecture integrates most body functions into one large central-control module, including a message-routing central gateway *(Fig. 2)*. In addition, this centralized architecture often requires a large number of low- to mid-bandwidth control buses like local interconnect network (LIN) and controller area network (CAN). It's becoming common to have more than a dozen LIN networks and more than eight CAN networks in a centralized BCM-plus-gateway application.

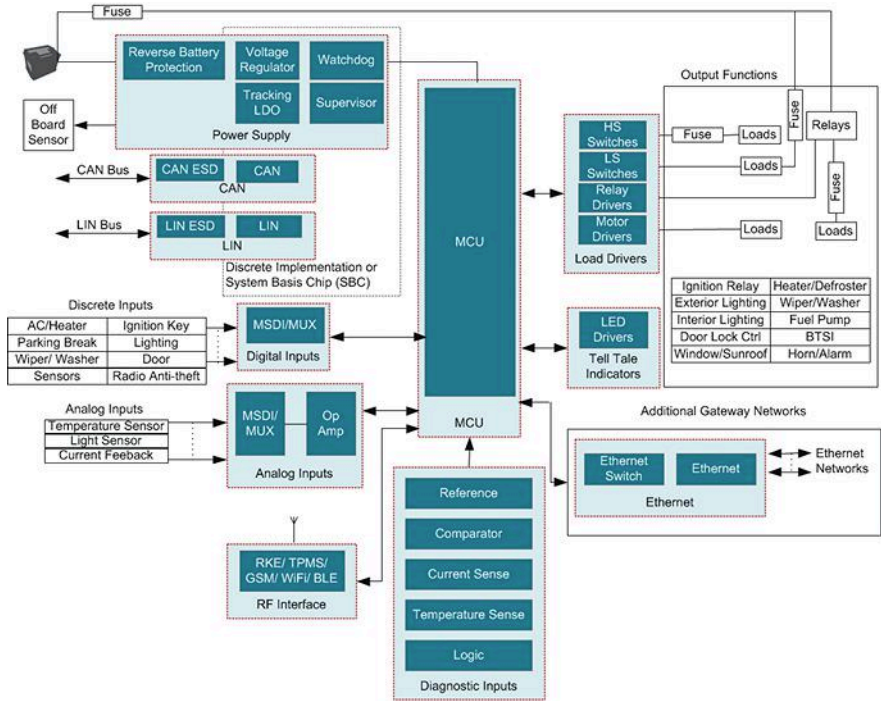

1. The next generation of automotive body control modules (BCMs) provides a wide range of connectivity.

It's also common to see interfaces like Ethernet used as the high-bandwidth interface in centralized BCM-plus-gateway modules. Ethernet enables the distribution of large amounts of data around the vehicle between different domain controllers or to external tools.

Other body-domain architectures decentralize the body domain, with multiple smaller distributed BCMs around the vehicle. For example, in the three-BCM architecture *(Fig. 3)*, one BCM is placed in the rear of the vehicle, a second goes in the center of the vehicle, and a third is in the front of the vehicle. Placing multiple modules around the vehicle enables you to locate them closer to the inputs that need monitoring and the loads that need controlling. This can often result in simpler wiring, with only power and networking shared between control modules.

By handling a smaller number of inputs and a smaller number of loads, these modules can also be simpler from both a power and networking perspective, containing just a single regulator and a single interface, such as LIN or CAN.

Populating and depopulating transceiver integrated circuits and using single, dual, and quad transceivers to optimize solution size scales the number of interfaces on the BCMs so tthat hey can be used across different car platforms or trim levels.

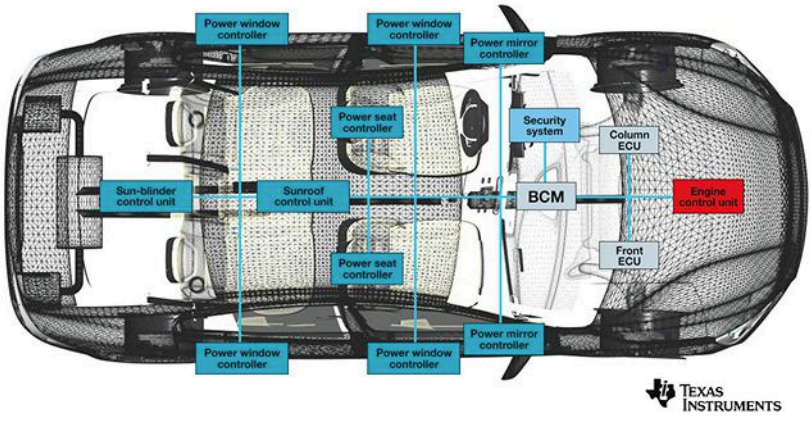

2. The body-domain architecture integrates a majority of body functions into one large central-control module.

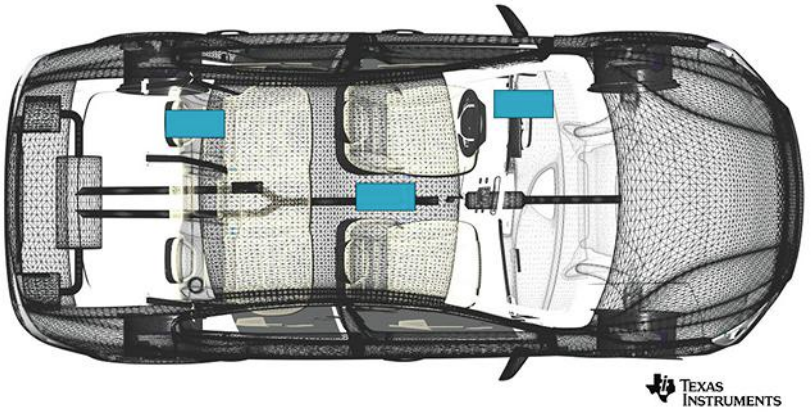

3. A three-BCM architecture has one BCM in the rear of the vehicle, a second in the center, and a third in the front. Placing multiple modules around the vehicle places them closer to the inputs that require monitoring and the loads that need controlling, resulting in simpler wiring.

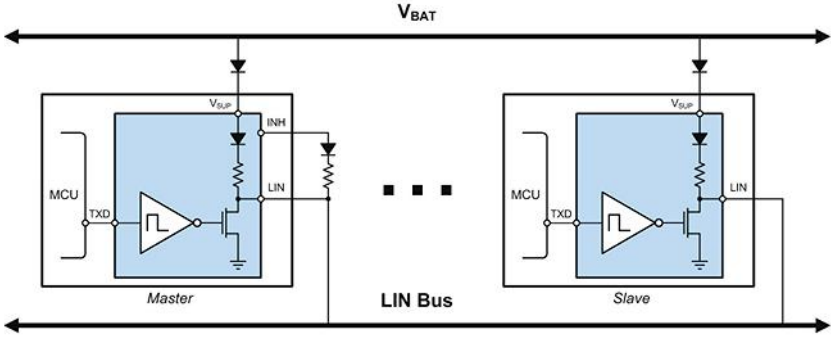

4. LIN targets human-machine interface applications where latency on the order of tens of milliseconds is more than tolerable.

VEHICLE INTERFACES

To optimize the complexity and cost of in-vehicle networking, a number of interfaces have been developed over the years.

LIN

In direct response to the growing integration of small electronic comfort features into vehicles, the first LIN standard was released in 1999. The idea was to offer a single-wire, low-cost, low-speed interface that could be implemented on

a small microcontroller with a simple universal asynchronous receiver transmitter peripheral *(Fig. 4)*.

LIN targets human-machine interface applications where latency on the order of tens of milliseconds is more than tolerable. The master module is responsible for polling the status of each slave on the network on a pre-defined schedule. With one central master node on the network handling all scheduling, syncing, and clocking, the slaves can be as low cost as possible. Master multi-slave polling networks lower the total system cost. However, the delay for an event to be processed can be a maximum of the time it takes to run through the full schedule table.

Since there are master and slave modules, some of the key transceiver features will depend on the role the transceiver will have on the network. For example, most slave modules need to respond to a LIN wakeup signal broadcasted over the bus. This can require a low-power wake receiver and an inhibit (INH) output pin for enabling and disabling a local regulator. It would allow slave nodes to be in the lowest possible power state, with only the transceiver monitoring the LIN bus and reducing system supply current while the vehicle is off.

On the other hand, a LIN master module is responsible for broadcasting the wakeup signal and therefore doesn't need a low-power wake receiver and INH output pin. BCMs often connect to large numbers of LIN clusters, thus necessitating many LIN master transceivers. For BCM applications like this, it's important to have transceiver families with single, dual, and quad product offerings to easily scale designs in a size-optimized manner.

In order to earn OEM approvals, other important features include the ability to handle radio-frequency (RF) immunity, transient, and electrostatic discharge, and to have low emissions. Depending on the vehicle type, you may also need different bus-fault protection voltages. For example, 12-V passenger vehicles require bus-fault protection voltages of up to 40 V, but 24-V commercial vehicles like semi-trucks and buses require higher bus-fault protection voltages of up to 60 V or more.

CAN

CAN started in the 1980s and has continued to evolve. Used by OEMs worldwide, CAN is a tried-and-true distributed control, serial, two-wire differential network standard developed for automotive and industrial applications. A differential topology leads to a more robust interface that's less susceptible to RF interference and has lower emissions due to the balanced differential signaling *(Fig. 5)*.

CAN is a multi-master protocol with nondestructive, bit-wise arbitration. Any CAN node on the network can access the bus when it's idle; if multiple nodes try to do so at the same time, the higher-priority message will win arbitration without any data errors. Other important features of the CAN standard include error checking, fault confinement, and clock synchronization.

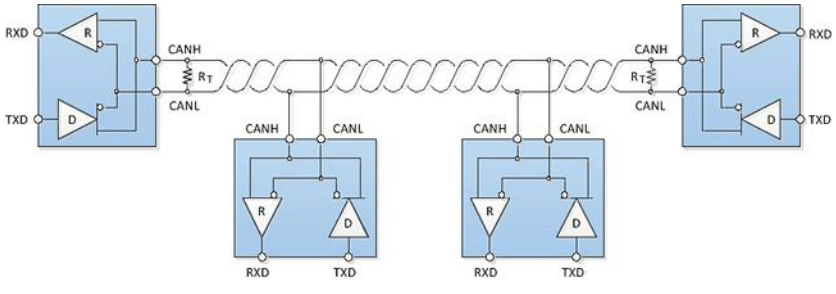

5. CAN is a multi-master protocol with nondestructive, bit-wise arbitration.

Take the path of least resistance

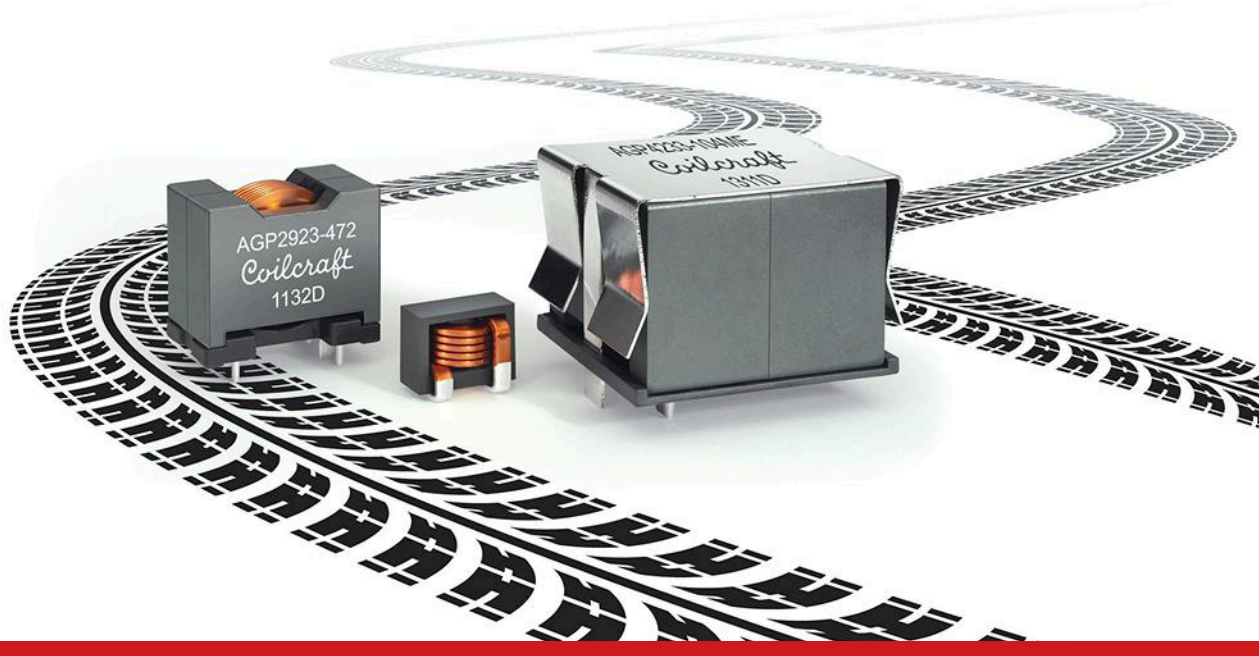

Increase the efficiency of your power electronics systems with our ultra-low DCR power inductors

Our AGP Family of high current power inductors feature flat wire construction for exceptionally low DC and AC resistance, letting your power systems run cooler and more efficiently.

They're offered with inductance values as high as 470 µH and current ratings up to 108 Amps.

AEC
Q200 Grade1

And all AGP Family inductors meet stringent AEC-Q200 Grade 1 quality standards, making them suitable for automotive and other harsh-environment applications.

Visit **www.coilcraft.com/AGP** to learn more and arrange a test drive. Free samples are always just around the corner!

