Electronic Design®

THE AUTHORITY ON EMERGING TECHNOLOGIES FOR DESIGN SOLUTIONS

JULY/AUGUST 2020
electronicdesign.com

Formula E and IndyCar on the Auto Tech

FAST TRACK

p12

1.800.344.4539
DIGIKEY.COM

Digi-Key ELECTRONICS

10.6 MILLION+ PRODUCTS ONLINE | 1,300+ INDUSTRY-LEADING SUPPLIERS | 100% AUTHORIZED DISTRIBUTOR

ECIA MEMBER
Supporting The Authorized Channel

BROADCOM OPTOCOUPLERS

Take the Risk Out of High Voltage Failure!

Only Optocouplers Meet or Exceed International Safety Standards for Insulation and Isolation.

Stringent evaluation tests show Broadcom optocouplers deliver outstanding safety performance and exceptional high voltage protection. Alternative galvanic isolation technologies such as capacitive and magnetic isolators do not perform anywhere near the levels of insulation and isolation provided by Broadcom optocouplers.

GO with OPTO! 50 YEARS
BROADCOM® OPTOCOUPLERS

IEC 60747-5-5 Certified

To learn more visit
broadcom.com/optocouplers

BROADCOM®
connecting everything®

IN THIS ISSUE

FEATURES

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 Endeavor Business Media, 1233 Janesville Ave., Fort Atkinson, WI 53538. Paid rates for a one-year subscription are as follows: $60 U.S., $90 Canada, $120 International. Periodicals postage paid at Fort Atkinson, WI and additional mailing offices. Printed in U.S.A. Title registered in U.S. Patent Office.

Altech®

Your Source for PCBs & Test Measurement

Printed Circuit Board Terminal Blocks with ...

Push-in TECHNOLOGY

- Unique ergonomic 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

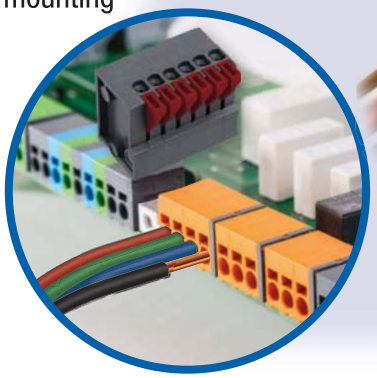

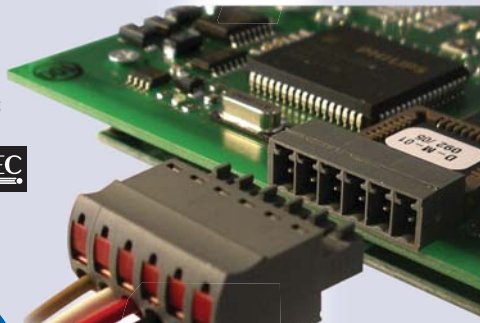

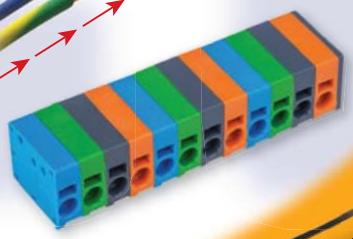

Also Available

PCB Fixed Terminal Blocks
PCB Plug-In Connectors

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

Well Established TEST & MEASUREMENT

Alligator Clips
- Various clamping ranges
- Stainless steel versions
- Soldering connections for “do it yourself” connections

Test & Measurement Sets
- 0.64, 2mm or 4mm System
- 4mm Retractable Sleeve System
- 4mm Safety System

Plunger Style Test Clips
- Measuring “difficult to access” testing points
- Various designs for different applications
- Handle with multiple connection points

Measuring Leads
- Wide variety of colors, length and connections
- Double insulation
- Highly flexible also at low temperatures

Test Probes
- Slim and compact body for easy access
- Stainless steel tips
- Spring-loaded tips

Test Plugs
- Large color spectrum
- Various connection types
- Different cross sections

Oscilloscope Test Probes
- 10:1 probe ratio
- Exchangeable spring tip
- Low Input Capacitance

Test Jacks
- Large variety in colors
- Tin/ nickel/ gold plating
- Screw-in and press-in versions

Altech Corp.®

Your Source for Automation & Control Components

info@altechcorp.com • 908.806.9400

www.altechcorp.com

ON ELECTRONICDESIGN.COM

Life at the Sensor Edge: The Increasing Demands of Autonomy

With autonomous-driving capabilities advancing at a lightning-quick pace, what are the latest trends pushing the electronics and sensor technologies that enable full self-driving functionality?

https://www.electronicdesign.com/markets/automotive/article/21134972/life-at-the-sensor-edge-the-increasing-demands-of-autonomy

Why Bulletproofing the Li-Ion Supply Chain Could Be a Lifesaver

Lithium-ion batteries are critical to today's power infrastructure, from mobile devices to cars, and concerns continue to escalate about their supply line.

https://www.electronicdesign.com/power-management/whitepaper/21135518/why-bulletproofing-the-liion-supply-chain-could-be-a-lifesaver

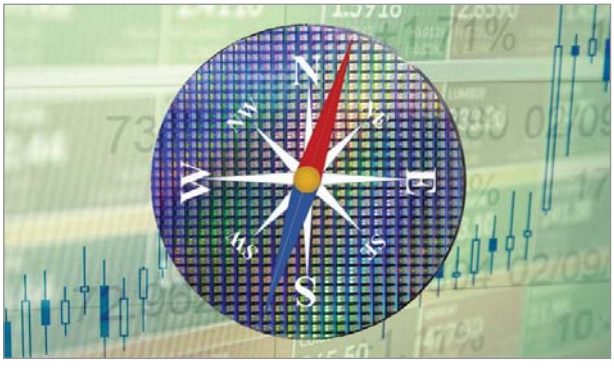

Emerging Memory Market is Finding its Direction

Emerging memories have been relegated to niches for decades, but that scenario has taken an abrupt turn. Recent market dynamic changes are suddenly opening the door to their widespread use.

https://www.electronicdesign.com/industrial-automation/article/21135237/emerging-memory-market-is-finding-its-direction

A False Foundation

Asimov's *Foundation*, which centers around Hari Seldon and his science of psychohistory, is coming to Apple TV, but will it touch on conspiracy and artificial intelligence?

https://www.electronicdesign.com/altembedded/article/21137005/a-false-foundation

FOLLOW US ON FACEBOOK AND TWITTER

(ElectronicDesign) (@ElectronicDesgn)

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.

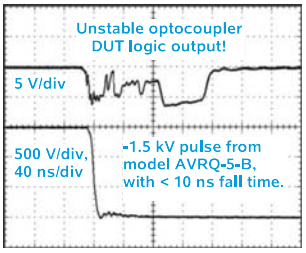

- Kilovolt amplitudes (±1, ±1.5, -2 kV)
- Transition times down to 10 ns, dV/dt rates up to 120 kV/us
- Switchable daughterboards to handle a variety of DUT package styles
- GPIB, RS-232 ports standard
- Ethernet / VXI optional

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!

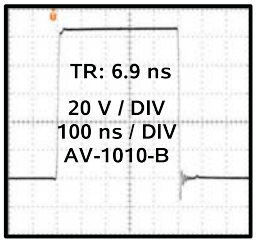

Model AV-1015-B:	50 Volts, 10 MHz
Model AV-1010-B:	100 Volts, 1 MHz, 25 ns to 10 ms, 10 ns rise time
Model AV-1011B1-B:	100 Volts, 2 ns rise time
Model AV-1011B3-B:	30 Volts, 0.5 ns rise time

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

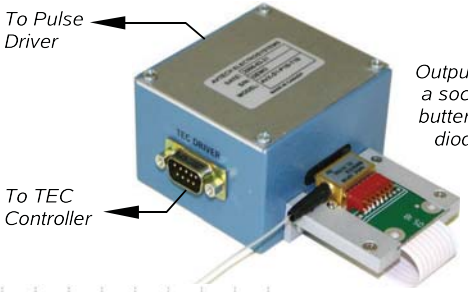

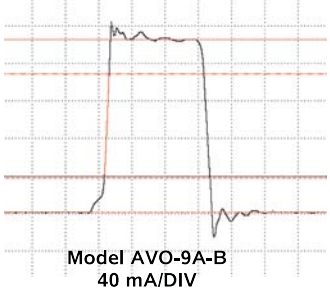

Model AVO-9A-B
40 mA/DIV
1 ns/DIV

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:

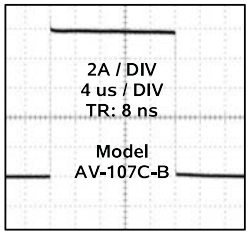

Series	I, V	PW	TR
AV-107	2 - 20 A, 60 V	0.2 - 200 us	10 - 30 ns
AV-106	5 - 100 A, 100 V	0.5 us - 1 ms	50 ns - 1 us
AV-108	12.5 - 200 A, 100V	2 us - 1 ms	5 - 15 us
AV-109	10 - 100 A, 5 V	10 us - 1 s	10 us
AV-156	2 - 30 A, 30 V	1 us -100 ms	0.2 - 50 us

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

ElectronicDesign

JULY/AUGUST 2020

EDITORIAL

SENIOR CONTENT DIRECTOR: **BILL WONG** bwong@endeavorb2b.com
SENIOR EDITOR: **DAVID MALINIAK** dmaliniak@endeavorb2b.com
SENIOR STAFF WRITER: **JAMES MORRA** jmorra@endeavorb2b.com
CONTRIBUTING EDITOR: **LOUIS E. FRENZEL** lfrenzel@endeavorb2b.com
ASSOCIATE EDITOR/COMMUNITY MANAGER: **ROGER ENGELKE** rengelke@endeavorb2b.com

ART DEPARTMENT

GROUP DESIGN DIRECTOR: **ANTHONY VITOLO** tvitolo@endeavorb2b.com
ART DIRECTOR: **JOCELYN HARTZOG** jhartzog@endeavorb2b.com

PRODUCTION

GROUP PRODUCTION MANAGER: **GREG ARAUJO** garaujo@endeavorb2b.com
PRODUCTION MANAGER: **DEANNA O'BYRNE** dobyrne@endeavorb2b.com

AUDIENCE MARKETING

USER MARKETING MANAGER: **DEBBIE BRADY** dmbrady@endeavorb2b.com

FREE SUBSCRIPTION / STATUS OF SUBSCRIPTION / ADDRESS CHANGE / MISSING BACK ISSUES
OMEDA **T** | 847.513.6022 **TOLL FREE** | 866.505.7173 **FAX** | 847.291.4816 electronicdesign@omeda.com

SALES & MARKETING

REGIONAL SALES REPRESENTATIVES:

AZ, NM, TX: **GREGORY MONTGOMERY** **T** | 480.254.5540 gmontgomery@endeavorb2b.com

AL, AR, SOUTHERN CA, CO, FL, GA, HI, IA, ID, IL, IN, KS, KY, LA, MI, MN, MO, MS, MT, NC, ND, NE, OH, OK, SC, SD, TN, UT, VA, WI, WV, WY, CENTRAL CANADA: **JAMIE ALLEN** **T** | 415.608.1959 **F** | 913.514.3667 jallen@endeavorb2b.com

CT, DE, MA, MD, ME, NH, NJ, NY, PA, RI, VT, EASTERN CANADA:
ELIZABETH ELDRIDGE **T** | 917-789-3012 eeldridge@endeavorb2b.com

AK, NORTHERN CA, NV, OR, WA, WESTERN CANADA: **STUART BOWEN** **T** | 425.681.4395 sbowen@endeavorb2b.com

INTERNATIONAL SALES:

GERMANY, AUSTRIA, SWITZERLAND: **CHRISTIAN HOELSCHER** **T** | 011.49.89.95002778
christian.hoelscher@husonmedia.com
BELGIUM, NETHERLANDS, LUXEMBURG, UNITED KINGDOM, SCANDINAVIA, FRANCE, SPAIN, PORTUGAL:
LARA PHELPS lara.phelps@husonmedia.co.uk
ITALY: **DIEGO CASIRAGHI** diego@casiraghi-adv.com
PAN-ASIA: **HELEN LAI** **T** | 886 2-2727 7799 helen@twoway-com.com
PAN-ASIA: **CHARLES LIU** **T** | 886 2-2727 7799 liu@twoway-com.com

ARTICLE REPRINTS: reprints@endeavorb2b.com
LIST RENTALS/ SMARTREACH CLIENT SERVICES MANAGER: **MARY RALICKI** **T** | 212.204.4284 mralicki@endeavorb2b.com

DIGITAL

SENIOR DIGITAL INNOVATION & STRATEGY DIRECTOR: **RYAN MALEC** rmalec@endeavorb2b.com

DESIGN ENGINEERING & SOURCING GROUP

VICE PRESIDENT, DESIGN & ENGINEERING: **TRACY SMITH** **T** | 913.967.1324 **F** | 913.514.6881 tsmith@endeavorb2b.com
VICE PRESIDENT OF MARKETING SOLUTIONS: **JACQUIE NIEMIEC** jniemiec@endeavorb2b.com

