Electronic Design®
THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS
SUMMER 2023
electronicdesign.com
ROBOT SOFTWARE
BUILDS ON
COMMUNICATION FRAMEWORK
SELECTION
STARTS HERE
Digi-Key ELECTRONICS
Find it at digikey.com

SELECTION
STARTS HERE
More products. More name-brand suppliers.
More new product inventory. We simply have more of what you need.
Find it at digikey.com or call 1.800.344.4539.
Digi-Key ELECTRONICS
Digi-Key is an authorized distributor for all supplier partners. New products added daily. Digi-Key and Digi-Key Electronics are registered trademarks of Digi-Key Electronics in the U.S. and other countries. © 2023 Digi-Key Electronics, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA
ECIA MEMBER
Supporting The Authorized Channel

ebmpapst

engineering a better life

Axi*Force* – Experience the future of *fan technology.*

The AxiForce tubeaxial fan series is ideal for keeping consistent, optimized temperature in control units for automation and other highly-modern technologies. Due to its variable installation, high cooling capacity in the smallest of spaces, and interactive integration into the device logic, it has already proven to be indispensable in modern automation. Now with the addition of our 80 mm, IP68 rated models AxiForce is even more versatile for the future.

Contact sales@us.ebmpapst.com for more info.

IN THIS ISSUE

FEATURES

8 **ROS and DDS: Making the Most Out of Your Software Framework**
ROS 2 has become popular beyond robotics, but it has significant technical limitations. Learn how to break past these limits using the powerful framework on which ROS 2 is built.

14 **Overcome the Challenges of Using Sub-Milliohm SMD Chip Resistors (Part 1)**
Today's low-ohmic-value current-sense resistors have driven downward from milliohms to hundreds of microhms. Working with such low ohmic value, designers face challenges as well as opportunities throughout their design and manufacturing processes.

20 **Enhanced 3D-Printing Technique Yields Custom MEMS Sensors**
Researchers have used a specialized form of additive manufacturing to create MEMS sensors that can be fabricated as custom, single-quantity units.

22 **Three Major Design Pitfalls Plaguing New Analog Signal-Path Designers**
Wouldn't it be great to not repeat the same amplifier application errors many new designers fall into?

30 **What's the Difference Between CAN and SPE in the Automotive Industry?**

34 **11 Myths About 3D Flash Memory**

36 **What are the Differences Between MicroSD and MicroSD Express?**

COLUMNS & DEPARTMENTS

6 **Editorial**
The Chatbot Challenge for Embedded Programmers

39 **Ad Index**

40 **Lab Bench**

FOLLOW US ON FACEBOOK AND TWITTER

(ElectronicDesign)

(@ElectronicDesgn)

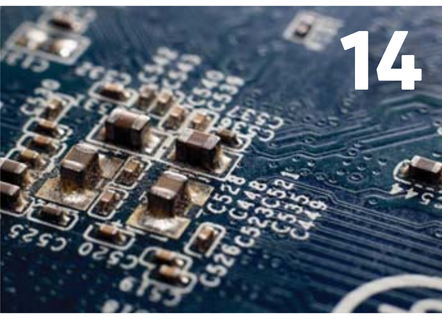

Cover images: 62617220 © Andrei Kolchin | Dreamstime.com loops7 136300536 | iStock/Getty Images Plus

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 (USPS Permit 172-080), (ISSN 0013-4872 print, ISSN 1944-9550 online) is published quarterly (Spring, Summer, Fall, Winter) by Endeavor Business Media, LLC. 1233 Janesville Ave., Fort Atkinson, WI 53538. Periodical postage paid at Fort Atkinson, WI, and additional mailing offices. **POSTMASTER:** Send address changes to Electronic Design, PO Box 3257, Northbrook, IL 60065-3257. **SUBSCRIPTIONS:** Publisher reserves the right to reject non-qualified subscriptions. Subscription prices: U.S. ($75.00 per year); Canada/Mexico ($112.50 per year); All other countries ($150.00 per year). All subscriptions are payable in U.S. funds. Send subscription inquiries to Electronic Design, PO Box 3257, Northbrook, IL 60065-3257. Customer service can be reached toll-free at 877-382-9187 or at electronicdesign@omeda.com for magazine subscription assistance or questions.