Coilcraft®

WWW.COILCRAFT.COM

coilcraftdirect.com
No min order. Next day delivery.

Classical CAN, as the original International Organization for Standardization (ISO) 11898 CAN standard is often referred to now, has a maximum data rate of 1 Mb/s and a maximum data payload of 8 bytes. As with the LIN bus, fault protection, transient protection, electromagnetic compatibility (EMC) performance, and packaging options are key differentiators.

CAN with Flexible Data Rate

Continuing along the trend of increasing bandwidth for in-vehicle networks, in 2012, CAN in Automation (CiA) proposed updates to the CAN standard to allow for faster data rates and larger payload sizes. Now known as CAN with Flexible Data Rate (CAN FD), the updates were ratified into new versions of the ISO 11898 standard.

CAN FD makes three major changes to the CAN standard:

- The maximum data rate increased from 1 Mb/s to 5 Mb/s.
- The maximum payload size increased from 8 bytes of data to 64 bytes of data.
- Updated algorithms for cyclic redundancy checking.

CAN FD is still backwards-compatible with classical CAN *(Fig. 6)*. It has benefits of the increased data rate and payload size.

Taking advantage of these benefits requires new CAN FD controllers for handling the updated protocol and new CAN FD transceivers to handle the more stringent physical-layer requirements. Nevertheless, a 5X increase in bandwidth and much lower effective overhead with the larger payloads have led OEMs worldwide to quickly adopt CAN FD.

As with classical CAN, bus-fault protection, transient protection, EMC performance, and packaging options are still key differentiators.

CAN with Partial Networking

As worldwide vehicle-emissions requirements become more stringent, OEMs are exploring ways to reduce vehicle weight and energy consumption. To address lower system-wide vehicle power consumption, a third version of CAN was created, CAN with Partial Networking (CAN PN).

CAN PN added the ability to have a subset of nodes on a CAN network

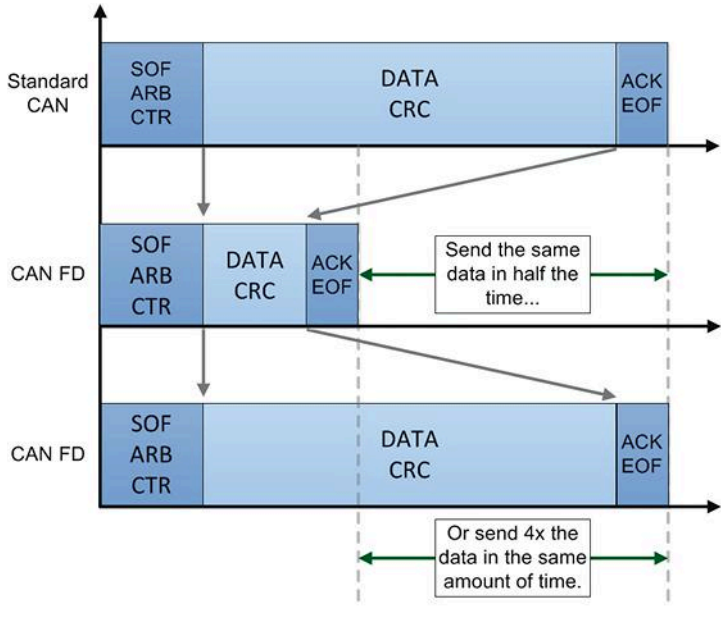

6. CAN FD, which is still backwards-compatible with classical CAN, has an increased data rate and larger payload size.

EXPERTS IN CARRIER GRADE DATA LINE PROTECTION SOLUTIONS

In-Stock and Available for Same-Day Shipping!

CPX SERIES - DATA SIGNAL SURGE PROTECTION

- Compact, 1RU modular chassis fits up to 16 ports
- Front access, modular for multi-protocol installations
- 10/100 Ethernet, GbE and T1/E1 fused and non-fused modules
- Coordinated Level 1 and Level II lightning protection
- UL 497A and UL 497B models
- Telcordia standard GR-1089-CORE
- Dual lug grounding studs

TRANSTECTOR
an INFINITE brand

transtector.com | +1 (208) 635-6400

Superior Protection | Patented Technology | Guaranteed Performance

CAN is set up as a broadcast topology: All nodes are required to check the validity of a transmitted frame and respond with an acknowledge bit. In practice, this means that if one node on a network sends a message, every other node on the network is required to wake up and error-check that message.

active while allowing other nodes on the network to remain in low-power mode. This is different from classical CAN and CAN FD, since all nodes on a network are either active or in low-power mode.

CAN is set up as a broadcast topology: All nodes are required to check the validity of a transmitted frame and respond with an acknowledge bit. In practice, this means that if one node on a network sends a message, every other node on the network is required to wake up and error-check that message.

Allowing nodes to remain in low-power mode until specific identifiers or payloads are broadcasted over the network enables OEMs to create a large variety of partial networks. Programming each node with only the messages important to it allows more nodes to remain in low-power mode, and thus reduces the vehicle's current consumption.

CAN PN has the same requirements as CAN and CAN FD, but also requires an accurate low-power oscillator and built-in digital logic for decoding broadcasted CAN messages.

ETHERNET

Unlike LIN and CAN, the Ethernet standard started as an industrial protocol and was well established worldwide before ever being considered for the automotive space. When 1- to 10-Mb/s data rates were no longer sufficient, automotive OEMs and suppliers looked to Ethernet. The Institute for Electrical and Electronics Engineers (IEEE) has a family of Ethernet standards called IEEE 802.3.

OEMs' desire to minimize the weight and emissions of any protocol used in automobiles led to the creation of two new standards specifically for automotive applications. Using new encoding schemes and transmitters and receivers capable of full-duplex operation reduced the wiring from two or four twisted pairs to a single twisted pair *(Fig. 7)*.

The first automotive standard released was the 802.3bw standard, which is also known as 100Base-T1. It has that

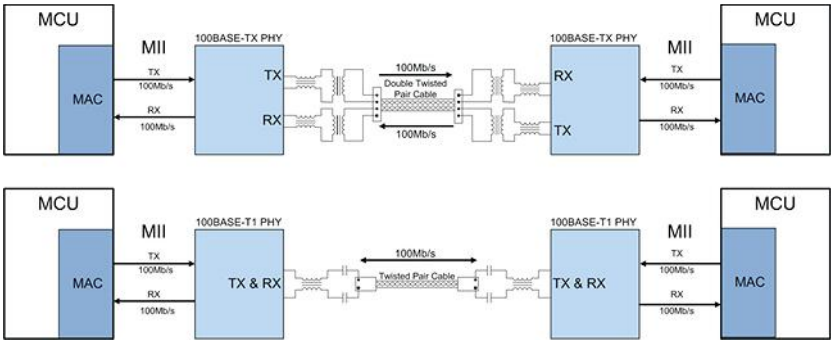

7. Ethernet normally uses four wires, but automotive Ethernet employs only two.

NEED IT NOW?

RELIABLE RF SURGE PROTECTION

CLICK | SHIP | PROTECT

Quality PolyPhaser RF Surge Protection Products Available for Online Purchase Today!

Our expanded eCommerce capabilities now offer:

- A broad range of quality RF surge protection products
- Same-day shipping
- Expanded product data sheets and technical information
- 24/7 online support

Learn more at polyphaser.com
Call us +1 (208) 635-6400

PolyPhaser
an INFINITE brand

name thanks to its 100-Mb/s data rate and ability to communicate full-duplex over a single unshielded twisted-pair cable (T1). The standard specifies that the transceivers must be able to support a minimum of 15 m; today, many transceivers support 30 m or more.

To further increase the bandwidth of the bus by another 10X, the IEEE 802.3 working group ratified the 802.3bp standard in 2016. This standard is also referred to as 1000Base-T1, since it allows for 1,000 Mb/s over a single unshielded twisted pair of cable.

One limitation of Ethernet is that it is a point-to-point standard; an Ethernet switch in the physical layer is necessary if one node needs to communicate with multiple other nodes. But because 100Base-T1 and 1000BASE-T1 are ac-coupled to the bus, the fault-protection voltage doesn't need to be as high as LIN and CAN *(Fig. 8)*. Furthermore, low emissions, high RF immunity, and the ability to communicate over longer distances are differentiators for Ethernet transceivers.

JOHN GRIFFITH is an experienced analog engineer with a demonstrated history of working in the semiconductor industry. He is skilled in semiconductor product support and definition, and PCB design and layout. John holds a Master's focused in Electrical Engineering from the Rochester Institute of Technology.

LIN	CAN	Ethernet
Low-bandwidth, single-wire interfaces	Mid-bandwidth interfaces, pair cabling	High bandwidth, pair cabling
Best for: Non-safety-critical information	**Best for:** Time-critical and safety related applications	**Best for:** Audio, video and high-bandwidth applications

8. LIN, CAN, and Ethernet provide complementary functionality.

HELPING YOU MINE

TRANSFER

ACCEPT

AND TRADE BITCOIN

MPD CAN BE FOUND IN CRYPTOCURRENCY APPLICATIONS AROUND THE WORLD

WWW.BATTERYHOLDERS.COM

IMS

Connecting Minds. Exchanging Ideas.

IMS2019
2-7 JUNE 2019

BOSTON 2019

THE HUB OF MICROWAVES

BOSTON
CONVENTION &
EXHIBITION
CENTER

BOSTON, MASSACHUSETTS

Register Today!

For more information visit,
www.ims-ieee.org

IEEE MTT-S

Get the entire pulse of the Microwave and RF industry in just a few days:

- The world's largest Microwave/RF component test & measurement, software and semiconductor exhibition featuring over 600 companies
- Over 75 technical sessions, workshops and panel sessions
- RF Bootcamp, a special course designed specifically for newcomers to the microwave world
- 5G Summit featuring experts from industry, government and academia
- Student forums such as design competitions, paper competition, PhD Initiative and Project Connect
- Special interest groups such as Women in Microwaves and Young Professionals
- Countless networking opportunities
- Special focus on startup companies with the first ever Startup pavilion, Startup panel session and Next Top Startup competition

IMS2019 and RFIC 2019 Keynote Speakers

RFIC Plenary Session Speaker
Sunday, 2 June 2019

Dr. Greg Henderson – *Senior Vice President Automotive, Communications and Aerospace & Defense, Analog Devices, Inc.*

"The Digital Future of RFICs"

RFIC Plenary Session Speaker
Sunday, 2 June 2019

Dr. Ir. Michael Peeters – *Program Director, Connectivity + Humanized Technology, imec*

"Do the Networks of the Future Care about the Materials of the Past?"

IMS Plenary Session Speaker
Monday, 3 June 2019

Dr. William Chappell – *Director of the Microsystems Technology Office (MTO), Defense Advanced Research Projects Agency (DARPA)*

"The Mind and Body of Intelligent RF"

MTT-S Awards Banquet Speaker
Wednesday, 5 June 2019

Dr. Ryan C.C. Chin – *CEO and Co-founder, Optimus Ride Inc.*

IMS Closing Session Speaker
Thursday, 6 June 2019

Dr. Dina Katabi – *Andrew & Erna Viterbi Professor of Electrical Engineering and Computer Science at MIT, Leader of NETMIT research group at CSAIL, Director of the MIT Center for Wireless Networks and Mobile Computing*

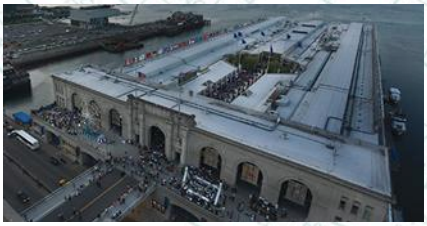

The IMS Welcome Reception will be held on Monday, 3 June 2019 at the Seaport World Trade Center Headhouse in Boston. The reception immediately follows the conclusion of the IMS Plenary Session.

For the latest on IMS and Microwave Week visit www.ims-ieee.org

TAMMY CARTER | Senior Product Manager, Curtiss-Wright Defense Solutions

GPU TRENDS: The Quest for Performance, Latency, and Flexibility in ISR Systems

Employing strategies such as GPUDirect, PCIe Device Lending, and implementing SISCI API can help system integrators optimize ISR solutions.

For military intelligence, surveillance, and reconnaissance (ISR) applications, such as radar, EO/IR (electro-optic/infrared), or wideband ELINT (electronic intelligence), the ongoing problem is how best to handle the expanding "firehose" of data, fed by an increasing number of wide-bandwidth platform sensors. To handle this massive inflow of data, and the complex algorithms required to process it, state-of-the-art computational engines and data-transport mechanisms are essential.

Deployed High Performance Embedded Computer (HPEC) systems designed to support these applications typically have a heterogeneous architecture of high-performance FPGAs, GPUs, and digital signal processors, or DSPs (today, often Intel Xeon-D based modules). GPUs provide a large number of floating-point cores tuned for complex mathematical algorithms, which makes them ideal for processing the complex algorithms used in ISR applications. In comparison, a single Intel Xeon-D processor can provide a peak throughput of ~600 MFLOPS, while NVIDIA's Pascal P5000 GPU sports 6.4 TFLOPS of peak performance.

TIGHTER INTEGRATION

Today, ISR system integrators have three main goals: minimize latency, maximize system bandwidth, and optimize configuration flexibility within their given SWaP constraints. To address these issues, leading COTS vendors of OpenVPX modules are seeking ways to provide closer integration between the compute elements.

In the beginning, sensor data preprocessed by the FPGA had to be copied to the CPU, which subsequently copied it to the GPU for further processing. Then, NVIDA introduced GPUDirect, which added the capability to move the data directly from an FPGA or network interface, such as Mellanox Infiniband, to a GPU. By eliminating extra copies, both latency and backplane utilization were decreased.