ENDEAVOR BUSINESS MEDIA, LLC

331 54th Ave N., Nashville, TN 37209 USA | www.endeavorbusinessmedia.com
CEO: **CHRIS FERRELL**
CRO/CMO: **JUNE GRIFFIN**
CFO: **WILLIAM NURTHEN**
COO: **PATRICK RAINS**
CHIEF ADMINISTRATIVE AND LEGAL OFFICER: **TRACY KANE**
EVP KEY ACCOUNTS: **SCOTT BIEDA**
EVP SPECIAL PROJECTS: **KRISTINE RUSSELL**
EVP GROUP PUBLISHER – DESIGN & ENGINEERING, ENERGY, BUILDINGS & CONSTRUCTION: **REGGIE LAWRENCE**

Electronic Design | Machine Design | Microwaves & RF | Hydraulics & Pneumatics | Source ESB | Source Today | Evaluation Engineering

More new products in stock

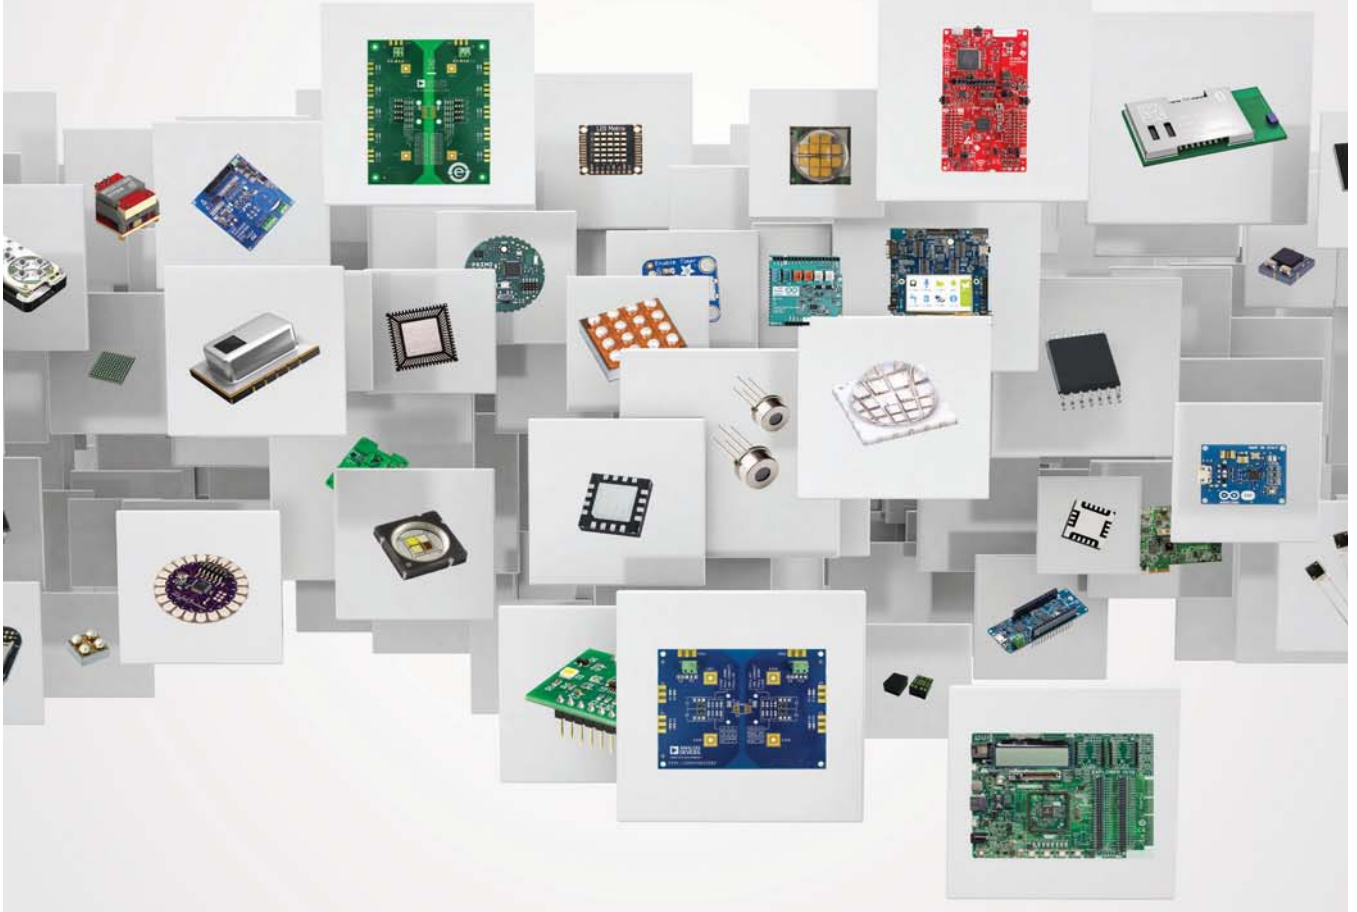

ORDER WITH CONFIDENCE

MOUSER ELECTRONICS

Editorial

WILLIAM WONG | Senior Content Director
bwong@endeavorb2b.com

Evaluating Technical Worth DURING A PANDEMIC

COVID-19 has really done a number on the economy, leading to reevaluation of the need and worth of just about everything. The pandemic has forced the cancellation of almost every tradeshow, and large gathering, especially those not employing masks and social distancing, have caused major spikes in infections and deaths.

Cancellation of large events and closing of things like schools has forced many to turn to teleconferencing and video conferencing. Virtual conferences are now the name of the game, along with webinars. The number of webinars that we here at Endeavor Business Media are hosting has increased as companies scramble to get their message out to those with dollars still left to spend.

Distance learning has changed from the exception to the rule, while other areas that were doing some remote interaction like telemedicine are seeing massive growth in only a few months. Tools such as Zoom and Microsoft Teams have becoming ubiquitous rather than a growing niche.

Unfortunately, going from an in-person interaction to a virtual one isn't always easy. I was able to help change the Mercer Science and Engineering Fair I work with into a virtual event within a couple weeks. I knew the tools and reconstructed the infrastructure to handle judging via video conferencing. It not only required getting judges and students together on a Zoom conference, but also having students upload photos and other materials for the judges to examine. The effort would not have succeeded without those involved tackling new procedures and dealing with cantankerous computers, cameras, and content.

Distance learning is neither cheap nor easy, and the results aren't always great if the sufficient effort and proper methods aren't employed. Likewise, those using and supporting these solutions need to keep in mind that the different approaches may be better or worse compared to the alternatives. Often, it's easier for attendees to watch recorded sessions at their leisure rather than try to match the schedule at in-person events. In theory, someone utilizing recorded sessions could see more sessions because they no longer must include travel time and scheduling is obviously easier.

The question is: How valuable are remote interactions versus in-person meetings? Obviously, remote meetings are the norm these days because of the pandemic. Do the advantages of remote meetings offset disadvantages compared to in-person meetings, and should the cost of using remote access be higher or lower?

Many encounter the "free" versions of these tools or the cost is absorbed by the provider. A related issue is that some revamped in-person conferences that charged for attendance were now "free" as they went virtual. Switching back to a paid conference and still remain virtual could be a challenge. Likewise, telemedicine is having issues with some insurance companies wanting to pay less for remote sessions, even though they may actually take more time for both parties.

These cost issues aren't unique to one environment. Users may not be aware of other possible costs. For example, telemedicine security requirements and associated software will likely be higher than hosting a webinar that's open for registration to anyone. There are even opportunities such as the ability to more easily track attendees and/or provide incremental pricing—for instance, making keynotes free, having a fixed set of sessions for a fee, and maybe additional fees for specific training sessions.

One certainty is that the focus on remote-access solutions has skyrocketed and innovation is coming in this area. It means examining new options for marketing and learning, as well as in areas such as development, service and support. Determining the worth of the remote-access technology and presentation of content in a remote fashion is something we will all have to do as we move toward a new norm. ed

The 2020 Mercer Science and Engineering Fair judging was conducted using video conferencing.

The higher standard of high voltage DC/DC power supplies.

UMR Collection & PMT Series

Cost effective, analog controls and monitors, and the best performance...

...for what you need now

SPS Series

Fully digital, extremely customizable, with serial communications...

...for making the future now

Outputs up to 6kV @ 30W

Available as biasing or capacitor charging modules

Contact DTI today to discuss your high voltage design.

DEAN TECHNOLOGY

+1.972.248.7691 II www.deantechnology.com

News

ANALOG DEVICES TO BUY
Maxim Integrated in More Than $20 Billion Deal

Analog Devices has agreed to buy rival Maxim Integrated for more than $20 billion in stock to expand its scale and broaden its product portfolio in the market for analog semiconductors. The all-stock deal would value the combined company at more than $68 billion, creating an analog semiconductor giant that will better compete with market leader Texas Instruments. Analog Devices said the deal would scale up its footprint in fast-growing segments such as industrial, automotive, data centers, and healthcare.

Under the terms of the proposed agreement, Maxim shareholders would receive 0.63 of Analog Devices stock for each share they own. Once the deal closes next year, shareholders in Analog Devices would own about 69% of the resulting company, with Maxim's stockholders owning the other 31% of the firm.

Both Analog Devices and Maxim are major players in the market for analog semiconductors that are key components of nearly all electronic devices, from smartphones and thermostats—to cars and the factory equipment that manufactures them. These types of chips handle the fundamental chore of translating electrical signals—including temperature, pressure, sound, and other signals relayed from the sensors slapped on a device—into digital data.

According to IC Insights, sales of analog chips accounted for around 13%, or $55.2 billion, of the global chip market in 2019. Overall sales of semiconductors slipped 12% to $412 billion.

The chips are also used in other areas, including power management for smartphones and other consumer electronics to boost the energy efficiency of the device and protect the battery from malfunctioning, overheating, and bursting into flames. Analog Devices also offers radio-frequency, millimeter-wave, and other analog chips for clarifying and conditioning radio signals used by cellular networks, including 5G.

Analog Devices and Maxim also compete with Texas Instruments in microcontrollers and embedded components that serve as the brains of millions of devices on factory floors, in cars, and household devices, ranging from thermostats to televisions to washing machines. Other chips Maxim sells are designed to protect devices, like connected door locks, from attackers trying to steal secret keys, passwords, or other data on the devices.

ADI said the deal would combine Maxim's strength in automotive and data-center segments with its footprint in the industrial and communications markets. Once the transaction closes, the combined company will sell more than 50,000 products, allowing it to offer more complete solutions. The company said it will serve over 125,000 customers and be positioned to capture a bigger portion of the overall analog market.

"We see lots of opportunity in both the short term to be able to take the existing portfolio of Maxim products to more and more customers," Vincent Roche, chief executive officer of Analog Devices, said on an analyst conference call to discuss the deal. "And over the longer term, to be able to produce more complete suites of product offerings for our customers by combining the complementary portfolios," he added.

Analog Devices said the deal would strengthen its lineup of analog chips, which generally sell for less than $1 each and often less than 50 cents. But these types of chips are also some of the most profitable in the industry because they can be manufactured on cheaper processes. Analog Devices' gross margin—the percentage of sales remaining after the cost of goods is subtracted—came to 70% last year. Maxim's gross margin was 65%.

Analog Devices would also swoop up thousands of Maxim's analog engineers in the deal. Between them, the companies employed more than 10,000 engineers and invested $1.5 billion in research and development in 2019. Competition for analog engineering talent has been escalating in recent years, the company said. Once the deal closes, Analog Devices believes that it will be in a better position to lure the chip industry's top analog engineers.

Overall, the deals have transformed it into the No.2 player behind Texas Instruments in the market for analog ICs, which is projected to grow from $55 billion in 2019 to $60 billion by 2023. ■

Advanced Magnetics for ADAS

From high-current, high-efficiency power inductors to filter components for a variety of communications buses, Coilcraft has the magnetics for all of your Advanced Driver Assistance Systems

Coilcraft offers a wide range of AEC-Q200 qualified products engineered for the latest advanced driver assistance systems, including high-temperature, high power density power inductors for radar, camera and LiDAR applications.

Our compact, low-profile WA8351-AL ultrasonic sensor transformer offers excellent temperature stability up to 125°C and high performance for time-of-flight (TOF) sensing.

Also choose from our broad selection of common mode chokes and filter elements for a variety of communications buses.

To learn more about our advanced solutions for ADAS and other automotive/high-temp applications, visit us at **www.coilcraft.com/AEC.**

Coilcraft®

WWW.COILCRAFT.COM

WILLIAM G. WONG | Senior Content Director

Automotive Tech is OFF TO THE RACES

COVID-19 may reduce the number of races, but not the amount of new technology for competitions like IndyCar and Formula E.

Today, fans in the stands for races from NASCAR to Formula E varies from none to a fraction of the crowds that are usually drawn to see cars and drivers trying to best the competition. As for the races themselves, though there are less of them, they should be just as fast and exciting as before.

Changes in these competitions are incremental, but they highlight safety and performance improvements. Here, we'll take a look at IndyCar, since we're a sponsor of DragonSpeed team and were able to get an inside look, as well as peek into Formula E courtesy of Envision Virgin Racing. The cars look similar with many common safety features. However, the underlying motive technology is radically different.