Printed in the USA. Copyright 2023 Endeavor Business Media, LLC. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopies, recordings, or any information storage or retrieval system without permission from the publisher. Endeavor Business Media, LLC does not assume and hereby disclaims any liability to any person or company for any loss or damage caused by errors or omissions in the material herein, regardless of whether such errors result from negligence, accident, or any other cause whatsoever. The views and opinions in the articles herein are not to be taken as official expressions of the publishers, unless so stated. The publishers do not warrant either expressly or by implication, the factual accuracy of the articles herein, nor do they so warrant any views or opinions by the authors of said articles.

Altech®
FULL LINE OF QUALITY
PCB Terminal Blocks
Many Styles & Pin Spacings
FIXED PCB TERMINALS
SCREW TERMINAL BLOCKS
Wire Guard Type
Lift Type
Excenter Type
SPRING TERMINAL BLOCKS
Push-In Type
Tension Spring Type
Lift Type
PLUGGABLE MULTI-CONNECTOR
SCREW TERMINAL BLOCKS
Wire Guard Type
Lift Type
Excenter Type
SPRING TERMINAL BLOCKS
Push-In Type
Tension Spring Type
PIN STRIPS
Time Saving
Push-In Technology
Easy wire insertion
No tool required !
• 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
c US
IEC
Altech Corp.®
Your Source for Automation & Control Components
AltechCorp.com | 908.806.9400

A BETTER WAY TO UL & CSA™

Experts who help you from START TO FINISH

CE too!

Design Consulting • Compliance Assistance
Product Testing & Certification Lab

CertifiGroup.com • 800-422-1651

FREE UL-CSA-CE Compliance Whitepapers

POWER SUPPLIES

HEAVY-DUTY

INDUSTRIAL & RAILWAY

Cost-effective solutions to tough design challenges from a few watts to multi-kilowatts

ABSOPULSE ELECTRONICS LTD.

www.absopulse.com

Tel: +1-613-836-3511

Since 1982

ElectronicDesign.

EDITORIAL

Senior Content Director: **Bill Wong**, *bwong@endeavorb2b.com*

Senior Editor: **David Maliniak,** *dmaliniak@endeavorb2b.com*

Managing Editor: **Roger Engelke,** *rengelke@endeavorb2b.com*

Editor-at-Large: **Alix Paultre,** *apaultre@endeavorb2b.com*

Senior Editor: **James Morra,** *jmorra@endeavorb2b.com*

Technology Editor: **Cabe Atwell,** *catwell@endeavorb2b.com*

CONTRIBUTING EDITORS

Lee Goldberg, Bill Schweber, Murray Slovick, Steve Taranovich

ART DEPARTMENT

Group Design Director: **Anthony Vitolo,** *tvitolo@endeavorb2b.com*

Art Director: **Jocelyn Hartzog,** *jhartzog@endeavorb2b.com*

PRODUCTION

Production Manager: **Brenda Wiley,** *bwiley@endeavorb2b.com*

Ad Services Manager: **Deanna O'Byrne,** *dobyrne@endeavorb2b.com*

AUDIENCE MARKETING

User Marketing Manager: **Debbie Brady** *dmbrady@endeavorb2b.com*

Article Reprints: *reprints@endeavorb2b.com*

SUBSCRIPTION SERVICES

OMEDA, 847-559-7598 or 877-382-9187 *electronicdesign@omeda.com*

LIST RENTAL

List Rentals/Smartreach Client Services Manager: **Mary Ralicki**
mralicki@endeavorb2b.com

SALES & MARKETING

Gregory Montgomery, *gmontgomery@endeavorb2b.com*
AZ, NM, TX

Jamie Allen, *jallen@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

Elizabeth Eldridge, *eeldridge@endeavorb2b.com*
CT, DE, MA, MD, ME, NH, NJ, NY, PA, RI, VT, EASTERN CANADA