Such an approach works well until the amount of incoming data overwhelms the system, such that one batch of data hasn't completed processing before the next batch of data arrives. This can result either from the transport systems being overwhelmed (I/O bound) or the GPU not completing the calculations in the required time frame (compute bound).

When using GPUs, the limiting factor is often the I/O, and this is usually addressed by employing either a round-robin distribution of the incoming data, and/or pipelining the processing stages. Unfortunately, as sensor data continues to increase, it's become apparent that new techniques are required.

PCI EXPRESS

In OpenVPX systems, the standard interface between the FPGAs, GPUs, and CPUs is PCI Express (PCIe)—it offers the fastest path to and from the processor, and by definition, connects to other devices via the expansion plane. Offloading the Ethernet with the PCIe connection reduces latency and increases throughput.

Based on the original PCI parallel bus design, PCIe is controlled by a single "master" host called the root complex that scans the bus to find and enumerate all connected devices. When a PCIe switch is used to connect multiple devices to a root complex, it's called a transparent bridge (TB), and all devices operate in a single address space. With a TB, two root nodes (processors) can't be connected because there will be a memory address conflict.

When a PCIe switch port is configured as non-transparent bridge (NTB), a root node doesn't look to enumerate devices beyond that switch port.

So, when either of the two processors enumerates their NTB port, the port requests memory on that processor. The NTB port provides the common memory address translation to either side.

MULTICASTING

The next step to reducing latency is to multicast the incoming data to multiple GPUs. Until recently, multicasting GPUs in deployed ISR systems has been impractical because previous PCIe implementations were hampered by the traditional design limitations of one-to-one connections. The problem of how to multicast data across multiple GPUs using PCIe links increases the already complex task of programming the PCIe interface and setting up the bridges and root complexes.

Included in Curtiss-Wright's OpenHPEC Accelerator Suite, Dolphin Interconnect Solutions' PCI Express (PCIe) Fabric Communications Library provides the software needed to squeeze every ounce of performance from the PCIe interface. It also eases PCIe programming by abstracting the otherwise time-consuming code to simple APIs.

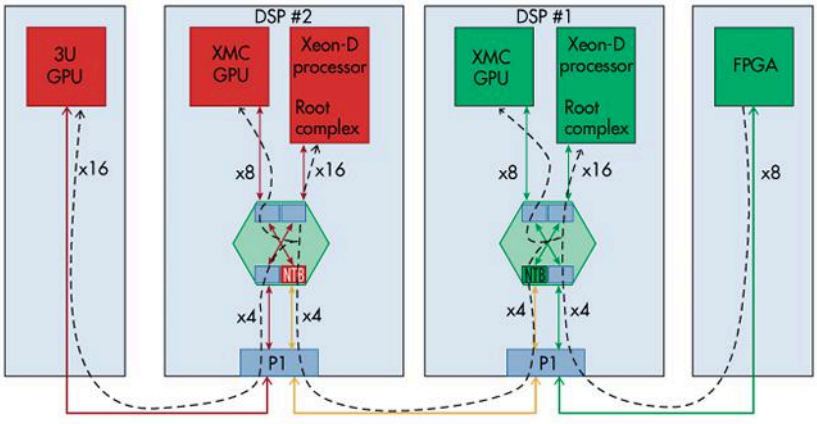

Here's an example of a 3U configuration.

The Software Infrastructure Shared-Memory Cluster Interconnect (SISCI) is a well-established API and is the fastest way to exchange data. It can be used for Programmed IO (PIO), where a pointer is used to directly access the remote memory with the lowest latency, or as remote direct memory access (RDMA) where the DMA controller of the PCIe-NTB copies data from remote to local memory with the highest bandwidth.

One of the latency-reducing features of SISCI is reflective memory/multicast. The PCIe switch will send out data on all connected ports simultaneously, meaning all nodes will receive data virtually simultaneously when connected to a single switch. When multiple bridges are used, each switch hop will add a tiny delay to the data delay.

In a typical backplane-based ISR system, the CPU is tied directly to a GPU that resides on its XMC mezzanine card, or in a separate slot in the chassis. For example, dual Xeon-Ds on a 6U OpenVPX board are connected via a PCIe switch. Likewise, a 6U GPU card might have two GPUs connected via a PCIe switch to each other and the backplane. Using this two-board 6U DSP and GPU system "slice," one Xeon-D can control both GPUs, or each of the Xeon-Ds can control one of the GPUs. With a three-board combination of one DSP board and two GPU boards, one Xeon-D controls the upstream GPU card while the second Xeon-D controls the downstream GPU board.

If the system is based on the smaller 3U OpenVPX form factor *(see figure)*, the latency issue becomes even more acute because there might only be one x4 PCIe link available to the CPU board. The solution becomes even trickier if the 3U system doesn't include a central switch.

CREATING FLEXIBILITY IN COMBINED SYSTEMS

ISR systems can run multiple modes, but now system integrators also want the ability to combine systems, such as radar and EW, into a single processing box. In such a configuration, the data flows and the computation requirements will vary greatly between each system's modes. Since the GPUs and FPGAs are attached with cables, or connected via the backplane, flexibility can be tricky. Some flexibility might be achieved by reconfiguring the PCIe bridge ports, but what if a CPU could borrow a GPU, FPGA, or even non-volatile memory attached on the PCIe bus?

With PCIe Device Lending by Dolphin, devices can be borrowed over the PCIe network without any software overhead. Device lending is one simple way to reconfigure systems and reallocate resources. The lending function makes the devices available on the network for temporary access. The borrowing function searches the available devices, and then the selected devices can be borrowed temporarily as required. When the work is complete, the devices will be released.

At the GTC 2019 conference, Dolphin demonstrated its proof-of-concept software library for creating GPU-oriented applications with GPU Direct using

capable GPUs and commodity NVMe disks. The memory-mapping capabilities of PCIe NTBs are used to set up efficient I/O data paths between GPUs and disks that are attached to different root complexes, allowing multiple GPUs to access a remote disk.

While GPUs have clear advantages for a variety of high-performance ISR applications, some designers are concerned with the lack of determinism exhibited by these devices. The good news here is that NVIDIA is now working to improve determinism on their GPUs. In CUDA, NVIDIA will expose APIs to describe tasks and the dependencies between tasks, thus providing more control over them. At the system level, NVIDIA is working on a timing-triggered scheduler.

WHAT ABOUT LEGACY CODE?

While new systems can readily benefit from these recent advances in GPU architectures, there is, of course, lots of legacy software still in use. Helping to protect the vast investments in this legacy code is OpenACC. It's a standard programming language that makes optimizing and porting older existing serial code to a multicore processor, like the Xeon-D or a GPU, a very simple process.

All that's required with OpenACC is a hashtag and a definition of the platform that the code will run on. In turn, the code will be able to run on the targeted board. This makes it possible to have optimization pragmas that enable code to be parallelized in small incremental stages (like one loop at a time), essentially allowing the system to take "baby steps" when optimizing/porting the existing code. While not as efficient as handwriting the code in CUDA, the results are impressive, and over the course of a few weeks can deliver great results for optimizing and porting years-old legacy code to a GPU.

When searching for the best possible signal-processing performance from a heterogeneous system, it's good to consider using GPUDirect to cut out the "middle man" CPU in order to reduce latency. The next step is to explore multicasting from the front end to multiple GPUs and CPUs, even if the data must pass through several bridges.

For maximum system flexibility, integrators should check out a library system of devices that can be borrowed using PCIe Device Lending. Keep software on time and under budget by using the SISCI API instead of writing all the PCIe code from scratch. For fast-turn porting of legacy, play with pragmas in OpenACC. These strategies can be helpful guides for any ISR system integrator seeking improved performance, latency, and flexibility. ed

NEERAJ DANTU | Applications and Systems Engineer, Octavo Systems LLC
MASOOD MURTUZA | Manager Package Engineering, Octavo Systems LLC
ERIK WELSH | Applications and Systems Manager, Octavo Systems LLC
GENE FRANTZ | Chief Technology Officer, Octavo Systems LLC
www.octavosystems.com

Path to Systems:
A SiP of Reliable Advantage—Systems Under Test

Fourth in a five-part series, this article presents how the methodologies of testing and characterization of system-in-package technology enables higher-quality, more-reliable products.

In parts 1, 2, and 3 of the series *(go to www.electronicdesign.com),* we talked about how the fundamental shift of focus toward systems makes system-in-package (SiP) an essential step forward in electronic design. Through SiP, a complete system can be packaged into a standard ball-grid-array (BGA) package the size of a nickel, such as the Octavo Systems' OSD335x in *Figure 1.*

We also discussed the advantages SiP brings to the design and manufacturing of semiconductors and how they can be used to tip the economy of scale toward low-volume opportunities. In this article, we will examine testing methodologies for SiP and how it simplifies testing and characterization of devices. This gives customers of SiP a more-reliable, higher-quality product to put into an end system.

There are several advantages to manufacturing a product or device using a system component (i.e., a SiP) rather than using discrete components. Here, we will focus on the advantages found in the device testing process and methodology. SiP components are uniquely suited to take advantage of the semiconductor manufacturing processes while simultaneously utilizing system-level testing normally reserved for the end product or device.

1. The OSD335x is an example of a complete system fitting into a tiny BGA package.

While Moore's Law has driven ever-increasing transistor densities and device complexity, semiconductor testing has kept pace by focusing on extremely high-volume manufacturing. This necessitates moving quickly through the manufacturing learning curve to reduce the cost of a product as quickly as possible. As a result, the semiconductor test methodology and equipment used during manufacturing have not only made semiconductor component testing extremely cost-effective, but have also resulted in higher quality and more reliable devices.[1]

With this as a backdrop, let's first explore some basic testing concepts of both semiconductor devices and systems. That follows with a look at a SiP as a device under test (DUT) and application of these concepts.

TESTING BASICS

Generally, the further down the manufacturing process an unusable device travels, the more money is lost.

It's therefore important to test as much as possible early in the manufacturing process. In addition, as devices become more differentiated, they become more expensive to test as the volume of the devices drops and the process becomes more manual.

For example, the component tests for a generic CPU are much more automated than system-level tests on specialized equipment using that CPU. As such, each subsequent round of testing should focus on aspects that weren't tested or couldn't be tested earlier in the process. This applies to both component-level testing of semiconductor devices as well as system-level testing required for products or devices.

Traditionally, semiconductor testing is focused on components rather than systems since most semiconductor devices are used as discrete components within a larger system. Component testing is primarily divided into two parts: wafer-level testing and packaged-device testing.

In a wafer probe,[2] the equipment will initiate different tests by probing each die in order to determine good die that will be packaged for further testing. Once the semiconductor die is packaged, the test equipment can perform the next set of tests, which includes external circuits, in order to finalize the component to be sold.

Due to the volume of products manufactured in the semiconductor world, the component-level testing uses automated test equipment (ATE) *(Fig. 2)*. Such ATEs use input vectors, patterns of 1s and 0s, or voltages designed to test for and uncover faults in the DUT. This is done by comparing the actual output vectors of the DUT to the expectations of the design specification, datasheet, and/or statistical characterization of known-good devices. ATEs are designed to use different test fixtures depending on the testing process and device to provide high test throughput and minimize the idle time of the ATE. For example, load boards enable automatic handling of packaged devices.

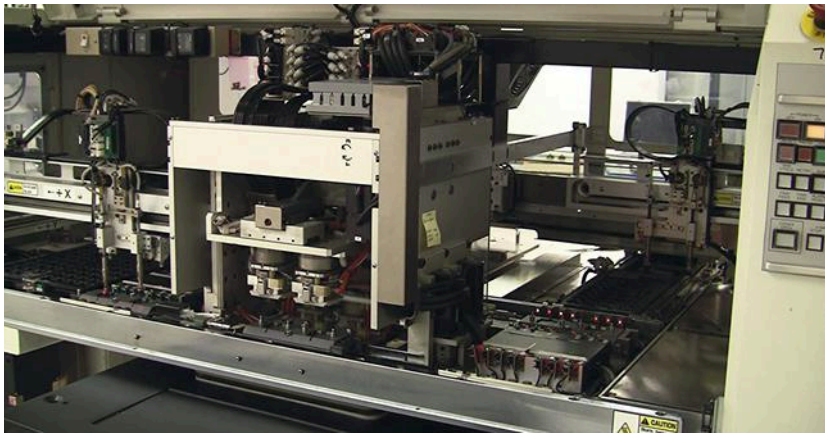

2. A typical automated test equipment (ATE) setup. *(Courtesy of Advanced Semiconductor Engineering Inc.)*

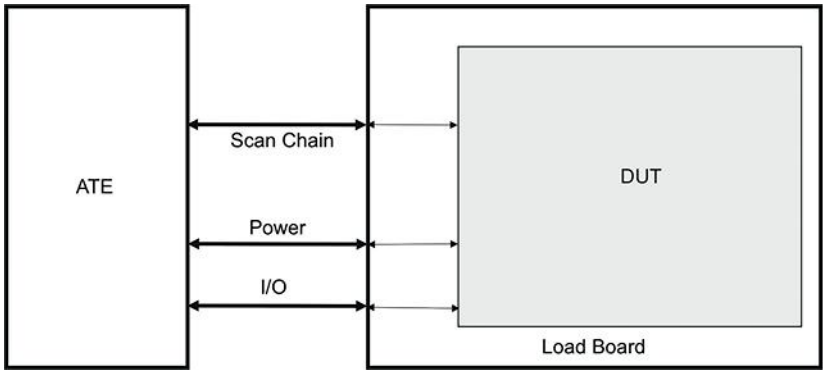

3. ATE test setup with load board and DUT (top); flow graph of how the ATE tests the DUT (bottom).

Figure 3 shows a simple block diagram of an ATE interacting with a packaged DUT via a load board and a flow graph of how it tests the DUT. While the ATE can have access to some or all of the pins of a given device, the primary interaction between the ATE and complex DUTs, such as a processor or system-on-chip (SoC), is generally via a JTAG (Joint Test Action Group) Test Access Port (TAP) and boundary-scan chain infrastructure *(Fig. 3, top)*. These elements inside a DUT can be accessed via a JTAG interface, which allows the testing of a device with different methodologies, such as automatic test pattern generation (ATPG) to detect stuck-at faults, built-in self-test (BIST) to detect faults in both memory and logic,3 or functional tests like IO loopback *(Fig. 3, bottom)*.