COMMON PLATFORMS TIGHTEN THE COMPETITION

Most of the racing events specify a good bit of commonality to keep the competition tight. The chassis, tires, and engine of an IndyCar are standard; teams are only allowed to make limited modifications to specific parts of the systems. Often the components come from third parties, which supply all of the teams with essentially identical components.

For example, teams lease engines from Chevrolet and Honda. The 2022 configuration is a 2.4-liter, twin-turbocharged V-6 engine that's projected to generate over 900 horsepower or more than a 10% improvement compared to the current 2.2-liter platform. The engines run at a maximum of 12,000 RPM. The suppliers

1. IndyCars are now equipped with Red Bull Advanced Technologies' Aeroscreen to protect drivers from flying debris.

provide an engineer that oversees each engine, keeping track of its operation to the point of removing the engine control unit (ECU), which means the engine can only be run when the engineer is present. An engine will be replaced if it doesn't meet the specifications.

For Formula E, there's a similar consistency. However, a standard battery system is used as well as a specification on how much power can be expended during the race. More on that later. Formula E teams do have more leeway on the motor that they supply, although there are regulations here as well. At one point, a team utilized a pair of motors, which is no longer allowed.

Most race cars have a host of sensors available to the driver. These days, the information is provided wirelessly to the pit crew, which may not be physically located in the pit area. The number and placement of sensors is often arbitrary during testing, but there are typically limitations during practice and the race. For example, the rotation speed of the tires in Formula E isn't allowed, although it's possible to estimate this by the speed of the engine. Obtaining tire speed using markings on the tire plus an external camera isn't allowed either.

A digital twin for a car has moved from the crew chief's head into a shared model for the pit crew and team to analyze and provide the driver with insight into the car's capability and performance. Most people viewing a race will never see most of the team other than the pit crew changing tires, but the team makes a big difference in how well a driver and car will perform.

All of these regulations are designed to make the skill of the drivers and crew more important to winning versus building a faster car.

SAFETY AND OPTIMIZING AN INDYCAR

IndyCars already have already been equipped with a Suspension Wheel/Wing Energy Management System (SWEMS). It consists of multiple cables connected to keep objects attached to the frame should an accident occur. In addition, multiple interlocks incorporated on the fueling systems help prevent accidental spills.

One standard piece of equipment recently added to IndyCars is a windshield and titanium frame that protects the driver in case of an accident. Driver Justin Wilson died in 2015 when he was struck by a piece of equipment from a wreck that occurred in front of him. The cockpit is now almost completely enclosed with Red Bull Advanced Technologies' Aeroscreen to protect from debris *(Fig. 1)*.

One place where teams can customize their car is the suspension system. These mechanical tweaks are usually refined during pre-race practice and time trials. Vent changes can also be made regularly, based on temperature info provided wirelessly to the pit crew. Changes can be made during the race, but these tend to be as minimal as possible for obvious reasons.

IndyCar races have pit stops for changing tires and adding ethanol to the tank. The aforementioned tweaks can be done, too, but normally a stop is as fast as possible. The pit crew may even peel off a layer of the windshield that's covered by multiple plastic layers should the top layer become pitted.

ELECTRIC RACING

Formula E races run for 45 minutes using a battery pack with a fixed amount of energy *(Fig. 2)*. The Gen2 Formula E cars will have a 54-kWh battery that's contained in a carbon fiber case. It's crash tested and has a safety check that provides a red/green status indication. A red error condition indicates that the driver needs to hop out of the vehicle without touching the car and ground at the same time. An overheated battery is one type of error condition.

The battery capacity is insufficient to run the engines at full race power for the duration of the race, something drivers must keep in mind even with regenerative braking. More than one car has been unable to complete a race due to insufficient power. Regenerative braking is performed using paddles on the steering wheel. Such braking is done via the rear wheels driven by the motor.

The front brakes are conventional hydraulics and thus don't provide regenerative braking. There are hydraulic brakes for the rear wheels; the regenerative braking is power-limited. Drivers also need to keep in mind that the battery is fully charged at the start of the race; thus, regenerative braking provides essentially no advantage at that time.

Drivers typically try to exit corners as quickly as possible. The driver will get a beep in their headset from the car to indicate that they should lift off the throttle.

2. Formula E cars use the same battery pack. Battery-management and regenerative-braking support can be key to a team's success.

Teams provide their own motor and inverters. The battery-management system is part of the battery pack, so teams can't modify it. They can tweak the suspension like IndyCars.

The Formula E Gen1 race was done using a pair of cars with a single pit stop to switch cars. The latest version Gen2 has no pit stops; however, there could be one in the future if fast charging becomes an option. Tires have less of an effect on the race results compared to other races, due to the length and speed of the races.

Teams get a minimal amount of practice time on a new track to limit the amount of analysis and simulation that teams can perform. Of course, that results in lots of late nights to analyze and incorporate sensor information from practice runs.

WATCHING THE RACES

If you haven't watched a race, then you may not know about the differences that go beyond what the cars look like. For instance, take street courses versus speedways. The former contains many tight turns and narrow roads. Passing is difficult and overall speeds are typically lower than other races. Speedways usually involve oval tracks that are wider and often banked, allowing for overall higher speeds.

How races are won is yet another matter. Races can be a fixed number of laps and the winner crosses the finish line first after completing the set number of laps. There are other ways to run a race, though. Endurance races like "24 Hours of Le Mans" or "12 Hours of Sebring" have teams trying to rack up the most laps in a set time. The alternative

Formula E's attack mode provides a 35-kW boost; however, it's enabled by driving through a specific area *(Fig. 3)*.

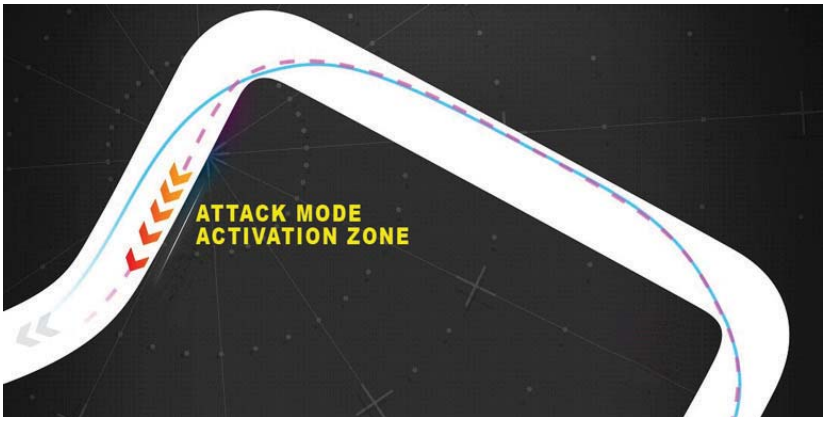

3. Formula E attack mode requires the car to move through a designated area that normally forces the driver to slow down or cover more ground to offset the later advantage.

"24 Hours of LeMons" takes this approach as well, but the price of the cars can't exceed US$500. Longer races usually employ multiple drivers that trade off during pit stops.

Formula E uses a variation of the fixed time limit. The next-to-last lap starts when the leader crosses the finish line after time expires. Hopefully, teams have reserved enough battery energy to finish quickly.

Competitions have tried to make things more interesting by essentially providing a turbo or high-performance mode that the drivers can exploit. This is normally a limited option that can be implemented in numerous ways. Formula 1 adds an overtake button, while IndyCar calls theirs "push-to-pass." This provides a power boost for a fixed time period, which is essentially an extra 40 horsepower for 10 seconds. Information about when a driver uses the feature is made available at the timing stand, but teams must use their pit-to-car radio to notify the driver.

Formula E's attack mode provides a 35-kW boost; however, it's enabled by driving through a specific area *(Fig. 3)*. The area is marked, making it easy for drivers and spectators to identify it. Driving through the area also takes longer to offset the advantage. The additional energy is available for four minutes from the time the car exits the area.

Using the high-performance mode makes for interesting tradeoffs—it can be initiated to stay ahead of someone else who has activated the mode or to pass those who haven't activated it. The high-performance mode usually has limitations like the duration or number of uses. The teams are often required to use this feature, although the advantage is one that they would typically want to use more.

The different race types, limitations, requirements, and winning methodologies impact how the car is driven and how the teams work together with their driver. Power utilization, tire wear, and other factors come into play as teams push the envelope. The details are often lost to casual observers. For example, Formula E cars have a limited amount of energy for the entire race and the engine efficiency changes based on speed. Run too fast and you run out of energy before the race is completed.

All of these differences are why drivers tend not to move between difference race platforms. There's a commonality between the cars due to physics, but the subtleties can make a big difference in how well a driver may perform. It takes time to learn a new environment and gain the muscle memory to quickly adapt to the changing race conditions. ed

CHARLES PAO | Senior Marketing Specialist, Sensor Fusion Business Unit, CEVA

Dealing with Magnetic Interference

Handling magnetic interference in sensors involves understanding your environment and developing algorithms to detect interference.

A while ago, I covered a glossary of terms you would need to know related to sensors, including their requirements, calibration, output vectors, and more. I'd like to take it a step further and discuss the most common question I get about sensors and how to work around it: Usually it's something like, "I want to use the magnetometer, but I'm not sure how to deal with magnetic interference."

SENSOR REFRESH: THE BASICS

For the purposes of this discussion, let's assume we're working with a 9-axis inertial measurement unit (IMU) comprised of a 3-axis accelerometer, gyroscope, and magnetometer.

It's important to understand that each sensor represents a part of the orientation puzzle, so let's have a quick review. The accelerometer can measure the direction of gravity relative to the sensor, and thereby estimate the tilt (pitch and roll). The gyroscope determines the change in angular position by measuring the angular velocity. The magnetometer measures the strength and direction of the magnetic field, which can sometimes measure the direction of North and thereby estimate the heading (yaw).

The following suggestions are based around some general principles when using an IMU:

- You can trust the gyroscope for relative orientation changes over short periods of time.
- You can trust the magnetometer in stable magnetic environments.
- You can trust the accel when it's being used in an application without constant acceleration.

And by "trust" these sensors, I mean you can use their output with a reasonable level of confidence.

HOW TO HANDLE MAGNETIC INTERFERENCE

One of the most common questions I receive from customers is how to deal with magnetic interference. The answer to this, as you might expect, depends

Design Note

System-Level Solution for High Density, Isolated Analog Output Module for Process Control Applications

Albert O'Grady, System Applications Manager

Introduction

When designing channel-to-channel isolated analog output modules for process control applications, such as programmable logic controller (PLC) or distributed control system (DCS) modules, the main trade-off is usually between power dissipation and channel density. As module sizes shrink and channel densities increase, the power dissipation per channel must decrease to accommodate the max power dissipation budget for the module. Higher channel density also means that there is less PCB real estate available for each channel.

System-Level Solution

Figure 1 shows the AD5758 and ADP1031 system solution that solves both the power dissipation and space challenges, allowing a high level of integration. This design note shows how an 8-channel module with less than 2 W of power dissipation for all channels can be produced in compact form factor, as shown in figure 2.

The ADP1031 solves the isolation and size challenges, providing 300 V basic isolation for power and data, while the AD5758 provides a low power dissipation, precision, configurable current, or voltage output channel. The AD5758 implements a technique called dynamic power control (DPC) to minimize power dissipation in the module under worst-case operating conditions when the part is configured for current output. It does this by continuously tracking the output voltage and reducing the power supply to the output driver to the minimum required to maintain the output load current—implemented via an integrated, programmable, high efficiency buck converter. When DPC is enabled for current output modes, the AD5758 automatically adjusts the DPC voltage to minimize power dissipation under all load conditions.

The ADP1031's design is optimized to provide efficient isolated power for the AD5758 under worst-case load conditions so that the total channel power dissipation is minimized. The high speed SPI channel integrated into the ADP1031 is also designed to reduce power dissipation when active and enters a low power state when inactive.

Solution Size and Power Dissipation

The high level of integration allows each isolated channel to fit in an area of less than 400 mm^2 on a two-sided PCB. This includes all associated passive components and isolation

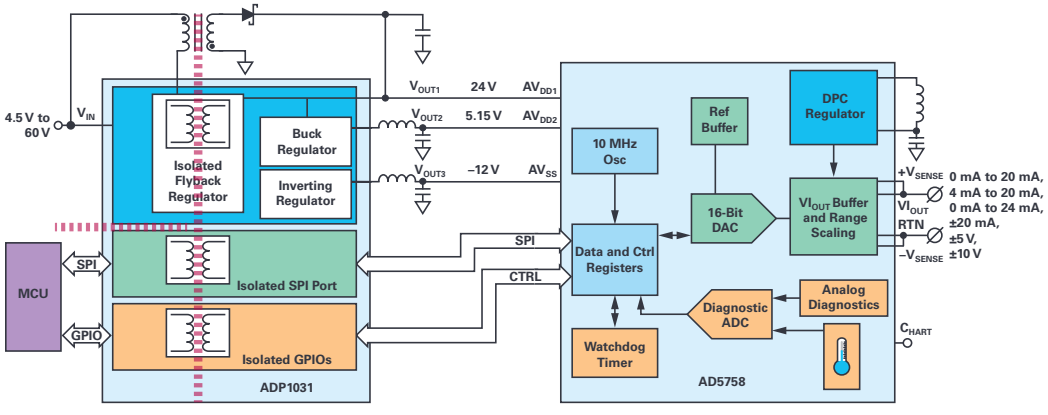

Figure 1. System-level approach enables the design of precision, high density, isolated analog output module for process control.