Stuart Bowen, *sbowen@endeavorb2b.com*
AK, NORTHERN CA, NV, OR, WA, WESTERN CANADA
AUSTRIA, BELGIUM, FRANCE, GERMANY, LUXEMBURG, NETHERLANDS, PORTUGAL, SCANDINAVIA, SPAIN, SWITZERLAND, UNITED KINGDOM

Diego Casiraghi, *diego@casiraghi-adv.com*
ITALY

Helen Lai, *helen@twoway-com.com*
PAN-ASIA

Charles Liu, *liu@twoway-com.com*
PAN-ASIA

DIGITAL

VP Digital Innovation Data: **Ryan Malec,** *rmalec@endeavorb2b.com*

ENDEAVOR BUSINESS MEDIA, LLC

30 Burton Hills Blvd, Ste. 185, Nashville, TN 37215 | 800-547-7377
www.endeavorbusinessmedia.com

ENDEAVOR BUSINESS MEDIA

CEO: **Chris Ferrell**
President: **June Griffin**
CFO: **Mark Zadell**
COO: **Patrick Rains**
CRO: **Reggie Lawrence**
Chief Administrative and Legal Officer: **Tracy Kane**

DESIGN & ENGINEERING GROUP

EVP, Design & Engineering Group: **Tracy Smith,** *tsmith@endeavorb2b.com*

Group Content Director: **Michelle Kopier,** *mkopier@endeavorb2b.com*

Electronic Design, Machine Design, Microwaves & RF, Power & Motion, SourceESB, Source Today, 3DX

Altech®
LARGEST Selection of
DIN ENCLOSURES
8 different series with over 240 parts
KS4400
Series
KU4000
Series
KO4070
Series
KU4100
Series
KO4700
Series
K70
Series
KS4400
Series
• DIN Rail and Panel Mount
• Fixed and Pluggable Terminals
• Integral PC Board Guides
• Many terminal to PCB
connection options
Multi-functional
Enclosure System
NEW
KV4600
Series
Altech Corp.®
Your Source for Automation & Control Components
AltechCorp.com | 908.806.9400

DC-DC Converters
Transformers &
Inductors

DC-DC Converters
2V to 10,000 VDC Outputs
1-300 Watt Modules
• MIL/COTS/Industrial Models
• Regulated/Isolated/Adjustable Programmable Standard Models
• New High Input Voltages to 900VDC
• AS9100D Facility/US Manufactured
• Military Upgrades and Custom Modules

Transformers & Inductors
Surface Mount & Thru Hole
• Ultra Miniature Designs
• MIL-PRF 27/MIL-PRF 21308
• DSCC Approved Manufacturing
• Audio/Pulse/Power/EMI Multiplex Models Available
• For Critical Applications/Pico Modules, Over 50 Years' Experience

VISIT OUR EXCITING NEW WEBSITE
www.picoelectronics.com

For full characteristics of these and the entire PICO product line, see PICO's Full line catalog at
www.picoelectronics.com

PICO ELECTRONICS, Inc.
143 Sparks Ave., Pelham, New York 10803
Call Toll Free 800-431-1064
FAX 914-738-8225
E Mail: info@picoelectronics.com

Editorial

BILL WONG | Senior Content Director
bwong@endeavorb2b.com

The Chatbot Challenge for Embedded Programmers

CHATBOTS ARE ONE form of generative machine-learning (ML) model designed to interact with people in a conversational mode through multiple interactions, unlike many other types of ML models that are presented with an input and provide an output. Chatbots such as ChatGPT, Google Bard, YouChat, JasperChat, Bing Chat, and now Amazon's Bedrock are just the tip of the iceberg, and much of the world is the beta tester for these tools.

Generative artificial-intelligence (AI) tools like chatbots are typically trained using massive amounts of data, enabling them to respond based on this information. The challenge for these tools is that they need to turn a question into a search and then provide a response-based analysis of the question as well as the data it was trained with. Generative AI goes a step further with the ability to accept additional questions and combine those with the prior query to provide a subsequent response.