Traditionally, when building real-world products or devices, a system-level testing (SLT) methodology[4] is used. Generally, this involves loading production software or firmware on to the system and testing the system-level functionality of each component by covering all of the intended use cases. In many cases, this process is more manual.

An interesting example comes from the early part of Gene Frantz's career on the manufacturing production line in the Texas Instruments calculator division. Instead of an ATE, many times the production tester was a human. The production test on the DUT (in this case, a calculator) consisted of a set of test vectors (an equation) entered into the DUT and the output test vector (resulting answer to the equation) being read by the tester (human) with a determination of pass or fail. Depending on the volume of production, this type of manual testing process is still used today for system-level testing.

System-in-package devices, on the other hand, fall somewhere in between. They incorporate many tested die themselves, as well as other semiconductor components connected together. As such, they can be considered a system, requiring system-level testing. However, the good die may not have been as fully tested as individual packaged devices and, therefore, can't be considered a component, requiring component-level testing in SiP. This leads to the question: How should a SiP be tested?

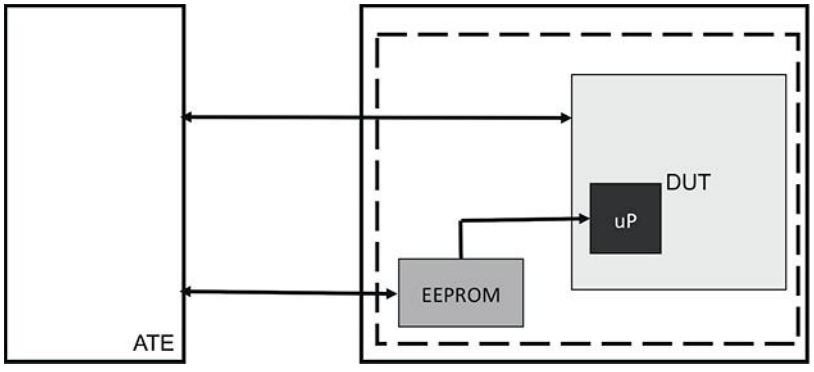

4. ATE setup to test a SiP (top); flow graph of the ATE testing the SiP (bottom).

HOW SHOULD THE SiP BE TESTED?

System-in-package devices can consist of multiple active and passive components. To test every component individually, all of the pins of each component need to be brought out to a pin/ball on the SiP. While this would enable full component-level packaged-device testing, it also drastically reduces the benefits of a SiP. Since manufacturing testing shouldn't reduce the design, cost, and size benefits of a SiP, a hybrid approach is required to verify and validate a SiP.

Component-level testing should be employed for all direct connections between die and the package pins/balls. This will ensure that there are no structural faults and provide a platform for characterization similar to what would be performed on standard packaged components.

On the other hand, system-level testing should be employed to validate and characterize all internal connections. This could include functional-level tests to validate connections between components, such as a processor and DDR memory. It could also involve stress tests to verify proper operation under heavy loads, e.g., confirming that a power-management IC (PMIC) can supply enough current to meet system needs. SiP devices may also use some pins/balls as test points (i.e., pins that should not be used functionally) to provide visibility on internal signals. For example, test points can be employed to monitor voltage rails only used within the SiP.

Using a hybrid testing approach for SiP devices provides several benefits. First, since components are already connected together within the SiP, those intermediate connections no longer have to be tested as rigorously as they would when testing the component individually. For example, if a PMIC provides a voltage rail that's used completely within the SiP, the PMIC output no longer must be fully characterized for all potential output loads since the output load is, by design, fixed.

Similarly, the input pins of all components attached to that voltage rail no longer need to be tested over a wide range of input voltages since that, too, is fixed. This allows for a reduction in the component-level tests required, since that connection can now be validated by a single system-level test. In addition, the form factor of a SiP makes it possible to utilize semiconductor ATEs, which can reduce the time needed for handling during many system-level tests and automate what might have otherwise been a manual process.

TESTING A SiP AS A SEMICONDUCTOR

Given the SiP form factor, it can be handled as if it were a semiconductor component. Socketed load boards (i.e., load boards that allow an ATE to connect to one or more DUTs) enable the ATE to send test vectors to the DUT and verify the resulting vectors of the test. In addition to validating the test vectors, the ATE can also perform open and short testing as well as voltage-level verification on the DUT. For example, the following diagrams describe how the OSD3358 SiP from Octavo Systems is tested as a semiconductor component rather than as a typical end system, such as a printed circuit board (PCB) or system-on-module (SoM).

As shown in *Figure 4 (top)*, the ATE interfaces with the SiP and non-volatile memory on the load board. The socketed load board provides the ATE with

Transient Immunity Testers from AVTECH

The Avtech AVRQ series of high-voltage, high-speed pulse generators is ideal for testing the common-mode transient immunity (CMTI) of next-generation optocouplers, isolated gate drivers, and other semiconductors.

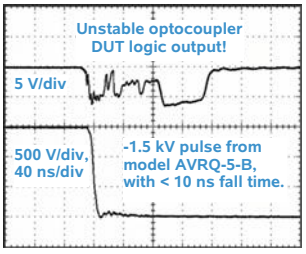

- Kilovolt amplitudes (±1, ±1.5, -2 kV)
- Transition times down to 10 ns, dV/dt rates up to 120 kV/us
- Switchable daughterboards to handle a variety of DUT package styles
- GPIB, RS-232 ports standard
- Ethernet / VXI optional

AVTECH

Nanosecond Electronics
Since 1975

Pricing, manuals, datasheets and test results at:
http://www.avtechpulse.com/semiconductor

30, 50 and 100 Volt Lab Pulsers

Avtech offers an extensive series of user-friendly 30, 50 & 100 Volt general-purpose lab pulsers. We can provide an alternative for the discontinued Agilent 8114A or HP214!

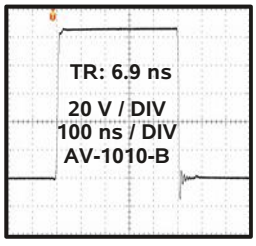

Model AV-1015-B: 50 Volts, 10 MHz
Model AV-1010-B: 100 Volts, 1 MHz, 25 ns to 10 ms, 10 ns rise time
Model AV-1011B1-B: 100 Volts, 2 ns rise time
Model AV-1011B3-B: 30 Volts, 0.5 ns rise time

Variable baseline and burst mode options and output currents to 8 Amps with accessory transformers.

AVTECH

Nanosecond Electronics
Since 1975

Pricing, manuals, datasheets and test results at:
http://www.avtechpulse.com/general

Nanosecond Laser Diode Drivers With Butterfly Diode Sockets

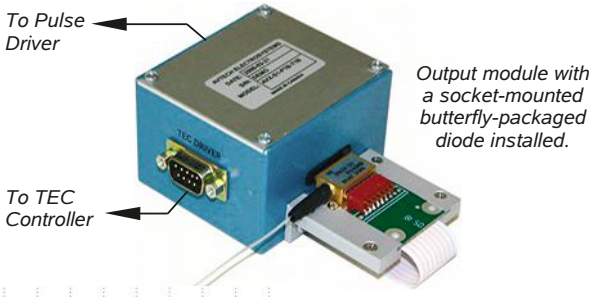

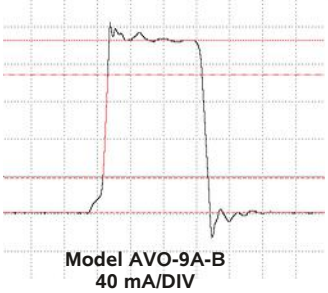

Model AVO-9A-B
40 mA/DIV
1 ns/DIV

Each of the 19 models in the Avtech AVO-9 series of pulsed laser diode drivers includes a replaceable output module with an ultra-high-speed socket suitable for use with sub-nanosecond rise time pulses. Models with maximum currents of 0.1A to 10A are available with pulse widths from 400 ps to 1 us. GPIB, RS-232, and Ethernet control available.

AVTECH

Nanosecond Electronics
Since 1975

Pricing, manuals, datasheets and test results at:
http://www.avtechpulse.com/laser

2 to 200 Amp Current Pulsers

Avtech has pioneered the design of user-friendly 2 to 200 Amp constant current pulsers. We offer over 35 models in five series for laser diode, air bag, squib, fuse and other applications. The basic specifications are as follows:

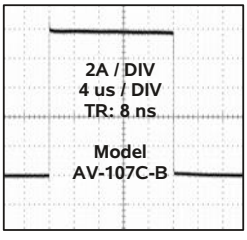

Series	I, V	PW	TR
AV-107	2 - 20 A, 60 V	0.2 - 200 us	10 - 30 ns
AV-106	5 - 100 A, 100 V	0.5 us - 1 ms	50 ns - 1 us
AV-108	12.5 - 200 A, 100V	2 us - 1 ms	5 - 15 us
AV-109	10 - 100 A, 5 V	10 us - 1 s	10 us
AV-156	2 - 30 A, 30 V	1 us -100 ms	0.2 - 50 us

Avtech has a long history of producing one-of-a-kind custom units.

AVTECH

Nanosecond Electronics
Since 1975

Pricing, manuals, datasheets and test results at:
http://www.avtechpulse.com/current

access to all of the SiP device's pins. *Figure 4 (bottom)* shows the flow graph of the test procedure. Open and short testing is first performed to eliminate DUTs (in this case, SiPs) with manufacturing faults that resulted in unconnected or shorted signals. The ATE then supplies power and verifies the voltage levels of the SiP. Finally, a software program contained in the non-volatile memory, which includes functional tests, is copied into the SiP during boot and is used to examine the rest of the SiP. This program interfaces to the ATE through the test vector input and output interface whereby the ATE can then send the commands to control which tests are executed by the SiP. *Figure 5* shows the OSD335x SiP undergoing test.

These tests might include verifying the hardware inside the SiP or validating various interfaces. After each test, the SiP can send an output via the vector interface that reflects the outcome of the test. Consequently, that part can be sorted and binned based on whether the device passed or failed the test. In general, the tests should be as short as possible and ordered based on which covers the most common failures, so that "bad" devices can be found as quickly as possible.

There are several advantages to this methodology. It's easy to update the non-volatile memory to add more tests or update existing tests. In addition, the quick feedback loop between the DUT and the ATE means that failures are identified faster and in a completely automated manner. Boot time is minimized since the DUT need only boot and load the tests one time. This can result in shorter test times, which reduces the cost of the device.

YIELD

While many testing benefits can be realized with SiP technology, one classical concern often raised is yield. The argument is, given that SiP can contain hundreds of components, each with their own failure rate, the final yield of a SiP is the product of the yields of individual components. For example, if you have five components within a SiP, each with a 95% yield, theoretically the resulting SiP yield is only approximately 77% ($0.95 * 0.95 * 0.95 * 0.95 * 0.95 \approx 0.77$). However, this isn't what's actually seen.

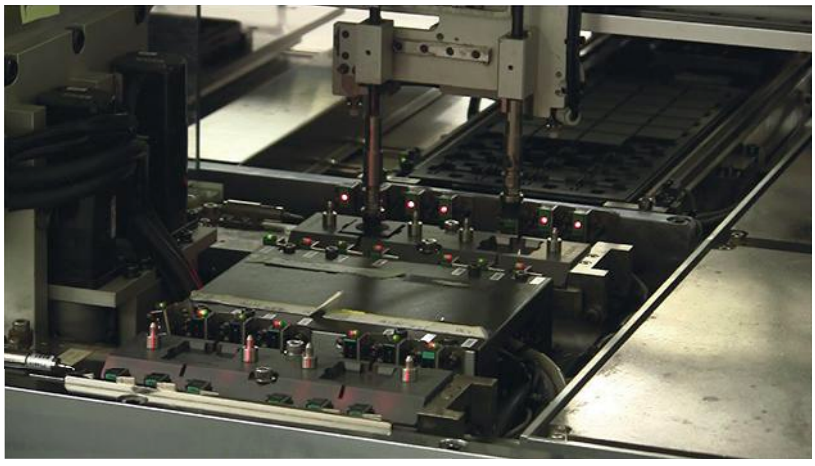

5. The OSD335x-SM SiP under test. *(Courtesy of Advanced Semiconductor Engineering Inc.)*

In reality, SiP yields are much higher due to a combination of factors. As discussed above, when discrete components are tested, all inputs and outputs must be fully tested. Furthermore, tighter test limits, or guard-banding, are applied to ensure that the component can be used in all extremes of the use conditions called out in the general-purpose specifications.

However, within a SIP, the use environment is known and controlled. For internal connections, variation in input voltage, current load, timing, and other variables are minimal. Therefore, devices that might have otherwise failed the normal component testing can be used within a SiP because it meets the requirements for the given SiP. Hence, SiP yields depend on the system use requirements, more like an end product, rather than just the individual component yields.

As more semiconductor companies look to make their devices "SiP ready," refinements in device testing will occur in wafer-level testing, enabling further improvements in SiP yields. Eventually, wafer-level tests will include any provision to make physical changes in chips, such as resistor trimming or fuse blowing, as well as additional test coverage to avoid complex testing in SiP form.