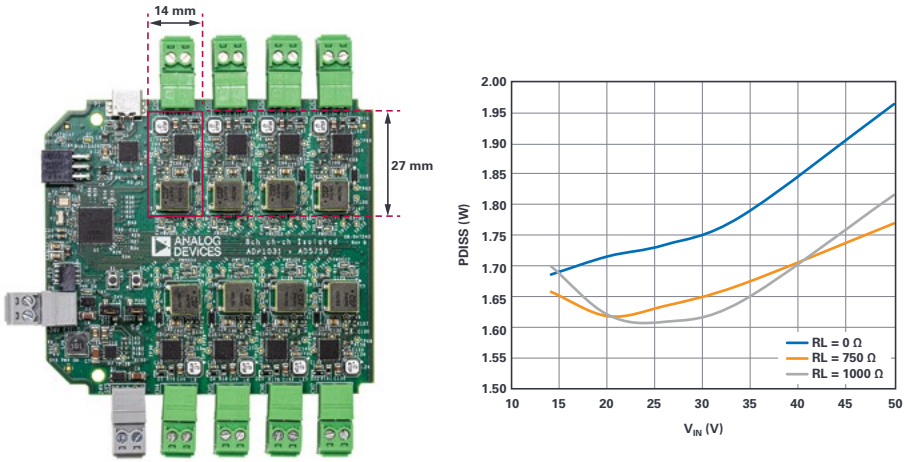

Figure 2. AD5758 and ADP1031 8-channel board with module power dissipation vs. supply voltage and load.

spacing. This solution achieves best-in-industry power dissipation of less than 2 W for all eight channels operating under worst-case power dissipation conditions.

EMC Robustness Performance

The AD5758 incorporates line protectors on all pins that may potentially be connected to screw terminals (VI_{OUT}, $+V_{SENSE}$, and $–V_{SENSE}$). The line protectors protect these pins from positive and negative voltages up to ±38 V by limiting the voltage internally to the V_{DPC}+ and AV_{SS} rails. If a voltage outside of these limits is detected on the VI_{OUT} pin, an error flag is set, which can be read back over the SPI port.

Extensive EMC testing has been completed on the AD5758 and ADP1031 system. See Table 1 and Table 2 for a summary of the test results.

Conclusion

The AD5758 and ADP1031 system-level solution enables a robust, compact, 8-channel, channel-to-channel isolated analog output module, which achieves best-in-industry power dissipation of less than 2 W for all eight channels operating under worst-case power dissipation conditions.

About the Author

Albert O'Grady is system applications manager with the Industrial Automation Technology Group at Analog Devices in Limerick, Ireland. His main focus is on industrial automation and process control applications. He can be reached at *albert.ogrady@analog.com*.

Table 1. Emissions Performance Summary

Test	Basic Standard	Frequency Range (MHz)	Limits	Measured Minimum Margin (dBµV/m)	Result
Radiated emissions	CISPR 11, Class B	30 to 1000	30 dBµV/m 30 MHz to 230 MHz; 37 dBµV/m 230 MHz to 1000 MHz	9.25	Pass

Table 2. Immunity Performance Summary

Test	Basic Standard	Test Levels	Performance Criterion	Result
Conducted Immunity	IEC 61000-4-6	10 V/m	A	Pass
Radiated Immunity	IEC 61000-4-3	10 V/m	A	Pass
ESD	IEC 61000-4-2	±6 kV contact	B	Pass
ESD	IEC 61000-4-2	±12 kV air	B	Pass
ESD	IEC 61000-4-2	±30 kV coupling	B	Pass
EFT	IEC 61000-4-4	±4 kV	B	Pass
Surge	IEC 61000-4-5	±4 kV	B	Pass

on the situation. So, let's go over a few scenarios:

Outdoor Use in a Rural Setting

If you're using the magnetometer in a rural setting, the changes and interference in the magnetic field should be few and far between, and you should be able to use its output to acquire long-term orientation information. Fusing this with information from the accelerometer and gyroscope produces a full picture. A key thing to keep track of is sudden changes in the magnetic reading and potentially rejecting them as interference.

Outdoor Use in a City

An outdoor urban setting has lots of potential magnetic interferences, from metal pipes to large motor vehicles. In short, it's harder to get consistent output from the magnetometer due to the number of potential interferences compared to rural settings.

Developing a robust magnetic-interference rejection algorithm is crucial, and ensuring that the magnetometer is trusted more when in stable environments is key. In those times when a magnetic environment isn't stable, you can rely on the short-term orientation accuracy of the gyroscope and correct its orientation over time when there's a more stable magnetic environment.

Indoor Use

Indoor applications have just as many possibilities for magnetic interference as an urban environment. Variations exist, however, between large buildings reinforced with concrete and steel compared to smaller buildings with more wood-based construction. The former will often contend with largely unstable magnetic fields, and even when reasonably stable, it can be dramatically warped to point in a different direction from the Earth's magnetic field.

Understanding the setting you're in and adjusting to it is important. For instance, in an indoor application with a warped magnetic field, it may be better to judge orientation based on relative differences over absolute direction.

This solution revolves around one of the principles we discussed above. In a less-warped environment, it also holds. Trust the mag when the magnetic environment is stable. Create a magnetic-interference rejection algorithm for improved consistency. Lean on data from the accelerometer and gyroscope at times when your magnetic environment is less stable. Is this sounding similar? It should! Because the key principles are the same.

To summarize:

- Trust the magnetometer in stable magnetic environments
- When not in stable magnetic environments, lean on the gyroscope and accel for orientation information
- Develop a magnetic-interference rejection algorithm to help detect and ignore sudden and potentially erroneous magnetic information

ADDITIONAL MAGNETIC CONSIDERATIONS (HARD- AND SOFT-IRON EFFECTS)

A small note about mounting magnetometers on devices with nearby hard-iron effects: Hard-iron effects provide an offset to the magnetometer readings, but they don't change the scale/sensitivity. This takes the sphere of magnetometer readings in all directions and shifts the origin.

The biggest issue with large hard-iron offsets is that you might saturate the supported range of the magnetometer with this offset. Soft-iron effects change the scale/sensitivity and skew. As a result, this can stretch, shrink, or turn the sphere into an ellipsoid *(see figure)*.

If you're new to sensor fusion or IMUs in general, a company such as CEVA can help. It offers a sensor fusion suite that enables companies and tinkerers alike to shortcut their design by using a number of convenient tools and sensor outputs.

CEVA has developed products with algorithms behind them to get the most out of sensors (including those tricky magnetometer situations). Calibrated sensor output, sensor fusion using various sensor combinations (AG, AM, AGM), and magnetic interference rejection help users navigate the IMU world. Additional features include dynamic calibration (adjusting sensor output over temperature and age), gesture recognition, and activity classifiers. ED

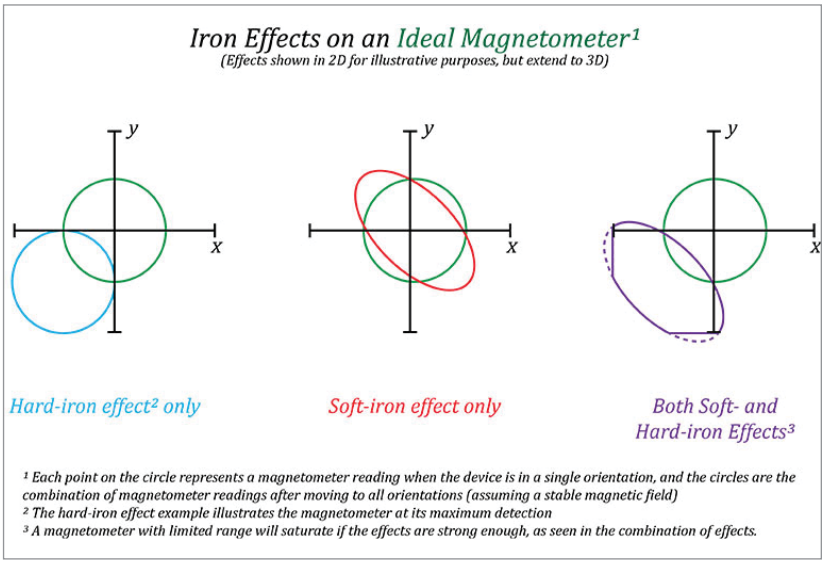

Iron effects differ on an ideal magnetometer.

OVIDIU MESESAN | Senior Engineer, Elma Electronic

11 Myths About Signal Integrity in High-Density Backplanes

With more connections being packed into backplanes, signal integrity becomes a concern. However, misconceptions have arisen around solving this problem. Elma Electronic's Ovidiu Mesesan sets the record straight.

As a signal travels across a network, it focuses only on what it sees in its path. And increases in data rates over the past few decades have made this path a bit more clouded. System infrastructure needs to accommodate such increases through better bandwidth management, better system interoperability, and higher data throughput. This is no short order.

Many of these challenges are being handled in the backplane, which manages much of the interconnect between various circuit boards or cards across a system. By cramming more electrical connections into the same footprint, high-density backplanes seem to invite signal-integrity (SI) issues, making some embedded designers skeptical about how to best move this influx of data within their applications. Here, we explore what can be accomplished using a high-density backplane, while preserving signal integrity across the network.

1. Signal-integrity analysis considerations are different in nature, depending on the backplane architecture/topology/technology.

Signal integrity typically looks at the same parameters of the backplane channels—insertion loss, return loss, operating margin, crosstalk, etc.—irrespective of the building blocks of the channel, how the system is architected, or the intended data rates. The acceptable limits of these parameters are what vary, depending on the protocol/data rate.

Although SI is not a function of form factor, considerations do exist that may be somewhat affected by topology. Solutions to optimize signal integrity require versatility in their approach in order to properly measure different backplane parameters (topology, lengths, geometry), especially with increasing system density.

2. Data provided by the connector and PCB manufacturers have no bearing on the results of pre- and post-layout signal-integrity analysis of channels.

It's essential that the data provided by both connector manufacturers and PCB manufacturers first be vetted for, and then incorporated into, the analysis performed. For instance, this data is critical to the optimization and analysis of the PCB stackup, trace geometry parameters, as well as physical and electrical models for the connectors *(see figure)*.

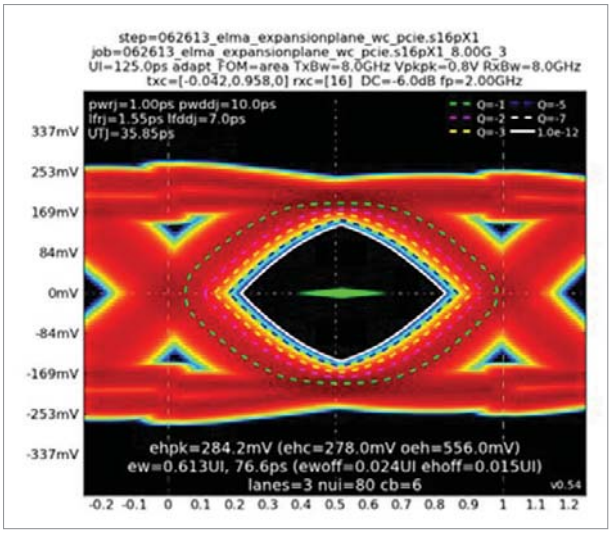

Channel performance depends on minimizing crosstalk across the connector.

3. Links in a backplane-based system are independent of one another, and therefore don't affect overall performance.

Each link in these types of systems will contribute to system performance and potentially affect the performance of otherwise good links, causing bottlenecks in bandwidth across the entire system via excessive crosstalk. Provisions to simulate and measure all problematic (worst case) links within a design need to be incorporated into a system-level signal-integrity analysis.

4. The behavior of a complete channel is the sum of the behaviors of its individual sub-sections.

At higher data rates, the behavior of a complete channel isn't always accurately represented by just a concatenation of indi-

vidually simulated/measured blocks. Or in the words of Aristotle, "The totality is not, as it were, a mere heap, but the whole is something besides the parts." Cascading the models of individual blocks is warranted, only if the points at which the cascading is done are those with very low reflections, which in most cases (like connector interfaces) is patently false.

That's why co-design, or cooperation at the layout stage, is an important aspect of an embedded system. Developing a system in parallel can reduce signal-integrity issues, yielding better return loss and, more importantly, much better insertion loss.

5. The type of weave used on the layers that carry high-speed signals has no bearing on signal integrity and, therefore, isn't a critical consideration.

The weave does impact signal integrity as the data rate increases. With the advent of protocols that require lower intra- and inter-pair skew, fiber-weave-induced skew may make or break the performance of some of the longer links in a system. Careful consideration must be given to the compromise between a spread weave choice and the resin content for a given laminate material.

The long-held assumption that the dielectric material surrounding the traces in the backplane or daughtercard PCBs is almost homogenous and isotropic in all directions is not only demonstrably false, but dangerous to make for high-data-rate transmission channels.