To many, this may appear to be "real AI." However, it's a far cry from the "positronic brains" of movies like *I, Robot* based on science fiction from authors like Isaac Asimov. The current crop of chatbots are improving, but they typically go "off the rails" into nonsense or bad responses after more than half a dozen consecutive interactions. Likewise, bias in data, training, and implementation can affect the quality and reliability of the responses.

One aspect of generative AI tools is "explainability," or the ability to describe why a response was provided. Other ML models can provide similar information, which can be invaluable in determining the usefulness of the results. Still, that leads to the issues about using chatbots for embedded programming chores.

So, what can chatbots do for embedded programmers? They're being tasked to do many things from finding errors to writing programs to finding examples of other programs that might be useful. The latter is a relatively safe way to use chatbots. It's the reason so many big names in the search arena are interested in them, as they can improve the search results in a more user-friendly way. From our perspective, it does the same and we can still evaluate the identified software based on our expertise.

A similar observation can be made when using a chatbot to identify problems in code, although many issues should be considered. They range from the amount of false positives to the scope of the query.

The chatbots' interactive nature and ability to explain its reasoning can make them better than conventional static-analysis tools. However, the quality, scope, and efficiency make conventional tools better for now. Many will be augmented using chatbot technology, but currently this is more in the experimental stage.

Creating and modifying programs also can be performed by some chatbots—this is where developers need to step carefully. Keep in mind that these tools are utilizing training content like the hoard of programs stored on github. They can range from trivial and poorly designed systems to massive, well-architected solutions.

Currently available generative AI is simply taking it all in with no real understanding of these differences. It may be nice to get a custom application to blink an LED, but it's probably not a good idea to have one of these tools generate the braking system for an airplane.

I'm both amazed and horrified by the interactions I've had with various chatbots. Unfortunately, oftentimes they're like dealing with search engine optimization (SEO) in that SEO tends to involve a black box with ever-changing secret algorithms.

My advice is to not ignore the technology, but be very, very careful in how you use it. ed

DOES YOUR APPLICATION HAVE LONG LIFETIME REQUIREMENTS?

Rochester is licensed to manufacture devices no longer produced by the original component manufacturer. We have manufactured over 20,000 device types. With over 12 billion die in stock, we have the capability to manufacture over 70,000 device types.

- *Products are manufactured using design information transferred directly to Rochester from the original component manufacturer (OCM)*
- *Offering ongoing manufacturing of stock products*
- *Offering build-to-order products Process flows include Commercial, Industrial and Military Temperature, MIL-STD-883, SMD/QML, and Space-Level S/V. IATF 16949:2016 certified for automotive production*
- *A range of industry standard packaging with a variety of lead finishes including Sn, SnPb, and RoHS.*

Authorized Distribution | Licensed Manufacturing | Manufacturing Services

NEIL PUTHUFF | Senior Automotive Applications Engineer, RTI

ROS and DDS: Making the Most Out of Your Software Framework

ROS 2 has become popular beyond robotics, but it has significant technical limitations. Learn how to break past these limits using the powerful framework on which ROS 2 is built.

The introduction and ongoing releases of the robot operating system, ROS 2 (now built on top of the DDS framework), has expanded its use beyond its original focus on robotic research. ROS 2 comes bundled with application packages and visualization tools, so it facilitates making robotics systems that can sense, map, and navigate their surroundings.

About DDS

Data Distribution Service (DDS) is an open-standard, data-centric communications software framework with more than a dozen commercial and open-source implementations. It provides low latency, extreme reliability, and a rich set of quality of service (QoS) controls to enable robust peer-to-peer communications in the most challenging of environments: contested battlefields, noisy industrial settings, wide-area networks, and remote systems with intermittent connectivity.

DDS has been used in thousands of critical systems, hundreds of autonomous-vehicle programs, and dozens of other frameworks and standards including ROS 2, AUTOSAR, and FACE.