CONCLUSION

By using a hybrid of component- and system-level testing methodology, SiP offers benefits to the system manufacturer as well as the end customer of a more robust and cost-effective product. This is achievable if SiP manufacturers perform system tests while utilizing the testing methodologies and infrastructure of semiconductor testing.

We have found this method to yield extremely high rates of pass at next-level system tests. This unique method of testing, in addition to all of the other benefits of using a SiP device discussed in this series, provides strong arguments that a SiP solution like that developed by Octavo Systems may be the best solution for any embedded product design. ed

REFERENCES:

1. https://www.electronics-notes.com/articles/test-methods/automatic-automated-test-ate/basics-primer.php
2. https://anysilicon.com/ultimate-guide-wafer-sort/
3. https://www.electronicdesign.com/test-amp-measurement/what-s-difference-between-atpg-and-logic-bist
4. https://semiengineering.com/toward-system-level-test/

Design Note

42V Monolithic Synchronous Step-Down Regulators with 2.5µA Quiescent Current and Ultralow EMI

Dong Wang

Introduction

High efficiency, low EMI step-down regulators are found throughout automotive, industrial, medical and telecom environments, where they power a wide variety of applications from a broad array of input sources. Particularly in battery powered applications, a significant amount of time is spent in standby mode, requiring all electrical circuits to operate with a low quiescent current in order to preserve battery run times.

The LT8606/LT8607/LT8608 are a series of monolithic step-down regulators optimized for applications with a wide input voltage range, low EMI levels and small solution sizes. All share the same thermally enhanced 10-lead MSE package and 8-pin 2mm × 2mm DFN package, enabling them to fit into tight spaces. They differ in their output current capabilities, as shown in Table 1.

The low I_Q of the LT8606/LT8607/LT8608 is indispensable in battery-powered applications where idle current must be kept low. They feature a Burst Mode® option, which consumes only 2.5µA quiescent current from the input source even while regulating the output voltage, maintaining battery standby time for as long as possible. The 3V~42V wide input voltage range satisfies the demanding requirements of industrial and automotive applications, which are distinguished by their lack of stable, high quality voltage sources. The devices come in the 10-lead MSE package and also include spread spectrum operation to meet ultralow EMI emission requirements.

Table 1.

Part#	Current Level	Package	Operation Mode
LT8606	350mA	MSE-10	Burst Mode Operation Pulse-Skipping Mode Spread Spectrum Mode Sync Mode
		DFN-8	Burst Mode Operation Only
LT8607	750mA	MSE-10	Burst Mode Operation Pulse-Skipping Mode Spread Spectrum Mode Sync Mode
		DFN-8	Burst Mode Operation Only
LT8608	1.5A	MSE-10	Burst Mode Operation Pulse-Skipping Mode Spread Spectrum Mode Sync Mode
		DFN-8	Burst Mode Operation Only

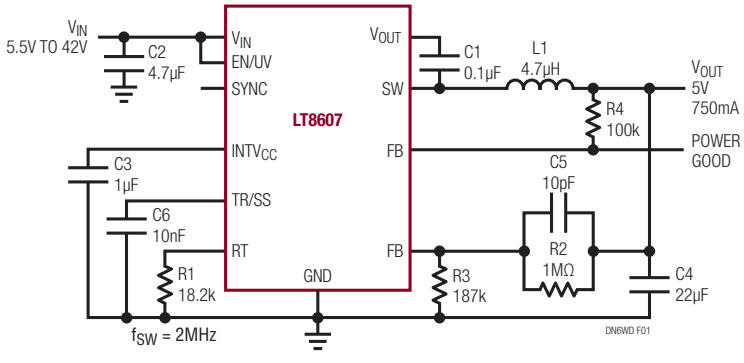

Figure 1. High Efficiency LT8607 12V to 5V Synchronous Step-Down Converter

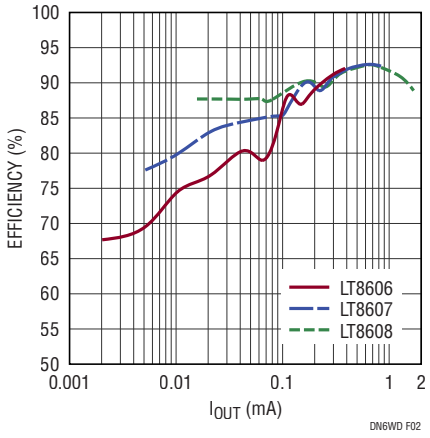

Figure 2. Efficiency vs. Load Current for LT8606/LT8607/LT8608 Based $12V_{IN}$ to $5V_{OUT}$ Step-Down Converter

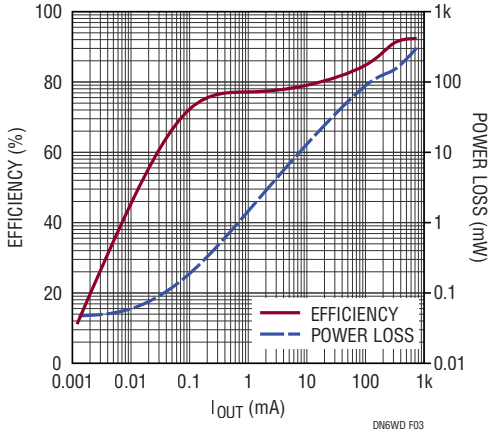

Figure 3. Efficiency and Power Loss vs Load Current for the Circuit in Figure 1

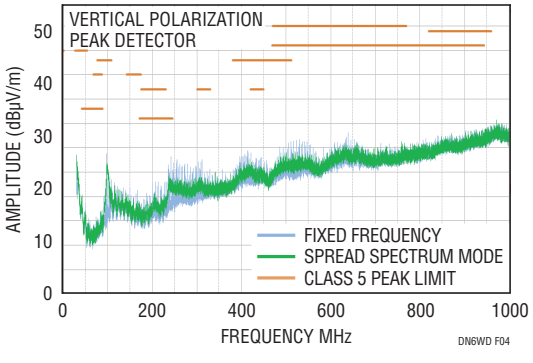

Figure 4. CISPR25 Radiated EMI Performance for the Circuit in Figure 1

Circuit Description and Functionality

Figure 1 shows a 5V output power supply based on the 10-lead LT8607 regulator. The input voltage extends up to 42V and the output is set to 5V at 750mA with 2MHz switching frequency. Only a few additional components are required for the complete solution, including inductor L1 and a few passive components. Figure 2 shows that this circuit can achieve 92.5% peak efficiency.

Burst Mode Operation Improves Light Load Efficiency

During light load operation and no-load standby mode, high efficiency and low idle current are very important for battery powered applications. The LT8606/LT8607/LT8608's 2.5µA quiescent current and Burst Mode operation option are perfect solution for these requirements. During light load and no-load conditions, an LT8606/LT8607/LT8608-based converter gradually reduces the switching frequency, which reduces switching power losses while maintaining low output voltage ripple. Figure 3 shows the light load efficiency of the solution shown in Figure 1.

High Switching Frequency with Ultralow EMI Emission

In addition to efficiency, EMI/EMC compliance is demanded in automotive, industrial, computational and telecom environments. A higher switching frequency allows a smaller solution size but often at the cost of increased EMI emission. The LT8606/LT8607/LT8608's integrated MOSFETs, built-in compensation circuit and 2.2MHz operation minimize solution size, but they also achieve excellent EMI performance, due to advanced process technology. Spread spectrum mode operation of the switching frequency can further reduce EMI emissions. Figure 4 shows the CISPR25 EMI test result of the solution shown in Figure 1.

Conclusion

The LT8606/LT8607/LT8608 are easy-to-use monolithic step-down regulators with integrated power MOSFETs and built-in compensation. They are optimized for applications with wide input voltage ranges and low EMI noise requirements. Their 2.5µA quiescent current and Burst Mode operation option makes them ideal solutions for battery powered step-down converters, significantly extending battery standby times. The 200kHz to 2.2MHz switching frequency range makes them suitable for most low power to micropower applications. Integrated MOSFETs, together with up to 2.2MHz switching frequency ability greatly minimize the total solution size. CISPR25 scanning results show their excellent radiated EMI performance, making them compliant with most stringent EMI standards.

Data Sheet Download
www.mouser.com/LT8607

DN577-48-03/19

Your Job just GOT a whole lot EASIER

NEW for 2019!

Customize your workspace with personalization features that will save you time:

- Customized Parts Library – build and organize parts into an easy-to-manage library
- Product Categories - tag only the one's that matter to you
- Part/Product Alerts – keep up-to-date on product availability
- Search "Authorized Only" Distributors – your reliable source for truly vetting authorized distributors

SourceESB has built-in engineering tools and functionality, designed just for you. If streamlining your workflow and increasing productivity are important, then you will appreciate the new features added to SourceESB in 2019.

SourceESB

SourceESB.com

The most trusted electronics parts search engine that serves the engineer from design to purchase.

WILLIAM WONG | Senior Technology Editor

The New "Shift-Left" Multi-Dimensional PCB Design-Verification Platform

Editor Bill Wong talks with Mentor's David Wiens about its unique "shift-left" PCB design verification platform for early prototyping and reduced re-spins.

A "shift-left" PCB design verification solution within the engineer's authoring environment is the industry's first of its kind, claims developer Mentor, a Siemens business. This new Xpedition platform helps the designer develop accurate virtual prototypes earlier in the design flow, to validate their designs with ease while reducing design re-spins.

To find out more about this latest tool, I talked with David Wiens, Xpedition product marketing manager for the Electronic Board Systems segment of Mentor.

What was the basis for this new technology development since Mentor currently has a broad portfolio of PCB design solutions?

This new PCB design platform is focused on multi-dimensional verification—a critical area that's lacking in sufficient technologies. The systems design industry continues to increase with global pressures to deliver advanced products in shrinking time-to-market windows. In our conversations with customers, PCB design managers are challenged by the lack of robust verification tools, or available tools that are too difficult to use. Increasing performance requirements plus the pressure to improve product quality are driving these teams to look at alternative approaches to their current verification methods with physical prototypes.

Best-practice design processes validate a "digital twin," or model, of the design, early and often to minimize re-spins and shorten design cycles. Our solution is a "shift-left" PCB design flow that allows design engineers and layout designers to validate within their native environment, removing the bottleneck of specialist reviews later in the design process.

We met with industry analyst Chad Jackson of Lifecycle Insights, who recently published a study on electronics simulation-driven design and its correlation to re-spins, costs, and time-to-market. His findings support our new technology, enabling PCB design teams to simulate earlier and often—minimizing critical issues in the early stages of product development.

What are some of these findings from the Lifecycle Insights study?

Lifecycle Insights reports that 58% of all new product design projects incur added costs and time delays. And 75% is worst case—where projects are cancelled for various reasons, including project delays or reallocated resources.

The study includes the high cost of design failure by looking at re-spins and re-designs due to errors found in physical prototypes and production boards. The average design re-spin averaged 8.5 days to complete at a cost of $44,000 (USD) per re-spin. The study reports an average of 2.9 re-spins per project. So this amounts to $127,000 or more per design—that's substantial. Today's high-end designs can approach six figures for a re-spin, and with steadily increasing design complexity, I foresee the number of re-spins and associated costs increasing year over year.

That's interesting. Tell me more about what Mentor is doing to address this issue.

It may be worth discussing the conventional PCB design processes first. Traditional processes leave validation for the prototype lab, so it's difficult to know the root cause of what caused the problem and options to remedy the problem. There are analysis tools in the market, and these are used right before prototype. But this approach relies on specialists trained to use these complicated tools, which can become a "bottleneck," particularly with time-to-market deadlines.

With the Xpedition design verification platform, validation is implemented as early as possible in the design process. This "shift-left" approach is applied right after an error could be made in the system definition, schematic, or layout stage. Design validation is handled by the author at that stage, such as the design engineer during the schematic stage, or the layout designer at the layout stage. We tightly integrated our validation tools with the authoring environment, making them easy to use.

FREE WEBINAR

March 21, 2019 | 2:00pm EST

Registration URL: https://www.electronicdesign.com/webcasts/simulation-focused-ultrasounds-medtech

SIMULATION OF FOCUSED ULTRASOUNDS FOR MEDTECH

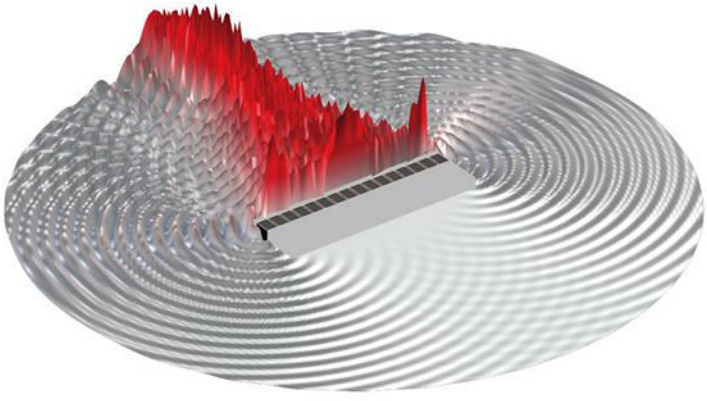

Focused ultrasound signal generated by 16 piezoelectric transducer elements.

Ultrasounds are widely used in several industries and use cases, like the nondestructive testing (NDT) of metallic parts or for medical imaging, one famous example being echography.

The advantage of ultrasounds is being able to reach a volume inside matter without affecting the surface and what is between the surface and the zone of interest.

In the medical technology (medtech) industry, high-intensity focused ultrasound (HIFU) surgical tools are designed based on focused ultrasounds and produce a localized elevation of temperature and necrosis of biological tissues, for instance, for the treatment of some types of cancer.

Simulating the acoustics and heat transfer phenomena involved in such tools allows engineers and researchers to select the combination of parameters that will deliver the right amount of energy in the targeted zone and limit the damage to the surrounding healthy tissues. There are numerous parameters, including the size of the transducer that transmits the ultrasounds, the frequency of the signal, and the duration of the treatment.