6. You can expect simulations to be close to measurements, even without validating your models.

Simulation at the pre- and post-layout stage is a necessary step. However, without using validated models, your simulations will only be as good as the assumptions built into those models. Vetting S-parameter models for passivity and causality, and verifying the validity of a particular laminate material's properties, are just two examples.

7. The values of copper foil roughness and resistivity given by PCB laminate manufacturers, and even those in empirical formulas, are pretty close to reality.

Apart from the values given in manufacturers' datasheets, typical models used in SI simulations for copper surface roughness include the Huray snowball model and Hammerstad. However, these datasheet values and theoretical models must be correlated with actual measurements of the copper foil as processed by the PCB manufacturer. Different PCB fab houses employ different methods and equipment for roughening the copper foil's surface for better adherence to the substrate. How well the theoretical models used in simulations reflect reality is something worth investigating, because it affects the simulation's accuracy *(see Myth 6)*.

8. Developing design rules to test signal integrity should be conducted up to the limit of the application.

To truly test for signal integrity, it's important to go beyond the scope of the application and account for worst-case scenarios to ensure your system, and your signals, will hold up under all circumstances. This "margin" allows for spikes in data transfer, higher than normal system loads, and less optimal designs of other components in the system (plug-in modules, mezzanine cards, power supplies, etc.).

9***. The mated interface to the backplane isn't important when considering the performance of a backplane PCB.***

To properly model the signal path, the mated connector interface absolutely must be considered. The signal budget is affected by the plug-in cards that mate to the backplane, the backplane itself, and the connectors used as interfaces between the plug-in cards and backplane. Moreover, it's not only the electromechanical properties of these mated connector interfaces that affect the signal budget, it's also the footprint used for these connectors both on the plug-in cards and on the backplane (things like pad and anti-pad size, trace routing geometry in the area where it connects to the footprints, etc.)

10. Any simulation software and any de-embedding approach would do a decent job at analyzing SI data and making accurate predictions (i.e., "all SI tools are created equal and it only depends on the user's ability of how accurately they make predictions about interconnect performance").

While one would be tempted to consider simpler and oftentimes cheaper solutions for SI, one must be aware that—as is the case with everything else in life—you get what you pay for. "Half-baked" solutions may get it right sometimes (even a broken clock is right twice a day), but investing diligently in proven simulation tools, test equipment, calibration, and de-embedding techniques goes a long way to achieving good correlation between simulations and measurements, which is key to a successful approach to SI.

11. At the end of the day, if my backplane meets the signal budget allocated to it, even within the slightest margin, everything will run smoothly.

As was mentioned in previous myths, one must approach even individual sub-component design with a holistic methodology. A poor (or marginally good) launch from a plug-in module into a backplane will only get worse. With the backplane being a passive element of an end-to-end channel, it can't improve the signal quality because it will add loss (different kinds of loss, too). Similarly, a backplane design with poor margins may actually result in a system failing. ed

KYLAN ALLISON | Process Engineer, Rohinni

Micro and Mini LEDs Lead to Design Freedom

Mini and micro LED technology is poised to take over as the next generation of lighting technology in a multitude of commercial applications.

Design freedom is the ability to create products without restriction, or the ability to bring future products to reality. Revolutionary products including lifelike displays, extremely thin consumer electronic devices, new lighting in automotive applications, and the presence of light in novel locations are made possible with the introduction of mini and micro LED technology.

Due to significant process advances regarding placement of mini and micro LEDs, product applications ranging from display backlights, computer keyboards, direct emission displays, automotive lighting, and consumer electronics can utilize light in new ways. Mini and micro LEDs eliminate countless design challenges and pave the way for unconstrained product design.

Thanks to the microscopic size of mini and micro LEDs, the lights can be placed virtually anywhere to fulfill design requirements. In terms of size, mini LEDs and micro LEDs are significantly smaller and thinner than their traditional LED counterparts. Mini LED components are between 50 and 300 µm, while micro LEDs are components smaller than 50 µm. In comparison, the smallest traditional, packaged LED is 650 × 350 µm.

With these shrunken, unpackaged versions of the traditional LED, it's now possible to create considerably thinner, downsized, and arrayed mini LED lighting technology. Any application that uses traditional packaged LEDs can utilize the smaller micro and mini LEDs, and novel applications are made possible due to the smaller components. Because mini and micro LEDs are a fraction of the size of traditional LEDs and involve less material, they're a fraction of the cost of the larger components and can be utilized in new ways.

SUBSTRATES

Design freedom is incorporated in many aspects of the construction of micro and mini LED applications, including the selection of a preferred substrate. Micro and mini LEDs can be placed on flexible, glass, and rigid circuits among other industry-available substrates.

Figure 1 illustrates the circuit flexibility for an application using mini LEDs. As a result, they bring flexibility to mechanical design as well as optical designs—LEDs can be placed exactly where intended while minimizing or eliminating light guides and other optics. This allows micro LEDs to be placed in areas where traditional LEDs are problematic.

1. Pictured is a 1-inch circuit with 1,300 mini LEDs placed on PET silver epoxy.

PLACEMENT

While many companies have created their own product demos using mini LED technology, most products have failed to enter mass production and are far from reaching the average consumer. For engineers to truly achieve design freedom with the new lighting technology, the primary technological roadblock that needs to be addressed is the accurate and fast placement of mini and micro LEDs. The placement solution that meets industry post-process yields and industry-demanded speeds allows for the mass production and broad adoption of micro and mini LED technology.

Today, packaged LEDs are placed using standard surface-mount technology (SMT) equipment, which takes much longer than an alternative solution to place mini and micro LEDs. A simpler solution that places mini and micro LEDs faster and accurately would be more cost-efficient due to the less time it takes to place the components.

Placement of mini and micro LEDs presents a major advantage for product designers because the components can be put virtually anywhere. In a product application, wherever the light is desired, the microscopic die can be placed and illuminated. For example, in a keyboard backlight, with mini or micro LEDs, a light guide plate is eliminated because the die can be placed directly behind each key.

Thinner products and more uniform light distribution are possible because of minimal diffusion or a light guide stack.

Mini and micro LEDs simplify the optical design requirements as the die are placed exactly where the light should go.

Mini and micro LED technology also offers several advantages in the process of manufacturing displays with the ability to place the mini or micro LEDs at the desired pitch, select LED components of the preferred size, and the ability to bin the components at specified wavelengths. Binning the small-scale LEDs at a fixed wavelength leads to optimal product performance in the matter of luminosity, uniformity, and color temperature. A variety of LED die size configurations is a key contributor to achieving the desired product design and performance.

DISPLAY ADVANTAGES

Display applications of all sizes boast of significant advantages when using micro or mini LED technology. Mini and micro LED backlights have individual die control. Even with the thousands of die in a backlight, each individual die can be controlled, offering greater resolution lighting and brightness control that directly leads to greater screen resolution for product applications.

Local dimming is a feature of LED displays that dims the backlight in different regions. Controlling each region individually improves display performance by enhancing the contrast ratio and making dark elements of an image appear darker.

Since mini and micro LED displays have a massive number of smaller-sized components involved in comparison to conventional LED displays, the technology brings a huge advantage to the fields of contrast, brightness, and black levels. Zones of the screen can now appear completely black, or "true black." Merging mini and micro LED technology with quantum dot abilities will enhance the brightness and color, leading to an overall better picture quality, thus surpassing the standards set by OLED technology.

For enhanced image quality, and to completely utilize the full capabilities of high dynamic range (HDR), local dimming is necessary. Televisions will be capable of dimming the mini and micro LEDs behind dark areas of an image to generate a contrast ratio of higher quality between darker and lighter regions of a picture. The advancement of mini and micro LED technology allows for a massive amount of controllable dimming zones; for example, a mini LED television may have approximately one thousand zones or more.

Micro LEDs enable higher image quality in direct emission displays because of the small components. Each pixel on a direct emission display is constructed of three subpixels—each composed of a red, green, and blue micro LED.

The narrow-pixel-pitch LED direct emission displays offer high luminous efficiency, reliability, brightness, contrast, and an especially fast response time. Coupled with other micro components, micro LED technology opens the door to greater design freedom in display applications.

In between each pixel on a direct emission display is a light dome to separate the pixels from bleeding into the adjacent pixel. That dome occupies space that can also be used to place an integrated circuit (IC) or other micro components like sensors. This option enables a simpler design package and a thinner product because components can be distributed to make a thinner stackup.

Soon, consumers will see indoor signage or electronic billboards as one of the first products to adopt the use of sensors coupled with micro LED technology. A sensor can detect when someone is standing near the display and is able to adjust the brightness or image. The overall design freedom enabled by coupling micro components with micro LED technology is creating smarter products, which can save power, create thinner products, and give designers more control.

Not only do display applications benefit from microscopic LEDs, but due to these characteristics, mini and micro LED technology also gives designers and engineers a license to create products as they imagine in other product categories. Other applications for this technology include input devices, computer keyboard backlights, display backlights, and automotive applications. *Figure 2* illustrates "LightThread," a Rohinni design that showcases the microscopic size of the LEDs. LightThread is less than 1 mm wide and can be used in applications ranging from clothing to general lighting.

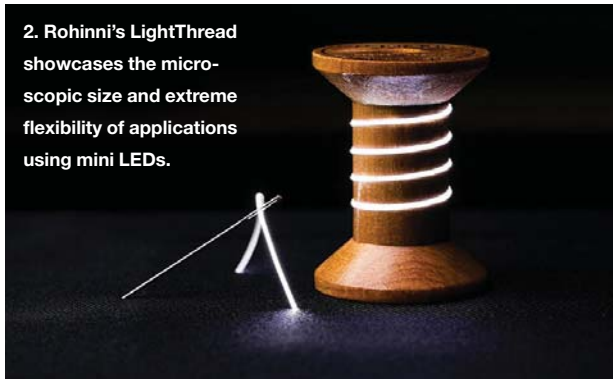

2. Rohinni's LightThread showcases the microscopic size and extreme flexibility of applications using mini LEDs.

POWER EFFICIENCY

Due to the sheer number of mini or micro LED die in the same area covered by traditional LEDs in an application, power driven to the components is more distributed. And the distributed drive architecture leads to more efficient lights.

Multiple micro or mini LED die in place of a single packaged LED die has better thermal characteristics because the LEDs aren't being driven as hard. Mini and micro LEDs can be driven at their maximum efficiency curve, instead of their maximum drive capability like traditional LED applications. In addition, the simplified optical stack helps minimize light loss since the light needn't be redirected or diffused—the light is going where it was intended and takes less power to do so.

DEBENDRA DAS SHARMA | PCI-SIG Board Member and Intel Fellow, Intel Corp.

What's the Difference Going from PCIe 3.0 to PCIe 6.0?

To meet rising demands for improved speed, cost, and power interconnectivity, the PCI-SIG continues to evolve the venerable PCIe architecture, which is looking at 64 GT/s for its next release.

The Peripheral Component Interconnect (PCI) architecture has provided I/O connectivity for computing, communication, and storage platforms for more than three decades. From its inception as a local bus interface for all types of I/O devices in the PC industry, it has evolved as a point-to-point link-based interface (PCI Express) to satisfy the I/O requirements across the cloud, enterprise, artificial intelligence, PC, embedded, IoT, automotive, and mobile market segments.

This is made possible due to the ability of PCI Express (PCIe) architecture to seamlessly deliver cost-effective, HVM (high-volume manufacturing) friendly, power-efficient, high-bandwidth, and low-latency solutions through six generations of technology evolution, doubling the data rate in every generation *(Fig. 1)* while maintaining full backwards compatibility with all prior generations to protect customer investments.

One salient and unique feature of the PCIe specification is that while it supports multiple data rates and multiple widths to back the different performance needs of different devices across a wide range of usage models, they interoperate with each other. This enables both silicon and platform developers to design and validate to one specification. Even though a multitude of form factors (e.g., M.2, U.2, CEM, various flavors of SFF) have evolved to meet the needs of diverse systems across the compute continuum, they all use the same silicon ingredients based on a common PCIe base specification.

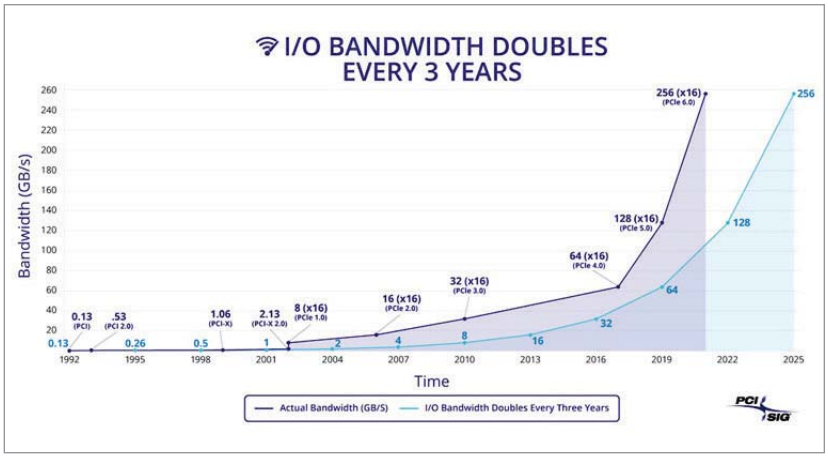

1. The PCI Express roadmap, demonstrating the doubling per-pin bandwidth every generation.

The success of PCI technology as a ubiquitous I/O interconnect is due to it being an open industry standard, backed by a robust compliance program to ensure seamless interoperability between devices from different companies. PCI-SIG, a consortium of more than 800 member companies, spread across the globe, owns and manages PCI specifications and runs the compliance program. PCI-SIG expects PCIe technology to continue to evolve to meet the diverse I/O needs across the entire compute continuum for many more years to come.