About ROS

The free, open-source ROS project is a one-stop shop for quickly creating robotics applications and systems. First released in 2010, the original ROS rapidly became popular in academia, eventually turning into the dominant framework for robotics researchers and educators.

The main value in ROS is in its tools and pre-built packages. Its main disadvantage was the middleware, which prevented it from being used in critical or constrained systems, in multi-robot swarms, or in applications with real-time constraints. As a result, ROS remained largely in academia for more than a decade.

ROS + DDS = ROS 2

ROS 2 is a redesign of the original ROS that should help solve emerging challenges in robotics. Built on top of the DDS framework, ROS 2 seeks to operate in constrained systems, multi-robot swarms, and production-grade platforms—an ideal marriage that combines the outstanding tools and packages of ROS with the “works everywhere” capa-

Kunnicha-Luengrungwaree_195994310 | Dreamstime

bilities of DDS. ROS 2 has helped ROS break out of academia.

The success of ROS 2 has led to the retirement of the original ROS following the 2020 "Noetic" release. All future ROS development is now on ROS 2, and all communications within ROS 2 flow via DDS. So, in effect, all ROS 2 applications also are DDS applications, and the ROS 2 ecosystem is a part of the DDS ecosystem. Read on to learn why this is an important distinction.

Is ROS 2 the Sum of its Parts?

Not quite. ROS 2 is a big improvement over the original version of ROS, but when ROS 2 is viewed from the field-proven perspective of DDS, it has some significant limitations.

However, those limitations can be overcome by directly accessing the DDS framework upon which ROS 2 is built, enabling system developers to get the full benefit of DDS while maintaining compatibility with ROS 2. Let's examine these limitations and how they can be resolved using this approach.

Performance

Because ROS 2 is layered on top of DDS, any data sent or received within ROS 2 must travel through these layers before reaching the underlying DDS framework. This takes a substantial amount of time.

By designing critical application components to directly use the underlying DDS API, engineers can eliminate many performance bottlenecks. For example, during the 2021 Indy Autonomous Challenge, a ROS 2 LiDAR device driver was modified to directly use the DDS API. This cut latency by up to 90% compared with a driver that used ROS API but didn't affect full interoperability with ROS 2 *(Fig. 1, page 10).*

This approach was employed to dramatic effect by the winning team of the Indy Autonomous Challenge, a $1.5M competition to autonomously race full-size Dallara IL-15 vehicles at the famous Indianapolis Motor Speedway. Not only did the approach eliminate the bulk of data latency, it also freed up two-thirds of the memory used by the ROS 2 driver, while maintaining drop-in compatibility with ROS 2.

Scalability

ROS 2 systems can quickly run into scalability issues because they create large numbers of data topics (unique, discoverable data flows) to support the implementation of ROS Parameters, Services, and Actions in ROS 2.

In operation, every ROS 2 application node creates more than a dozen unique topics for Parameters (even if you don't use them), and more unique topics are created for every ROS 2 Service, Message, and Action used in your applications. Consequently, it doesn't take long for a system to have hundreds or thousands of topics competing for space on the network.

An equivalent DDS system can avoid this overhead. First, common data types may be implemented using Keyed Topics, which enables large numbers of unique data sources to share common data flows while retaining their unique identity, easing congestion in large-scale systems.

Second, DDS can be used to create topic gateways to partition your large ROS 2 system into a tiered hierarchy, permitting only the necessary data to flow to other parts of your system. This comes in handy in situations such as creating a system of many mobile robots connected over a common wireless network.

Third, any applications implemented directly on DDS will not create the dozen-or-more topics created by ROS 2 for its parameter system, which directly reduces the number of topics on the network. The result is a system with reduced network traffic, faster startup, and greater scalability than the equivalent ROS 2 implementation.

Quality of Service (QoS)

QoS plays the primary role in assuring dependable system communications under real-world conditions. While DDS itself has a comprehensive set of QoS capabilities, the majority aren't accessible from within ROS 2. As a comparison:

ROS 2 has support for eight QoS categories, with basic on/off and sizing control for History, Depth, Reliability, Durability, Deadline, Lifespan, Liveliness, and Lease Duration.