In this presentation, Thomas Clavet from EMC3 Consulting will discuss how HIFUs can be produced and key points about how to model this multiphysics problem.

SPEAKERS:

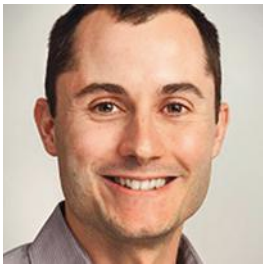

Thomas Clavet, *Principal Engineer, EMC3 Consulting*

Thomas Clavet founded EMC3 Consulting in February, 2014, to support companies (including SMEs, major industrial groups, and research laboratories) in their use of digital simulation for acoustics, thermal, mechanical, and flow (CFD) applications. Prior to becoming a COMSOL Certified Consultant, Thomas was a mechanical engineer in the nuclear industry and an applications engineer at COMSOL in the UK and Ireland, where he was able to meet many COMSOL Multiphysics® users from a variety of backgrounds and industrial sectors. Thomas is an Arts et Métiers ParisTech engineer and holds a master's degree in mechanics and numerical methods from the Royal Polytechnic Institute (KTH) in Stockholm, Sweden.

James Gaffney, *Applications Engineer, COMSOL*

James Gaffney works at COMSOL as an applications engineer for acoustics. He studied acoustical engineering at the University of Southampton, where he also earned his doctorate degree. His research involved predicting the fuselage installation effects from engine fan tones with analytical methods.

Hosted with
ElectronicDesign

COMSOL

Validating during the design process is often referred to as virtual prototyping, leveraging a "digital twin" model of the design. And by validating first-order issues during design authoring, the specialists can focus on more complex, complete multi-physics analysis for the challenging parts of the design versus the entire design. The net result is a shortened design cycle with reduced respins that lead to higher-quality products, validated by the respective design authors. Mentor is the first PCB design software company to provide this "shift-left" verification capability.

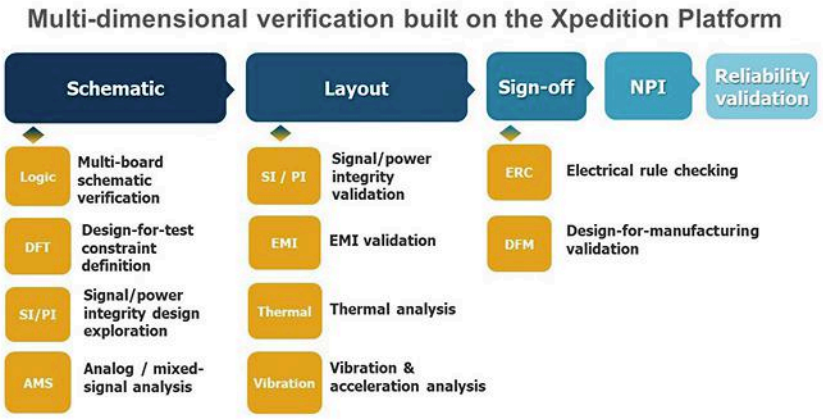

The Mentor Xpedition PCB design platform provides the industry's first "shift-left" upfront verification capability within the engineer's authoring environment to minimize design respins and improve overall product quality.

Can you elaborate on specific capabilities for upfront design analysis? For example, what can the design engineer do to validate the schematic design?

Based on our broad portfolio of analysis and verification technologies, the design engineer can identify the most common schematic integrity errors before layout. Traditionally, this is a manual, peer review process to catch errors, which is not only time-consuming but prone to miss mistakes.

We can conduct more than 150 voltage-aware rule checks automatically, on every net in the system design including across multiple boards—essentially hundreds of thousands of checks per design. Automated checking can uncover problems with improperly connected components, missing power or maximum power exceeded on a component, improperly placed diodes, pin voltage mismatching, missing connections, nets missing the receiver, wrong board-to-board connections, as examples. Even the smallest schematic error can lead to a major system failure.

Are there any customers who are realizing the benefits of this capability that you can talk about?

A major mil-aero customer uses Xpedition, and they were able to identify 13 critical errors that reduced their debug time from two months to two weeks, not including their time to actually re-design their system. By replacing the manual process with automated schematic analysis, this customer was able to review 100% of the nets for assured coverage.

Our tool provides color-coded results to highlight critical errors from possible defects—including basic warnings to improve the design. If an error is incorrectly identified, the user can hide it or omit it from future verification checks. Another benefit with this schematic analysis tool is seamless integration. The bill-of-materials and netlist data are automatically extracted and results reviews are cross-probed into the schematic. With a library comprising over six million parts, coupled with fast modeling services, tools adoption and ease-of-use are clear advantages. We've found among our customers that this has saved them an average of 18 days per project.

What are some of the other new capabilities for "shift-left" verification?

Integrated testability analysis. Test point assignment and test coverage are usually considered late in the PCB design process. Typically, the first or second revision of the product will be completed before testability issues are considered, but test strategy optimization is difficult at this stage since all design constraints were considered during layout. Therefore, testability considerations should be made during schematic capture, driving test point requirements during layout.

During layout, the test strategy is validated, so this capability can identify errors early and reduce errors associated with insufficient test coverage. By defining appropriate test coverage, resulting in not too many and not to too few test points, this solution makes design handoff to manufacturing much more seamless, reducing costs during functional testing. This "manufacturing-aware" feature delivers an efficient flow, from design conception to manufacturability.

You mentioned "multi-dimensional verification" earlier during our conversation. What do you mean by this?

The new Xpedition platform includes improved integrations of our existing tools for "shift-left" verification. These include automated component modeling for vibration and acceleration simulation, integrated electrical rule checking, DC drop for multi-board and rigid-flex designs, and concurrent design-for-manufacturing analysis during layout. Mentor provides a comprehensive set of verification technologies within the Xpedition flow. We believe this new platform will result in higher-quality products while saving time and costs by mitigating design re-spins.

High Voltage and High Current How You Need It

High Voltage
Diodes, Capacitors
& Rectifiers

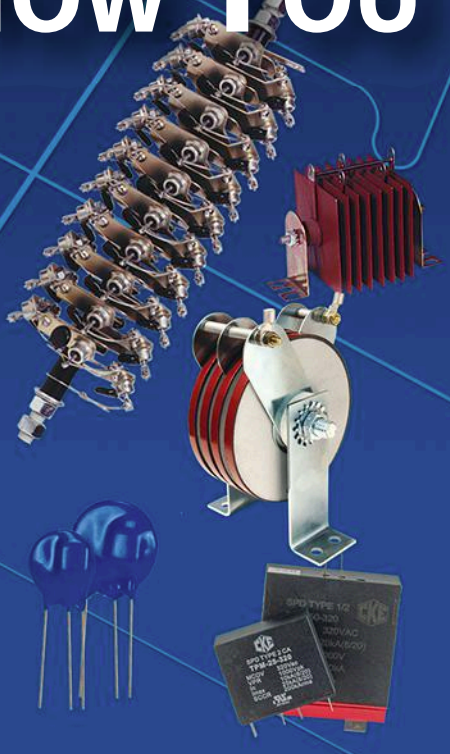

Suppressors
High Voltage Assemblies
High Power Assemblies

High Voltage
Power Supplies, Multipliers
& Test Equipment

HVPSI

Contact DTI today to discuss your high voltage or high power design.

info@deantechnology.com
www.deantechnology.com

DEAN TECHNOLOGY

STEPHEN McCANN | Chair IEEE 802.11aq Task Group, IEEE Standards Association
https://standards.ieee.org/

Simplified Wi-Fi Service Discovery

Connecting to a Wi-Fi network that provides the right service for your smartphone or IoT device has always been an issue. IEEE 802.11aq solves this problem.

Finding a Wi-Fi network to provide the correct service that your smartphone or IoT device is looking for has always been a problem. When you bring your new device into the office or home, it seems to always involve connecting to a nearby Wi-Fi network to try out the available features. As IoT devices advance and proliferate, this may become an issue for first-time adopters of IoT devices, especially when they don't have user interfaces.

To assist with this type of issue, IEEE recently published IEEE 802.11aq—"Pre-Association Service Discovery"—the latest standard in the IEEE 802.11 family. This article will explain a little more as to what it provides.

Wi-Fi is the premier local-area-network (LAN) technology for high-speed internet access for laptops, smartphones, and tablets, built using IEEE 802.11 standards. Most of us are familiar with the new faster versions of IEEE 802.11 that use the 2.4-, 5-, and 60-GHz industrial, scientific, and medical (ISM) bands, such as IEEE 802.11ac and the current IEEE 802.11ax under development. In addition, there are also updates, such as IEEE 802.11aq providing simple Wi-Fi service discovery, together with some MAC Address privacy enhancements.

PRE-ASSOCIATION DISCOVERY (PAD)

Let's start by looking at PAD, which is the primary feature of IEEE 802.11aq. It enables a protocol designed to discover services on a Wi-Fi network by an end-user device. PAD is an interworking function provided by a Wi-Fi network to allow a device, prior to association, to discover information concerning services that might be available to that device, once it is associated with the Wi-Fi network.

PAD provides methods for the device to gather information to aid in the decision to select a Wi-Fi network with which to associate. It's important to real-

ize that PAD doesn't provide an actual connection to a service, before association; just the information about the service's availability.

IoT SERVICE DISCOVERY

Your new IoT device requires a specific network service to control its power on and off feature. However, in your office environment, several Wi-Fi networks are available. Although the IoT device is smart, there's no user interface to assist an initial connection. How does that IoT device find a Wi-Fi network that has the specific service needed by the IoT device? These devices don't want to connect (associate) to every Wi-Fi network in range and then try to discover the service it requires. Indeed, without a user interface this may not be possible. The IoT device just wants to connect and connect in a manner that's seamless.

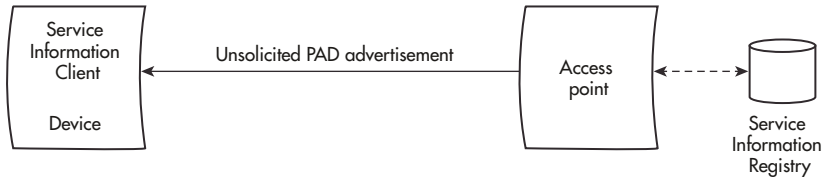

1. Pre-association Discovery (PAD) advertisement frames are transmitted over Wi-Fi in an unsolicited manner. The Service Information Registry contains information about supported network services.

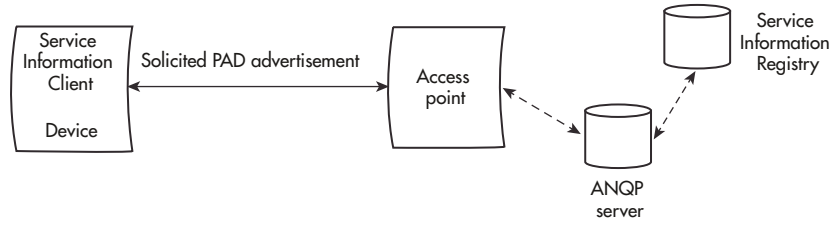

2. PAD supports a query/response protocol that allows a Service Information Client (operating at the device) to determine detailed information about available services in the SIR via an Access Network Query Protocol (ANQP) server.

BROADCAST SERVICE INFORMATION

The first PAD method allows the device to receive broadcast advertisements from an access point (AP) *(Fig. 1)*. The unsolicited PAD advertisements contain information about the services supported by the wireless local-area network (WLAN) to which the AP is connected. The device needn't send any uplink request for these advertisements.

When the device receives the advertisements, it passes information about supported Wi-Fi network services to the Service Information Client (SIC) that interfaces with higher layer applications. An application on the device that requires a service will then look in the SIC to determine if any Wi-Fi networks support the service that it requires.

REQUEST/RESPONSE SERVICE DISCOVERY

The next PAD method allows devices to perform a request of "what services can be reached" in a Wi-Fi network, before connecting (associating) *(Fig. 2)*. This enables the devices, in a pre-associated state, to request solicited PAD advertisements about a specific service. The request can also contain a more detailed query to discover additional information about that specific service. This method allows for more functionality than the unsolicited PAD advertisement described earlier, but also makes it possible for the device to operate a more directed search.

A device creates a service request message in the SIC, containing some information about the service such as a keyword or a well-known service hash, which is then sent from the device to the access point. From here, the request is forwarded to a Service Information Registry, via an Access Network Query Protocol (ANQP) server.

The SIR creates a service response message, containing more detailed information about any available services that match those of the request. If no service is available, this is also detailed in the service response message. The SIR relays the service response message back to the AP, which is then transmitted back to the device. The request and response are allowed to transit from the AP to the SIR even though the device isn't associated to the Wi-Fi network.

The service request and response messages can be considered as containers that enable requests based on upper layer service discovery protocols (e.g., UPnP, mDNS) to be transported between the device and AP.

PROVISIONING AND CONFIGURATION OF SERVICES

By introducing the SIR, PAD allows the provisioning and configuration of services within the Wi-Fi network. An AP can use the SIR as a proxy server that stores the service capability of the Wi-Fi network to which the access point is connected. Available services in the Wi-Fi network can be registered within the SIR, so that devices using PAD can then query for them at some later point. Therefore, the access point can quickly respond to "Service Discovery" queries without flooding the WLAN with network service discovery requests.

For example, a smart speaker is already connected to the Wi-Fi net-

work. It has a guest mode that allows any device to play music. A new device is brought in range and wishes to play music. Using PAD, the new device correctly discovers the Wi-Fi network that supports the smart speaker service for devices. Using the PAD methods, outlined above, it can be done in a seamless way without any user intervention.