This article delves into the details of PCIe technology covering its evolution from the fourth through the sixth generation, as a follow-up to our prior article in *Electronic Design* covering the first three generations.

THE FIRST THREE PCIe GENERATIONS AT 2.5, 5.0, AND 8.0 GT/s

PCIe technology started off in 2003 at a 2.5-GT/s data rate, supporting widths of x1, x2, x4, x8, and x16 for different bandwidth levels. The supported widths haven't changed through the six generations of evolution of PCIe architecture. PCIe 2.0 specification doubled the data rate to 5.0 GT/s in 2006. The first two generations of PCIe technology used 8b/10b encoding, incurring a 25% encoding overhead. This was needed to establish dc balance and for the additional encodings required for physical-layer packetization (e.g., to indicate start and end of different types of packets) and training handshake.

In terms of the PCIe 3.0 specification, a strategic decision was made to increase the data rate to 8.0 GT/s instead

FREE WEBINAR

August 13, 2020 | 2:00pm EDT

Registration URL: https://www.electronicdesign.com/resources/webcasts/webinar/21134595/simulation-of-heat-exchangers

Simulation of Heat Exchangers

Heat exchangers are used to transfer thermal energy, usually between two fluids. Several types of heat exchangers are available, and each type of heat exchanger can be sized to suit the desired application.

In this webinar, you will see how to simulate fluid flow and heat transfer in heat exchangers.

We will study how simulation can be used to evaluate and improve several properties, such as the overall heat transfer coefficient, pressure losses, total power exchange, compactness, and more.

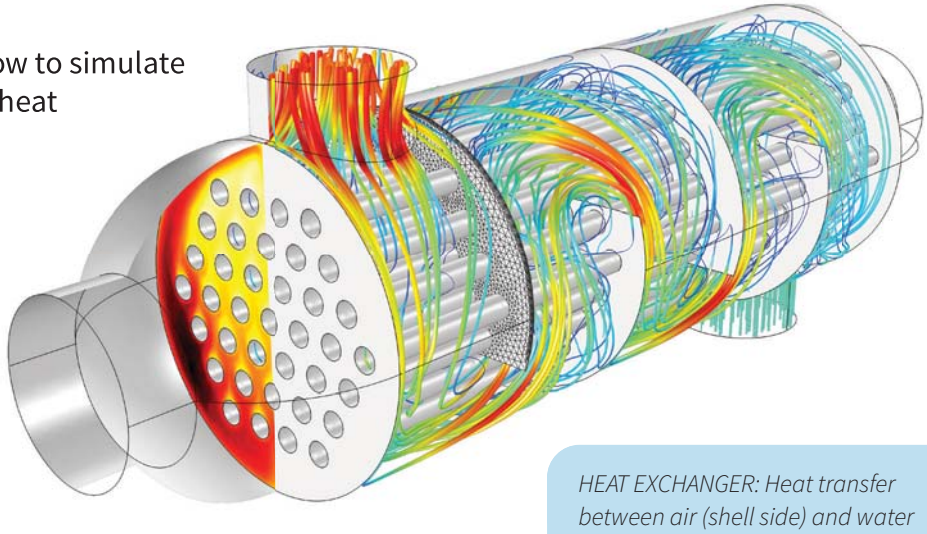

HEAT EXCHANGER: Heat transfer between air (shell side) and water (tube side) through steel walls in a shell-and-tube heat exchanger. Nonisothermal flow simulation results reveal flow velocity, temperature distribution, and pressure within the fluids.

You will also get a live demonstration of the model configuration using COMSOL Multiphysics®. Finally, specialized simulation applications will be used to compare the performance of heat exchangers as a function of, among other things, geometrical properties and the choice of a counterflow or parallel flow configuration.

SPEAKER: Nicolas Huc, Product Manager, Heat Transfer module, COMSOL

Nicolas Huc joined COMSOL France in 2004 and is currently the head of their development team. He is also the manager of the Heat Transfer Module. Nicolas studied engineering at ENSIMAG before receiving his PhD in living system modeling from Joseph Fourier University.

Hosted with ElectronicDesign

COMSOL

of 10.0 GT/s. This was augmented with a new 128b/130b encoding mechanism to double the bandwidth per pin over PCIe 2.0 (1.6 due to data-rate increase × 1.23 with the encoding overhead reduction). The new encoding mechanism ensured high reliability by adopting a fault model of three random bit flip detections and had several innovative approaches to perform physical-layer framing of packets while preserving the packet format delivered from the upper layers (Link and Transaction).

The decision not to double the data rate (to 10.0 GT/s) was made with data from extensive analysis to ensure that PCIe 3.0 can work on existing channels with the projected silicon and platform ingredient capabilities of the time of the specification's introduction (2010) and still meet projected power and cost constraints. Backchannel equalization was introduced in PCIe 3.0 specification to mitigate the effects of channel loss on the platform by silicon ingredients. The encoding and equalization mechanisms of PCIe 3.0 architecture was robust enough to sustain subsequent generations of speed increases.

In addition to greater speed, PCIe's evolution during this time supported features such as I/O virtualization and device sharing to accommodate the emerging trend of supporting multiple independent virtual machines and containers in a single platform. To support the performance needs of accelerators, it added features such as caching hints, atomics, and lower-latency accesses through enhanced transaction bypass semantics.

To help with the evolution of handheld segments such as smartphones and tablets, we enhanced our low-power states with deeper low-power states. This enabled devices to maintain their states for a quicker resumption of traffic while consuming power in the single digits of microwatts. This mechanism for lower power consumption during idle state along with the industry-leading power efficiency during active state (around 5 pJ/bit), plus a fast transition time between the two (1 to 100 µs), results in PCIe architecture being the interconnect of choice across low-power as well as high-performance segments.

OCTOBER 2017: PCIe 4.0 SPECIFICATION AT 16.0 GT/s

The journey to double the data rate from 8.0 GT/s to 16.0 GT/s took longer to ensure that platform and silicon ingredients could evolve in a cost-effective and power-efficient manner for a seamless technology transition. The channel loss budget was increased to 28 dB. The routing material continued to get better with newer and economically feasible materials such as Megtron-2 (4, and 6) with improved loss characteristics, along with improvements in packaging technology, making it feasible within the cost and power constraints of platforms with hundreds of lanes.

Even then, though, it wasn't enough to cover longer channel lengths such as 20 in. with two connectors. With board materials used in systems and available packaging technology, we can support about 15 in. of board trace with one connector and an add-in card. Retimers had to be formally specified as channel-extension devices. Retimers have the full physical layer and double the channel loss. Up to two retimers are allowed in a link, enabling longer-reach channels with PCIe architecture *(Fig. 2)*.

PCIe 4.0 technology allowed more outstanding transactions to sustain the ever-increasing bandwidth capabilities through scaled flow-control credit mechanisms and extended tags. It also enhanced the reliability, availability, and serviceability (RAS) features to enable migration to direct-connected PCIe storage devices through downstream port containment. Furthermore, systems can run non-destructive lane-margining capabilities without interrupting the system operation. As always, these enhancements will sustain a few generations of bandwidth increases.

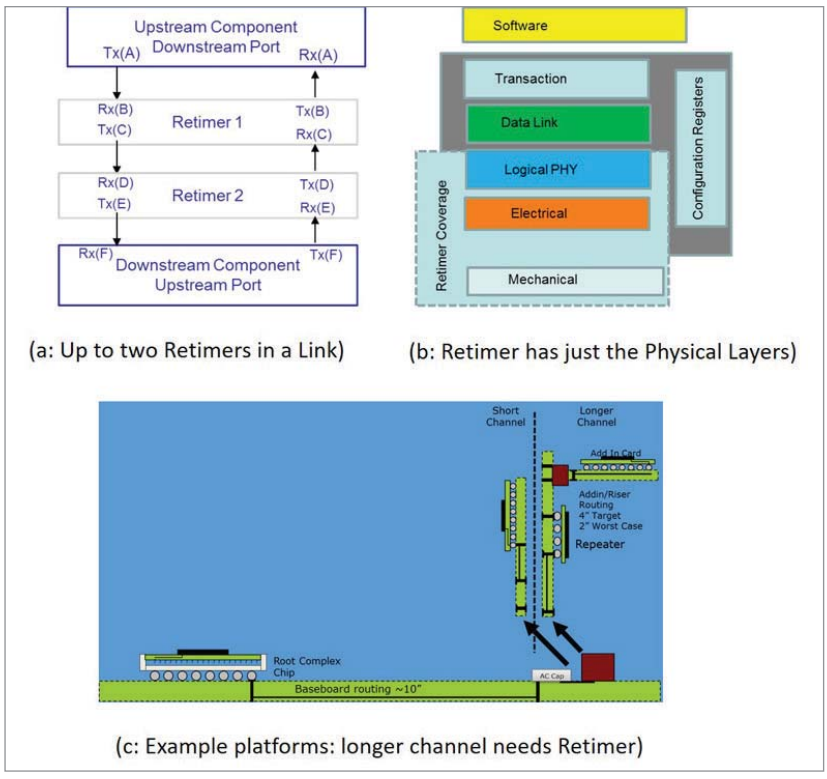

2. Retimers are used to extend the channel reach.

A question naturally arises, though: How did the systems deal with the I/O bandwidth needs when the technology transition from PCIe 3.0 to PCIe 4.0 specification took longer than usual due to the inflection points that had to be addressed? The answer is nuanced.

Platforms that debuted with PCIe 3.0 architecture had about 40 lanes of PCIe technology coming out of every CPU socket. Just prior to the transition to PCIe 4.0 specification, the lane count per CPU socket increased dramatically, reaching up to 128 lanes per CPU socket in some platforms. Thus, while the per-slot (e.g., x16) bandwidth didn't increase, the aggregate I/O bandwidth increased 3X in platforms, both in terms of the number of lanes as well as measured I/O bandwidth.

Storage tends to be an aggregate bandwidth driver, with each storage device connected to the system using a narrow link (e.g., x2 or x4). Thus, increasing storage needs were met with lane count increase. Networking is a single-slot usage; during this time, it transitioned from 10 Gb/s to 100 Gb/s and dual 100-Gb/s network interface cards (NICs). This bandwidth was accommodated by NICs moving from x4 to x16 (and 2 x16 for dual 100-Gb/s NICs) width—a great tradeoff from a power/cost/performance perspective. Accelerators and GPGPUs also evolved to make efficient data moves through proper transaction sizing and protocol hints.

Overall, the slowdown in speed evolution from PCIe 3.0 to PCIe 4.0 architecture was mitigated by a width increase, thanks to the flexibility offered by PCIe specification. The ecosystem evolved naturally, so the speed transition was accomplished in a cost-effective and power-efficient manner.

MAY 2019: PCIe 5.0 SPECIFICATION AT 32.0 GT/s

The past few years have seen a significant change in the computing landscape as cloud computing, edge computing, and applications such as artificial intelligence, machine learning, and analytics have led the demand for faster processing and movement of data. As the compute and memory capability increases at an exponential pace, we need to sustain I/O bandwidth doubling at an accelerated cadence to keep up with the performance of emerging applications.

For example, 400 Gb (or dual 200 Gb) networking needs a x16 PCIe at 32.0 GT/s to sustain the bandwidth. This required the release of a fully backwards compatible PCIe 5.0 in less than two years after PCIe 4.0 architecture—a significant achievement for a standard.

The evolution from PCIe 4.0 to PCIe 5.0 specification was primarily a speed upgrade. The 128b/130b encoding, which was the protocol support to scale bandwidth to higher data rates, was already built in with PCIe 3.0 and PCIe 4.0 specification. The channel loss was extended to 36 dB along with enhancements to the connector to minimize the loss with the increased frequency range. With the improvement in board material and packaging technologies, the channel reach is like PCIe 4.0 technology, using retimers to extend the channel reach.

One of the enhancements made with PCIe 5.0 architecture is the built-in support for alternate protocols. As PCIe technology has evolved to be the highest bandwidth, most power-efficient, and most widely deployed interface, some usages require additional protocols, such as coherency and memory to run on the same pins as PCIe architecture. For example, certain accelerators and smart NICs may cache the system memory and map their memory to system memory space for efficient data exchange and atomics in addition to PCIe protocols. Similarly, system memory is migrating to PCIe PHY due to the power-efficient high bandwidth and low latency solution it offers.

Other protocols are employed, too, such as symmetric cache coherency between components using PCIe PHY. Support for alternate protocols on PCIe PHY is provided to meet these user requirements, preventing the fragmentation of the ecosystem with different PHY for different usages.

TARGETED FOR 2021 RELEASE: PCIe 6.0 SPECIFICATION AT 64 GT/s

We continue the accelerated journey to double the bandwidth again in two years in a backwards-compatible manner. Applications such as AI, machine learning, gaming, visual computing, storage, and networking continue to demand bandwidth increases as we find ourselves in a virtuous cycle of more throughput driving new applications with greater capabilities.

Devices such as GPUs, accelerators, high-end networking (800 Gb/s), coherent interconnects, and memory expanders continue to demand more bandwidth at an accelerated pace. Constricted form-factor applications that can't boost width also demand increased frequency to deliver performance.

There are significant challenges to doubling the data rate beyond 32.0 GT/s using the NRZ (non-return-to-zero; i.e., the binary signaling where a 0 or a 1 is transmitted in every unit interval, or UI) method due to the channel loss. Thus, PCIe 6.0 will adopt PAM4 (pulse amplitude modulation, 4-level) signaling *(Fig. 3)*, which has been widely adopted by networking standards when they moved to data rates of 56 Gb/s and beyond. Using this method, two bits are encoded using four levels in the same UI, allowing PCIe 6.0 UI (and Nyquist frequency) to be identical to PCIe 5.0 architecture.

While PAM4 alleviates the channel loss due to running at half the frequency with two bits per UI, it's more susceptible to errors due to various noise sources caused by reduced voltage (and timing) ranges. This manifests as a higher bit error rate (BER), several orders of magnitude higher than the 10^{-12} BER for PCIe 1.0 through PCIe 5.0 specifications.

Another side effect is the correlation of errors due to correlated error sources such as power-supply noise as well as error propagation in the same lane due to the decision feedback equalizer (DFE). These effects are mitigated by deploying a forward-error-correction (FEC) mechanism, which has the drawback of reduced link efficiency caused by the FEC bits as well as the latency add for the encoding/decoding mechanisms. The stronger the FEC, the worse the performance characteristics. However, the effective bit error rate improves due to correction.

For example, some of the existing standards have a 11% bandwidth loss and a FEC latency of greater than 100 ns, which don't meet the bandwidth and latency demands of a load-store interconnect like PCIe technology. PCIe 6.0 specification development is following the guardrails in terms of key metrics delineated in the *table*. While these are challenging goals and haven't been solved before, we must meet these metrics to ensure that PCIe continues to be a high-performance interconnect.

With FEC, we need a fixed FLIT (flow control unit) size to apply the correction. With a fixed FLIT size, it's advantageous to have the error-detection mechanism (cyclic redundancy check, or CRC) operate on the FLIT. Since PCIe defines the data-link-layer packet (DLLP) and transaction-layer packet (TLP) of variable sizes, we define the payload to align to FLITs. Thus, a FLIT can have multiple DLLPs and TLPs, and a TLP/DLLP may span across multiple FLITs.

In this new mode, since the FLIT includes the CRC, the DLLP and TLP will not carry their individual CRC bytes as they did in prior generations. Furthermore, since FLITs are fixed size, there's no need to have a PHY layer framing token (4 bytes) for every TLP or DLLP. These savings help improve efficiency to overcome the FEC overhead.

The packet efficiency with PCIe 6.0 architecture exceeds that of prior generations for payloads up to 512 bytes. For example, a 4DW (Double Word, each double word is 4 bytes) request TLP will

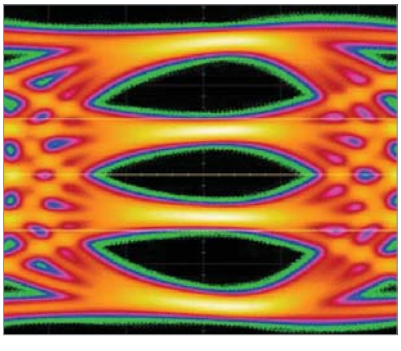

3. Eye diagram with PAM4 encoding: The three eyes in the same UI represent the four possible voltage levels, representing 2 bits.

METRICS FOR PCIe 6.0 DEVELOPMENT

Metrics	Requirements
Data rate	64 GT/s, PAM4 (double the bandwidth per pin every generation)
Latency	<10 ns adder for transmitter + receiver over 32.0 GT/s (including FEC)
Bandwidth inefficiency	<2 % adder over PCIe 5.0 across all payload sizes
Reliability	0 < FIT << 1 for a x16 (FIT = failure in time, number of failures in 10^9 hours)
Channel reach	Similar to PCIe 5.0 under similar setup for retimer(s) (maximum 2)
Power efficiency	Better than PCIe 5.0
Low power	Similar entry/exit latency for L1 low-power state Addition of a new power state (L0p) to support scalable power consumption with bandwidth usage without interrupting traffic
Plug and play	Fully backwards compatible with PCIe 1.x through PCIe 5.0
Others	HVM-ready, cost-effective, scalable to hundreds of lanes in a platform

have a TLP efficiency of 0.92 with FLIT-based encoding versus 0.62 in prior generations (with 128b/130b encoding and a 5% DLLP overhead). This results in an ~3X improvement in effective throughput (2X from data rate increase and ~1.5X improvement in TLP efficiency). As the TLP size increases, efficiency drops; for the 4-kB data payload size, it reduces to 0.98, in line with the bandwidth inefficiency provided in the metrics provided in the table. An example FLIT layout is shown in *Figure 4.*

Figure 5 demonstrates the tradeoffs associated with the raw burst error rate on a wire (error propagation to multiple bits is considered as one error) and the effectiveness of different FEC to handle that error. A single symbol error-correcting code (ECC) corrects one error burst, whereas a double symbol ECC corrects up to two error bursts.

While the length of the burst follows a certain probability distribution function, the ECC code is defined such that the probability of a burst exceeding the ECC capability is negligible. Simulations and silicon data have been used to make the tradeoff between the error rate, the nature of the burst, the channel constraints, and the silicon capability. PCIe 6.0 is targeting a burst error probability of 10^{-6}, which will result in a retry probability of the FLIT of around 10^{-6}.

Since PCIe has a low-latency link-level retry mechanism, we don't need to deploy a strong FEC, which will increase the latency and bandwidth overhead. A retry probability of 10^{-6} (or even 10^{-5}) is a reasonable tradeoff, resulting in the FEC latency adder of 1-2 ns in each direction. In the case of a retry (a 10^{-6} probability event), the FLIT

Table 6-1 Flit Layout in a x16 Link

Description		Lane															
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
128 UI	TLP Bytes [0… 223]	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
		16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
		32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
		48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
		64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
		80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
		96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
		112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
		128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
		144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
		160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
		176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
		192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
		208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
	TLP, DLP	224	225	226	227	228	229	230	231	232	233	234	235	DLP 0	DLP 1	DLP 2	DLP 3
	DLP, CRC, ECC	DLP 4	DLP 5	CRC 0	CRC 1	CRC 2	CRC 3	CRC 4	CRC 5	CRC 6	CRC 7	ECC 0A	ECC 0B	ECC 0C	ECC 1A	ECC 1B	ECC 1C

4. FLIT layout in a x16 Link: The first 236 bytes (0..235) are for TLP(s), the next 6 bytes are for data-link-layer payload (dlp0..5), the next 8 bytes for CRC (crc0..7) and the last 6 bytes are for ECC (ecc 0..1). The FEC is a 3-way interleaved ECC, each capable of correcting a single byte, with the interleaving shown in three colors.

is delayed by ~100 ns due to the round-trip retry mechanism. This is a reasonable tradeoff versus adding 100+ ns to every FLIT with a strong FEC and paying a high bandwidth penalty. Details of the PCIe 6.0 specification are available to members from the PCI-SIG website.

CONCLUSION

With a rich and successful history of navigating several technology transitions in a backwards-compatible manner spanning three decades, PCI-SIG is well-positioned to continue leading the changing computing landscape going forward. The power and promise of this open standards organization, backed by the combined innovation capability of 800+ member companies, makes our technology nimble, scalable, cost-effective, power-efficient, leading-edge, and multi-generational, with relevance across all market segments and usage models for the foreseeable future.

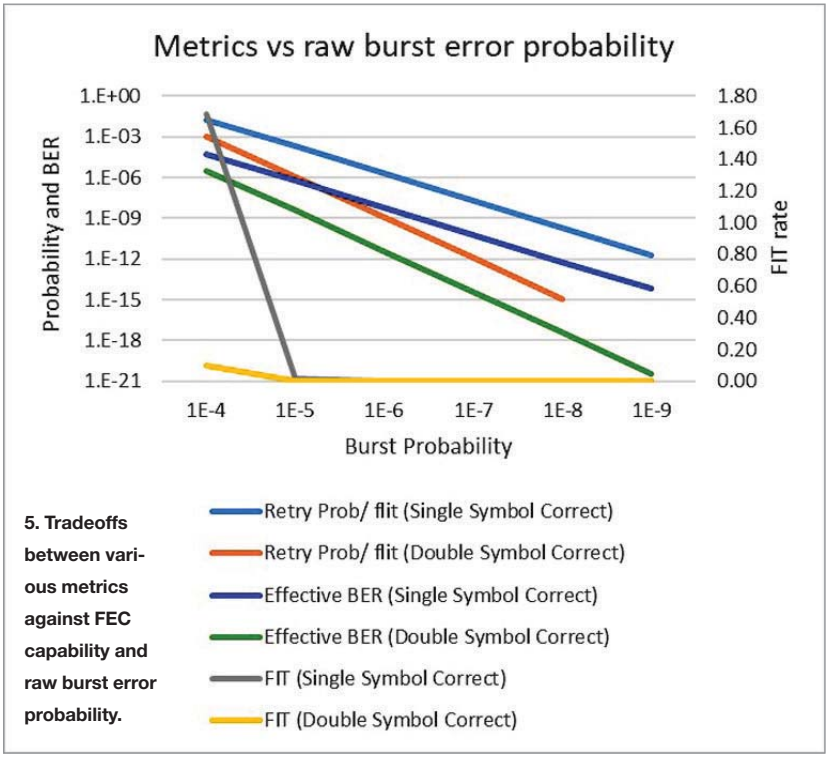

5. Tradeoffs between various metrics against FEC capability and raw burst error probability.

AKSHAY SARUP | Product Engineer, CXL and PCI
MARK OLEN | Product Marketing Group Manager, ICVS Division, Mentor, a Siemens Business

Validate Assertions in Packet-Based Protocol Designs Using UVM Callbacks

UVM callbacks help save time when multiple assertions must be validated in PCIe and other packet-based protocol designs.

Assertions bring immediate benefits to the entire design and verification cycle. To use assertions effectively in the verification cycle, they need to be exercised for checking legal design behavior and must fire when illegal behavior is encountered in the design.

When a large number of assertions need to be validated for firing under illegal circumstances, a callback mechanism is particularly effective for creating these scenarios. That's because they require minimal updates to the testbench environment and minimize the effort to create complex scenarios for validating the assertions.

By definition, a callback mechanism, or function, is a "call-after" function that alters the intent of the original function call. Callbacks are a common mechanism used in functional verification to modify the original content of a sequence item to induce a desired scenario. This enables more dynamic and fine-grained control in selecting which particular sequence item gets corrupted and, in a test, how many such corruptions occur in order to stress test the design. Since callbacks allow for the easy creation of nuanced and complex stimulus creation, they're instrumental to assertion verification.

This article describes how callbacks implemented in verification IP can be used for assertion validation in designs using PCIe and other packet-based protocols.

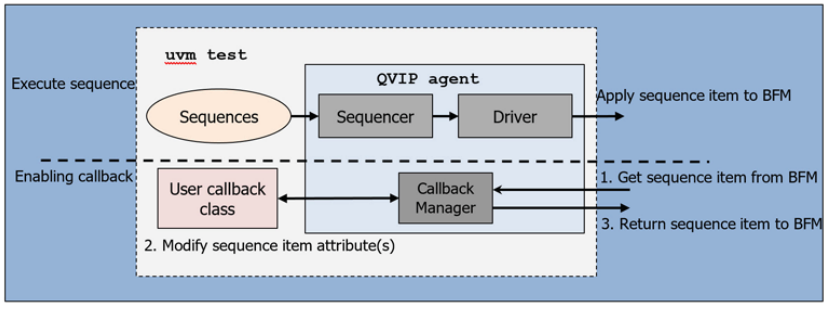

1. The basic sequence of events that take place when callbacks are enabled in an agent.

Assertion verification is usually an integral part of the verification IP development cycle *(Fig. 1)*. The first step involves coding the assertions. The second step validates the assertions by creating scenarios that fully exercise the assertions to ensure they don't fire under intended design behavior and do fire under erroneous scenarios.

Callbacks can be very beneficial in the scenario-generation step since the verification engineer doesn't have to write additional tests. Instead, they enable callbacks to modify the original stimulus to create interesting scenarios. This is achieved by simply extending the callback class to override the virtual method:

```
do_callback(...)
```

that gets called when a sequence item is executed by the bus functional model (BFM). This is particularly useful in packet-based protocols, such as PCIe, where packet fields need to be corrupted and callbacks provide fine-grained control to do so.