MAC ADDRESS RANDOMIZATION

By enabling and encouraging devices to transmit service discovery requests, before those devices have associated to Wi-Fi network, there's an increased possibility that the MAC addresses of those devices can be tracked. Users with IoT devices walking through a shopping mall or a sports stadium can have the MAC address of that IoT device logged, so that the same device can be recognized in the future. This type of information can then be sold to marketers and other third parties.

To address this privacy concern, IEEE 802.11aq also focuses on pre-association privacy. If a country was to adopt a requirement to provide privacy of device addresses, IEEE 802.11aq provides a standardized way to achieve this.

When an IEEE 802.11aq device connects to, or attempts to discover services, on a Wi-Fi network pre-association, it defines the addressing of its link layer for the particular connection. One such addressing scheme is to periodically and randomly change its MAC addresses, reset counters, and seeds prior to association. The device sets its MAC address in accordance with the policy of the network it's configured to connect.

BENEFITS OF IEEE 802.11AQ

- Enhance installation by allowing IoT devices to discover Wi-Fi networks with the services that are required.
- Increase the value of enterprise and public Wi-Fi networks to IoT home users, installers, and operators through a better experience.
- Improve the IoT ecosystem by enabling the operator to offer additional services.
- Provide enhancements to device privacy through the use of MAC address randomization. ed

STEPHEN McCANN has 29 years' experience in the telecoms industry, starting with DECT and the early days of GSM, then moving to Hiperlan2 and IEEE 802.11 WLAN technologies. He has been secretary of the IEEE 802.11 WLAN standards organization since 2006 and chaired the IEEE 802.11aq "Pre-association Discovery" project and its corresponding program within the Wi-Fi Alliance "Pre-association Infrastructure Discovery." Stephen has also chaired the GSMA Terminal Steering Group Wi-Fi group since May 2013.

MAHENDRA TAILOR | Technology Leader at Laird Connectivity

SIMPLIFYING BLE Mesh Integration

Bluetooth Low Energy (BLE) Mesh technology allows smart devices to maintain a cooperative wireless network that can be managed remotely, leading to the development of smart homes, building automation, and sensor networks.

Bluetooth Low Energy (BLE) is a wireless technology available on Android-, iOS-, Windows-, and Linux-based devices. It's designed to provide connectivity between a master controller, typically a smartphone or tablet, and smart devices.

Smart devices are everywhere and include everyday Internet of Things (IoT) objects such as lighting, heating systems, home appliances, access alarms, sensors, location tags, and wearables. Using BLE Mesh technology, these smart devices can connect to a wireless network and be managed remotely, enabling the development of smart homes, building automation, sensor networks, and smart industries and cities.

OVERVIEW

A BLE mesh is a network of as many as 32,766 smart devices, although most networks are much smaller. Messages sent and received within this network initiate predefined events that occur in the participating devices. Any BLE device running BT4.0 or higher has the potential to support BLE meshes, provided the proper firmware is installed.

The BLE mesh uses a managed flood technique for transmitting messages to connected devices, which are called nodes *(Fig. 1)*. It sends a message, called a BLE broadcast or advertisement, to those nodes that are within range. The receiving nodes rebroadcast, or relay, this message to other nodes within their range. Once the proper node receives the message, the predefined event is initiated.

Messages aren't rebroadcast unconditionally. The nodes rebroadcast until the message time-to-live (TTL) has expired, where the TTL defines the number of

relays, or hops, that a message can make. The TTL is decremented by one as it reaches each relay node. Relay nodes will only forward messages that have a TTL greater than one, to avoid an endless cascade of messages. Also, each node in the mesh maintains a message cache of its most recent messages, and incoming messages that match one already received are discarded without rebroadcasting.

BLE mesh networks using relay nodes are far more reliable than other types of networks that use nodes as routers. Since relay nodes receive broadcasts from any node within range, this redundancy keeps the network up even if an intermediate node fails or is removed from the network, while router nodes can be a single point of failure. Relay nodes also provide better scalability and improved performance.

Periodically, a heartbeat message can be sent from each node to indicate that it's still active. Using the heartbeat message to see how far each node is from other nodes, in terms of the number of hops needed to reach it, allows the TTL to be optimized and set to a value no higher than necessary.

TERMINOLOGY

While all nodes in a BLE mesh can transmit and receive messages, some nodes have additional specific purposes. A "relay" node can rebroadcast a message that it receives. A "low-power" node spends most of its time in a low-power state with its radio turned off to save battery life. A "friend" node works in conjunction with a low-power node by storing and then forwarding messages when it's polled by that low-power node.

A "proxy" node is a device that uses legacy Bluetooth connectivity to interact with other devices in the BLE mesh; in turn, it can forward messages to other nodes via other connections or adverts, if the relay feature is also enabled. Smartphones are expected to be the type of nodes that will use proxy nodes to convey messages simply because a phone can't use relay mode. A relay mode requires the radio to be switched on 100% of the time, which isn't going to be allowed by the phone manufacturer.

The Mesh spec allows a node to be in multiple modes, like relay, proxy, and friend. It's unlikely a node using low-power mode will use any other mode.

Messages are sent and received using a publish/subscribe paradigm *(Fig. 2)*. Publishing involves sending messages from one node to a group of one or more other nodes. Subscribing is configuring a node so that it can receive messages from 0 or more nodes. A node has a subscription list that contains the publish addresses to which it has subscribed, so that it can select or filter out messages that it receives.

Each message consists of an opcode and a payload, or content data. The opcode dictates the behavior at the receiving end. Outgoing messages on a BLE mesh are encrypted with both a network key and an application key. The network key provides security for all communication within the BLE mesh network. The application key delivers separation and authentication to ensure the correct model in the intended device is the one that's activated—for example, a message sent to turn on a light bulb should not be able to unlock a door. By the time the application receives its message, it has been completely decrypted and is given as plaintext.

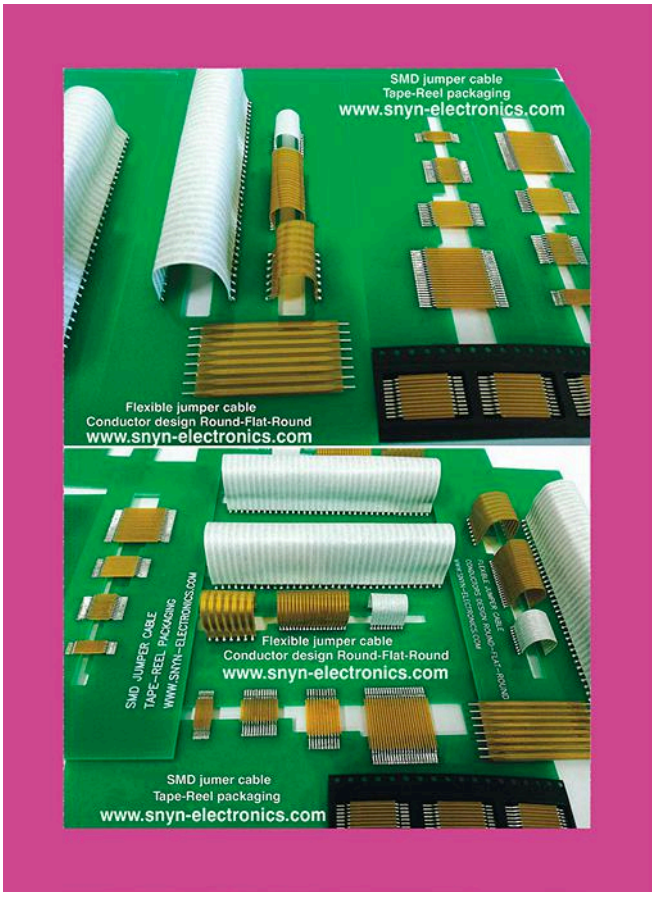

PROVISIONING A NODE

When a BLE mesh device is powered up for the first time, it's not "provisioned." That is, it doesn't have a node address or any other configuration information. To become part of the mesh, it must be given this information. Once the device is provisioned, typically by a smartphone, it's called a node, and has a unicast address, publish information, a subscription list, and network and applications keys, among other things. Provisioning updates the records in the device's persistent configuration database. Each node will have one or more elements, which in turn have one or more models, that expedite a particular behavior for that node.

The provisioner is the only entity in the mesh that's aware of all members in the network. It's also the only device that knows how to program the publication address and the subscription lists of every node. No individual node is aware of

BLE mesh is a managed flood network

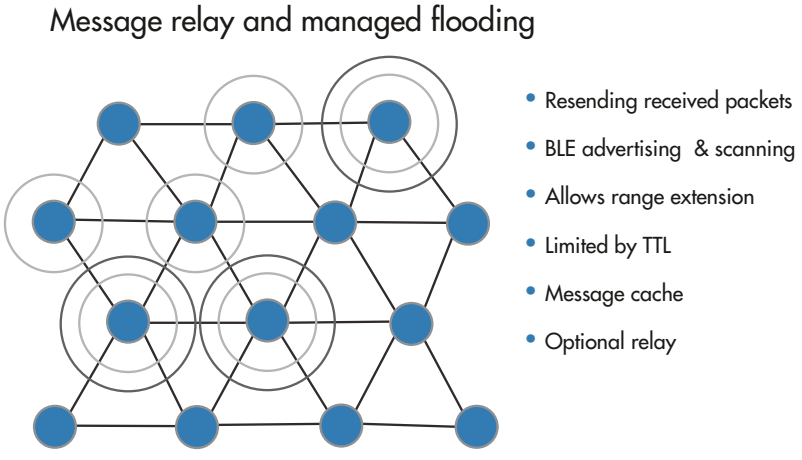

1. The BLE mesh takes a managed flood approach to distributing a system through the network.

the entire mesh. The provisioner isn't needed for the network to function, but it is needed to add or remove nodes from the network. In this way, there's no central point of failure that can bring down the entire network. However, if the provisioner device is damaged or lost, then no one else knows all of the nodes in the network; they must all be re-provisioned with another device.

As an illustration, consider an office with 50 lights and 20 switches. Only the provisioner is aware of the node addresses of all 70 devices. The provisioner can create group addresses for each area of the office, and then set the publication addresses in the switches and the subscription addresses in the lights based on the group addresses. When a light switch's state is changed from off to on, for example, it publishes a group address and the application-subscribed lights turn on.

IMPLEMENTING A BLE MESH

Implementing a BLE mesh may sound overwhelming, but the process can be greatly simplified. Manufacturers like Laird Connectivity offer BLE modules that jumpstart the mesh-development process. BLE modules have an embedded event-driven scripting language, smartBASIC, built into their firmware that's used by product developers to simplify BLE module integration.

By using only a handful of built-in functions and events, an application can be written for implementing a BLE mesh using a Laird BLE module. The first time the BLE mesh is started, the underlying stack will see that the device isn't provisioned, which starts the beaconing process. This alerts the user, who will provision the device, making it a node. Otherwise, the start function lets the node participate in the mesh network and listen for incoming messages, which it will either act on or relay onwards.

When an incoming message is fully decrypted by the underlying stack, it triggers a predefined event and the application script will have a user-written handler that defines the behavior for that event. Such a script examines the opcode in the message and determines what action should be taken, such as switching on a light, starting a washing machine, or setting the house alarm or even send the information to the cloud.

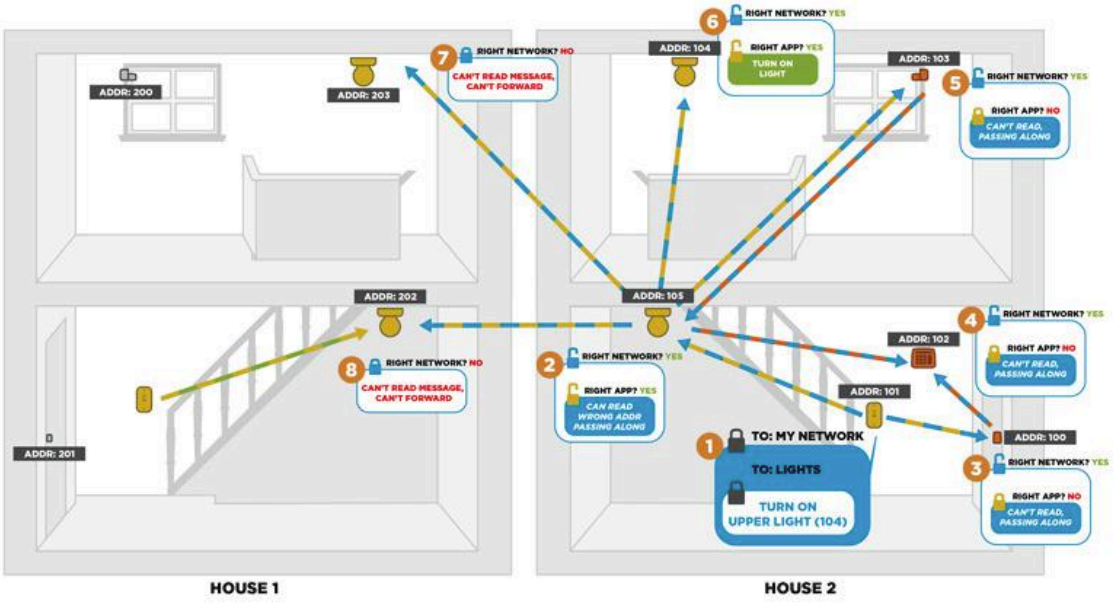

2. This is an example of the BLE mesh in action with neighboring networks.

Programmable BLE modules, like the BL652 and BL654 from Laird Connectivity, provide secure wireless engines for IoT networks. They have a wide range of configurable interfaces and an industrial temperature rating, and there's a fully featured development kit to start BLE development. The power-management-efficient BL652 and BL654, which measure 10 x 14 mm and 10 x 15 mm, respectively, provide hostless operation for automated use cases.