The validation engineer's job is to make sure that the assertions are validated in every relevant scenario. Callbacks can be of extreme importance when assertion validation is carried out by manipulating the fields of a sequence item or by initiating a sequence item in a state in which it's not allowed. Callbacks speed up assertion validation by providing the entire structure for such manipulation in a UVM testbench, with minimal updates and no interaction with the existing sequences responsible for generating the sequence items.

For a layered protocol like PCIe, where communication between the transmitter and receiver elements takes place via structures known as packets, the use of callbacks for assertion validation significantly improves efficacy. Using verification IP, the packet information can be modified via callbacks and all packet-based assertions can be easily verified by making use of this mechanism. Now let's look at a scenario from each of the transaction, data link, and physical-layer packets.

CASE 1: TRANSACTION-LAYER-BASED CALLBACKS

The high-level transactions occurring at the device core are known as transaction layer packets (TLPs). Verification IP, such as Questa Verification IP, allows all fields of a TLP (both request and completion TLPs) to be altered. If users want to validate an assertion related to the fields of a TLP, they can simply do so by corrupting the TLP fields via a callback.

For example, the PCIe protocol states that for a request of "length = 1 DW," the value of the Last Byte Enable field should be zero; where Last Byte Enable is the Byte Enable value for the last DWORD of the request. In simulation, if users want to inject this error to every Memory Write packet of length 1 DW, then they may do so as illustrated in *Figure 2.*

CASE 2: DATA-LINK-LAYER-BASED CALLBACKS

Data link layer packets (DLLPs) are used for a variety of purposes, such as ensuring the integrity of TLPs, flow control, and power management. Just as for TLPs, callbacks can be used to inject errors in DLLPs. For instance, lcrc is used to check the data integrity of TLPs and DLLPs. lcrc is appended in the TLP at the data link layer (DLL). If the value of the lcrc attached with a packet isn't the same as the calculated value, then it's a protocol violation. This incorrect behavior, or assertion firing, can be verified *(Fig. 3).*

CASE 3: PHYSICAL-LAYER-BASED CALLBACKS

Perhaps the most efficient use of callbacks comes from modifying ordered set fields. In the cases of DLLPs and TLPs, a packet can still be executed via a sequence once linkup is achieved. But using a sequence to inject an error into an ordered set before linkup can be more convoluted and more susceptible to errors because the ordered set rules change with each LTSSM (link training and status state machine) state.

Callbacks, on the other hand, allow the user to inject an error in a much more controlled way. For instance, if a control SKP ordered set is to be sent in place of a TS2 OS, one can easily use callbacks. Hence, this invalid protocol scenario can be easily validated via callbacks *(Fig. 4).*

As can be seen, when a large number of assertions are to be validated, callbacks save time by making sure the engineer doesn't have to code a new sequence for each scenario. Callbacks allow for the creation of nuanced and complex stimulus with ease and, thus, are instrumental for assertion verification.

```
function void tlp_ldwe_err_cb::do_callback(tlp_t tlp);
 req_cb_s req_st;
 pcie_transfer_type_e tlp_type;
 //Changing TL packet array into structure to locate LDWE field in TLP.
 req_st    = req_arr_to_st(tlp.tlp_packet);
 tlp_type  = req_st.transfer;

 if((tlp_type == PCIE_MWR_3DW) && (req_st.length = 'h1) )begin
   req_st.ldwe = 4'b1111;
   //Converting TL structure to packet array after changing Last Byte Enable
   //field
   tlp.tlp_packet = req_st_to_arr(req_st);
 end
endfunction : do_callback
```

2. Populating the do_callback method. "tlp" is the instance of the TLP sequence item.

```
function void tl_ep_seq_num_err_cb::do_callback(tl_to_dll_t tl_to_dll);
   tl_to_dll.lcrc = tl_to_dll.lcrc + 2'h2;
endfunction : do_callback
```

3. Populating the do_callback method. "tl_to_dll" is an instance of the TLP sequence item at the DLL.

```
function void pl_os_cb::do_callback(os_plp_t os_plp);
  pl_state_e main_ltssm_state;
  pl_substate_e main_ltssm_substate;
  ps_pllp_types_e os_type = os_plp.os_type;
  assert($cast(this.cfg,config_base));
  if (cfg.m_bfm.current_substate[0] == PCIE_RECOVERY_RCVR_CFG_SUB &&  cfg.m_bfm.g_current_speed_operation[0] == PCIE_GEN3 )
  begin
    if(os_type == PCIE_TS2_OS && err == 1'b0)
    begin
      os_plp.os_type = PCIE_CTRL_SKIP_OS;
    end
endfunction: do_callback
```

4. The do_callback method is populated to replace TS2 OS with CTRL_SKP in Recovery Rcvr-Cfg state at Gen3 speed.

Lab Bench
WILLIAM WONG | Senior Content Director
bwong@endeavorb2b.com

GAME OF DRONES RETURNS: Welcome to HoverGames 2

This year's HoverGames adds the i.MX8M-based NavQ mini vision system to the competition.

NXP's HoverGames is now in its second iteration. The developer kit in the competition centers around the RDDRONE-FMUK-44FMU flight management unit (FMU) that contains a Kinetis K66 SoC based on a Cortex-M4F. It's designed to run the open-source PX4 flight management software (FMS). This system has a good bit of head room, but a bit more horsepower is needed to apply machine learning (ML) to sensors like cameras or 3D image sensors. Quite a few of the entries in the initial competition paired the FMU with platforms like a Raspberry Pi.

HoverGames 2 takes the competition to the next level by standardizing on the NavQ stack *(Fig. 1)*. The stack includes three boards, starting with the i.MX 8M Mini system-on-module (SOM). The middle board adds an SD Card, Ethernet networking, a MIPI-CSI camera interface, and a MIPI-DSI display interface. The last board provides interfaces for drones or rovers that follow the PX4 standard like the NXP FMU.

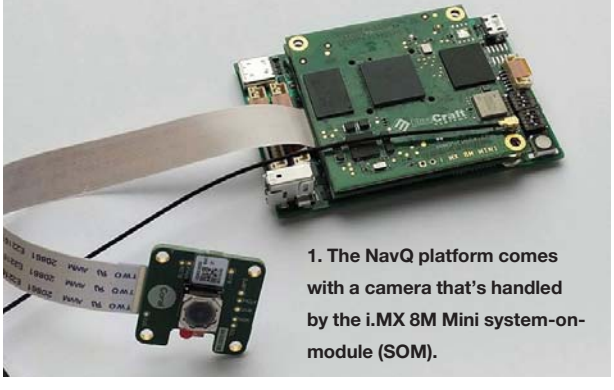

1. The NavQ platform comes with a camera that's handled by the i.MX 8M Mini system-on-module (SOM).

The i.MX 8M Mini includes a quad Cortex-A53 and Cortex-M4. The SOM adds 2 GB of LPDDR4 DRAM, 16 GB of eMMC flash memory, and a 32-MB QSPI flash. There's a PCIe M.2 interface as well. The SOM also supports Wi-Fi 802.11ac and Bluetooth 4.1.

The SOM takes the platform to the next level with ML support and the ability to run operating systems like Linux. ML models from frameworks like TensorFlow, Caffe, and PyTorch can take advantage of hardware acceleration, enabling real-time processing of sensor data as well as supporting other robotic chores like mapping and situational awareness.

In theory, the NavQ could run the PX4 FMS, but the FMU also has all of the sensors built in. The only sensor directly connected to the NavQ is the camera.

The winner of HoverGames 1 was Dobrea Dan Marius's Autonomous Human Detector Drone. It added a Raspberry Pi 3 Model B+ to the mix along with a camera. OpenCV was used to identify people by processing the video stream from the on-board camera in real time *(Fig. 2)*. The system was designed as a flying warning and risk-assessment tool. It could be used for search and rescue as well as assist in locating and tracking people in areas where there were fires. A secondary part of the project was to develop and test an ultrasonic obstacle avoidance system.

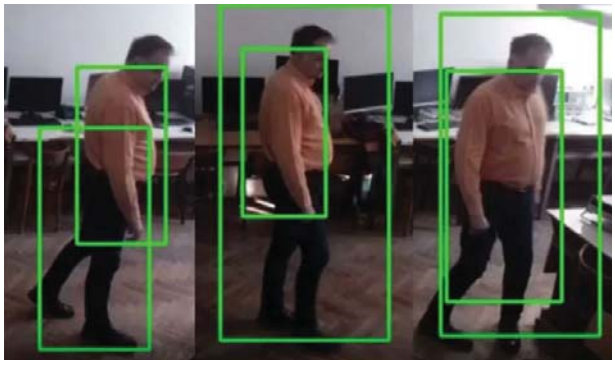

2. The winner of HoverGames 1 was Dobrea Dan Marius with the Autonomous Human Detector Drone that took advantage of image recognition. *(Courtesy of Dobrea Dan Marius)*

There were quite a few entries for the first HoverGames; you can find those online. Most, like the Human Detector Drone, have open-source code and schematics so that you can build on and extend the work done by the participants. I was impressed by the various tips and tricks that highlight design and implementation issues to help others, especially those participating in HoverGames 2.

The output from science and engineering competitions continues to amaze me. The combination of new hardware and sensors combined with machine-learning software is changing what's possible. ed

Advertisers Index

Tag-Connect

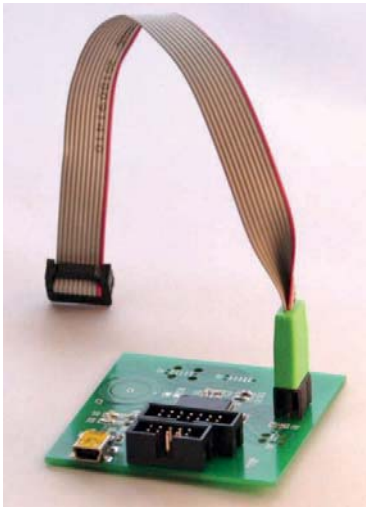

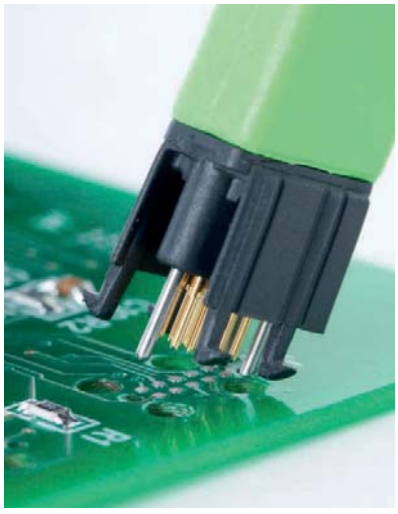

JTAG Connector Plugs Directly into PCB!!

No Header! No Brainer!

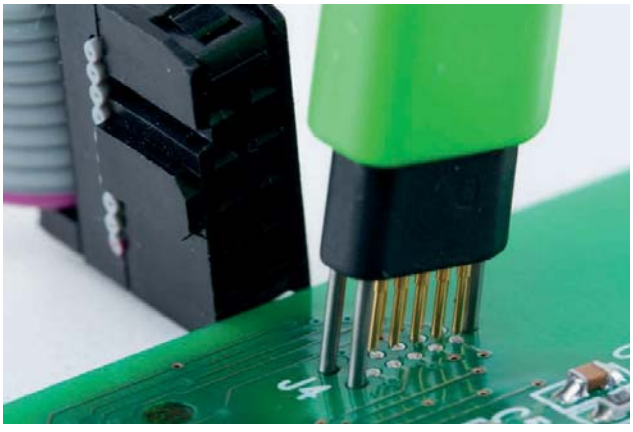

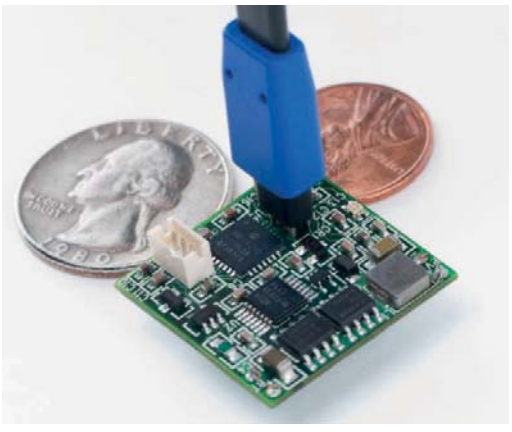

Our patented range of Plug-of-Nails™ spring-pin cables plug directly into a tiny footprint of pads and locating holes in your PCB, eliminating the need for a mating header. Spring-Pins tested to over 150,000 cycles!

Save Cost & Space on Every PCB!!

Solutions for: PIC . dsPIC . ARM Cortex . Atmel . JTAG . Altera . Xilinx . TI DSP . MSP430 . SPI / IIC . FTDI UART & More

Tag-Connector footprints as small as 0.02 sq. inch (0.13 sq cm)

www.Tag-Connect.com

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

High Performance MEMS IMU, Now in a Small, Robust Package for Industrial Applications

ADIS1650x

- Robust shock and vibration rejection in gyroscopes
- Excellent cross-axis sensitivity
- SPI interface for simple processor connection

Buy Now at mou.sr/ADIS16500

MOUSER ELECTRONICS