CONCLUSION

Implementing a BLE mesh network allows a user to connect IoT devices in a smart wireless network. The network is flexible, scalable, robust, and reliable. Using commercially available BLE modules greatly simplifies the implementation process, enabling these smart networks to be up and running more quickly. ed

MAHENDRA TAILOR is a Technology Leader for Laird Connectivity, where he has worked for 21 years. For nearly 15 years, he has been involved mostly in Bluetooth, which has resulted in many patents. Mahendra has worked in embedded and automation systems for over 35 years, after getting a degree in Control and Instrumentation Engineering from The City University in London. His goal is to make technology easy for others to embed into end products.

ROBERT NITSCHE | Director of Application Engineering, Plastic Logic
www.plasticlogic.com

Giving SMARTCARDS a New Face

That smartcard in your pocket can get even more interesting with electrophoretic displays. Here are some design tips to consider to make it happen.

Since their introduction in the 1970s, smartcards have become integral to a large proportion of financial transactions. Solution designers are on the verge of the next wave of evolution where active-matrix display technology is being integrated into a smarter new generation of smartcard to further improve security and usability for us humans.

Electrophoretic displays (EPDs) enable crystal-clear visibility of complex images, whether they're barcodes, card numbers, ticket IDs, or a mixture of graphics and characters. In this article, I'm going to outline some of the design considerations needed to accommodate a display in something so tightly defined by industry standards. Hopefully some of the information will also prove useful to other design challenges involving the addition of displays to smart portable devices.

IS THIS REALLY HAPPENING?

The simple answer is yes. My examples in this article relate to completed development projects for customers. Smartcards with active-matrix displays are allowing assets and packages to be labeled in more intelligent ways. They enable special active sensors to be included for applications like produce monitoring. And perhaps more relevant to you or me, they allow us to carry one solitary smartcard that can adopt the identity of the entire contents of your wallet. Or have ticketing cards that let you see the last 10 locations swiped.

WHAT A DESIGNER MUST CONSIDER

There's a well-defined standard that any designer needs to adopt to fulfill ISO10373/IEC 7816 for mechanical dimensions,

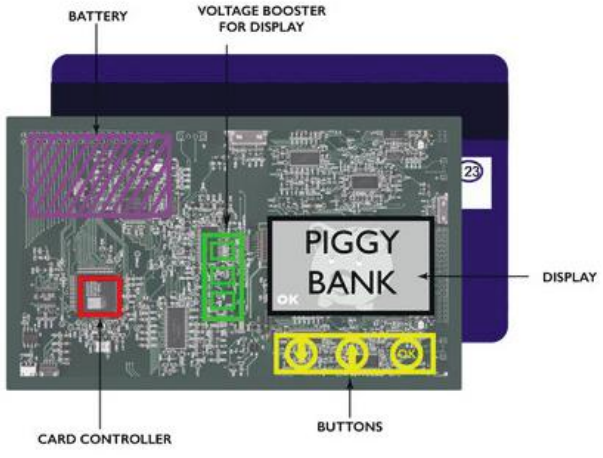

Smartcards can incorporate many thin technologies, from displays to buttons.

and more critically, the "bending" requirements for the smartcard in your pocket. It's challenging to meet these with any electronics product, but if you apply a display technology that is itself flexible, the outcomes are generally excellent. So, it doesn't necessarily present additional problems, just more of the same!

AND SIZE IS IMPORTANT

The application defines what size and geometry of display is required. We provide displays with an active area from 1.1 in. diagonal with a resolution of 150 ppi equivalent to 148 ×70 pixels, up to a 3.1-in. display with 105 ppi or 312 × 74 pixels. While the former can show a logo or multiple lines of a transaction, the latter can show your full credit-card number in its original size.

Display thickness is typically 480 µm, which is rarely problematic. However, it's possible to reduce this to as low as 300 µm to meet application demands (this is typically only needed where a smartcard is already very complex and crammed with other features). Some additional external components are required to update the display; a simple voltage booster *(see figure)* can be located near the edge of the card because it's relatively small and easy to track.

WILL IT BEND?

The all-important bend radius for smartcards isn't typically an issue even if you want a display to extend the full width of the card. This is because the bend radius for our EPD display is 30 mm for standard thickness and even lower for reduced thickness production units. More crucial is the IC display driver that's needed.

Much like the other critical processing elements of any smartcard, standard silicon devices are used with a package size of up to 10 × 1 mm based on the required number of outputs to drive a high-resolution display. The orientation and position within card should be carefully considered so that it's as close to the neutral axis as possible. A stiffener will typically be required to support the driver IC within the layers of the card.

The stiffener simply spreads the bending force around any components that can't flex, so that the overall bend radius of the card can still be achieved. There are various approaches to the way vulnerable components are positioned and grouped, and how stiffeners are used. Typically, stainless steel with a thickness of 0.1 to 0.2 mm is a tried and trusted material.

Furthermore, the stiffener doesn't have to go the full length of the card or even the display—it just needs to protect the driver IC so that it can't break or delaminate when subjected to bending forces. It doesn't add thickness to the display, since it's typically placed outside the active area where the display is much thinner.

SPI

A standard serial peripheral interface (SPI) is used to drive the display, so it's easy to integrate with almost any smartcard controller. Proximity to the display isn't particularly critical because EPD displays are deployed in applications that don't typically involve many display refresh cycles, making data corruption uncommon. An electrophoretic display only needs the transfer of a single image, not a continuous video stream.

Image data formats are 2-bit/pixel encoded bitmaps. Waveforms are already pre-programmed into the display to allow for different update modes. The user can select between a four grey-level update of about 0.7 seconds or a faster monochrome update of 0.3 seconds.

END OF LIFE AND POWER

It's easy to construct an NFC-powered smartcard that updates a built-in display only while being read or scanned—maybe to show that last transaction number or refresh date/barcode. An EPD display will then hold and display that image indefinitely without power, or until it's read or scanned again.

In more advanced applications, the use of a micro-battery, potentially rechargeable, gives the smartcard standalone interactivity. In other words, the user can interact with the card and select different information to be displayed as functions are accessed via buttons or a touch sensor. The typical power consumption for every display update is 5 mA (for less than one second), which needs to be factored in to any end-of-life projections for the application of the card.

Indeed, battery life is the biggest influence on end of life, but it need not be a major contributor. In many applications, operational life is projected at less than two years due to normal wear and tear and routine physical degradation.

Rechargeable battery options are readily available, so if longevity is required, such as when display updates are quite frequent in normal operation, wireless recharging solutions can be integrated. An example application for this would be dynamic tracking tags, used by logistics systems to track packages, parts, or assemblies; where human/legacy-machine readability is key, but tag information (barcodes and text) might be updated several times an hour.

PUSHING THOSE BUTTONS

Interactivity in smartcards needs something fingers can push *(see figure, again)*. Mechanical buttons blended into the card layers are relatively easy to accommodate physically, typically enabling scrolling up and down a pre-programmed list of options to change the mode of operation of the card, or show a variety of recent transactions. There are limits to the mechanical durability, but they're well-known and understood by manufacturers.

An obvious requirement that we've deployed for customers is a touch sensor layered on top of the display surface. This requires an additional film to be applied using an Optically Clear Adhesive, which adds at least another 50 µm to the overall thickness, plus 25 µm for the adhesive. It may be necessary to further protect the sensor in high-durability situations with additional hard coat or cover lens layers.

In such cases, total thickness may become a critical issue as the add up may be in the range of 200 to 300 µm, so total stack thickness (EPD display + touch solution) will already approach the ISO spec thickness limit for card applications of 0.84 mm max. Besides, the touch sensor will also have design impact on the power, as the power requirement and management will need more careful consideration. A battery will be critical, but the ability to, say, enter a PIN number or other numeric value might be hugely beneficial to the overall solution.

GOOD LUCK!

Smartcards literally changed the way we all make millions of transactions. Adding displays to show complex images is already extending smartcard deployment in real applications and I hope the tips above are useful if you're considering integrating EPDs in yours. Good luck with your next design.

INDEX

For more information on products or services visit our website www.electronicdesign.com. The advertisers index is prepared as an extra service. Electronic Design does not assume any liability for omissions or errors.

Lab Bench

WILLIAM WONG | Senior Technology Editor
bill.wong@informa.com

Radar Isn't Just for Cars and Aviation

Applications for low-cost radar systems now extend beyond the norm. Here, Technology Editor Bill Wong takes a look at a millimeter-wave development kit for industrial applications.

Millimeter-wave (mmWave) radar has found a home in many areas from self-driving cars and aviation applications, but low-cost radar systems are equally useful in a wide range of embedded applications. They can be used to detect people within a car to controlling entry doors more accurately than other methods such as ultrasonics.

I had a chance to check out Texas Instruments' IWR6843 60-GHz mmWave dev kit *(Fig. 1)* that targets industrial applications. TI has other kits that deal with different frequencies and apps, such as automotive. The chip, with its three transmitters and four receivers, uses a C674x DSP to handle signal processing. An ARM Cortex-R4F with ROM firmware handles object detection and interface control. Communication is possible with CAN-FD, UARTs, SPI, and I^2C. It also maintains six ADC channels for other chores. Internal memory features ECC protection.

What was equally impressive were the applications and application notes for the kit that are part of the mmWave Studio. The Studio is designed to present the mmWave sensor functionality through demos and system tools like the web-based mmWave Demo Visualizer and Sensing Estimators. The web-based tools connect to the hardware via a serial port, making it possible to be up and running in a very short period of time. The visualizer allows the configuration to be saved so that it can be used in an application.

Application development is done using TI's Code Composer Studio and the mmWave SDK. I didn't get too far with that yet due to time, but there are a number of projects to get started with from using the hardware with ROS (robot operating system)-based robots to gesture control.

The platform has many advantages over alternative range information systems, including the ability to detect and see through materials like glass and drywall. It can essentially ignore environmental factors such as smoke, rain, and lack of light. The system can generate a point cloud that's sufficient to distinguish people and other objects accurately.

Though the kit I tested uses a PCB-based antenna, developers can also take advantage of the antenna-on-package (AOP) solutions *(Fig. 2)*. It uses the same SoC I tested with antennas mounted on top of the package. This presents a number of advantages; in particular, all RF design/certification is handled by TI. Thus, a developer can use the AOP without having to address these issues.

There are benefits to using larger custom antenna arrays and often systems will need to employ multiple chips to handle such arrays. Nonetheless, a simpler, more compact solution is suitable for many applications.

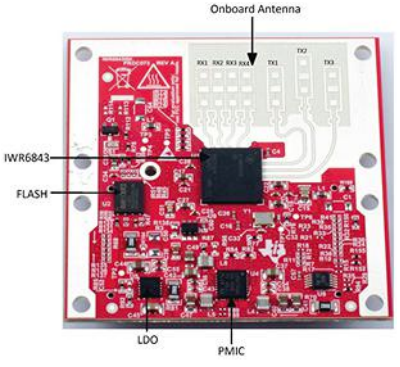

The IWR6843ISK is an antenna plug-in board for Texas Instruments' mmWave development kit.

The chip on the left is the IWR6843 antenna-on-package (AOP) that includes electronics and antennas for a complete millimeter-wave solution.

TI has taken a lot of the work out of developing an embedded radar solution—to the point that prototypes could be up and running in a very short period of time. They also provide direction and movement information that isn't possible with simpler proximity sensors, as well as operate in environments that would stifle or completely curb technologies like visual imaging systems. Checking out mmWave solutions is inexpensive and quick, so it's definitely worth looking into.

NEW! NEW! NEW! NEW!

pico® Technology

PicoScope®
9404-05 SXRTO
(Sampler eXtended Real Time Oscilloscope)

- 5 GHz bandwidth
- 1 TS/s effective sampling rate
- 4 channels, 12-bit resolution
- Under $15,000

A new format of broadband oscilloscopes

The PicoScope 9404-05 features four 5 GHz 12-bit channels, each supported by real-time sampling to 500 MS/s per channel and up to 1 TS/s (1 ps) equivalent-time sampling. Both the voltage and timing resolution specifications are characteristic of the highest performance broadband oscilloscopes.

To find out more visit www.picotech.com/A269

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

90% EFFICIENT, 100 A TO 400 A µMODULE REGULATOR COOLS FROM TOP, 4-SIDES, AND BOTTOM

LTM4700

- Single 100 A or dual 50 A µModule® regulator
- Scalable up to 400 A (eval board DC2784A-C)
- 12 V_{IN} to 1 V_{OUT}
  - 90% efficiency at 100 A
  - 47°C temperature rise with 200 LFM airflow, T_A: 25°C

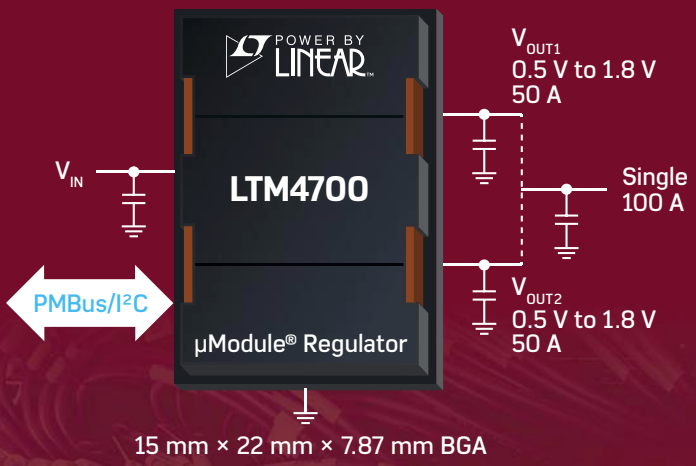

#PowerByLinear

mouser.com/LTM4700

MOUSER ELECTRONICS