OMG Standard DDS has 22 QoS categories. Here's a few highlights of what's not in ROS 2:

- *Entities and Keys:* These enable vast scale-up by sharing common data topics using unique "Keys" for each data source. This makes the system start faster and run more efficiently, and it reduces the load on the network.
- *Instance Management:* This allows applications to manage open-ended collections of objects (e.g., robot swarms) where all objects in the collection publish the same Topics but they're distinguished

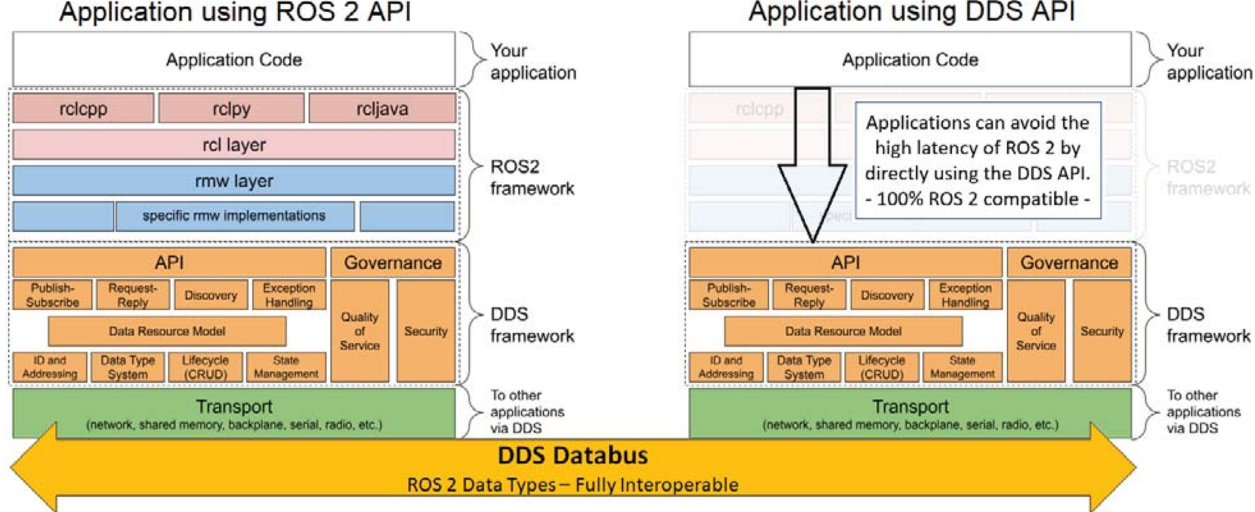

1. Applications that directly use the DDS APIs beneath ROS 2 eliminate the data latency of the ROS 2 framework, while maintaining full ROS 2 compatibility. *RTI*

by an Instance Key (typically the identifier of the object in the collection, e.g., a robot_id). Subscribers subscribe to the "collection" topic and are notified as new objects join or leave the collection or if any object ceases to be active. The data published by each object in the collection is maintained in a separate sub-stream (identified by the Instance Key) and the QoS is applied to each sub-stream separately.

- *Data Cache:* This allows the application to access the data cache of the subscriber (DataReader in DDS terms) directly. The cache provides zero-copy views of the data received synchronously or asynchronously with data reception. The data can be accessed multiple times (e.g., to perform aggregation or sensor fusion) without requiring the application to keep extra copies.
- *Ownership and Strength:* These provide for automatic failover when using redundant components. This also enables zero-downtime during system upgrades, maintenance, and testing.
- *Transport Priority:* This allows a system to give preference to a transport when it's available. For example, a mobile system using cellular communications can be made to automatically switch over to Wi-Fi when available.
- *Partitions:* These enable a network to be divided into isolated logical sections. Data flows within these sections can't interact with outside elements, thus providing system-scale encapsulation of related areas.

Vendor-Extended DDS offers even more capabilities to meet real-world networking challenges. For example, RTI Connext DDS has options for wide-area-network (WAN) connectivity, small sample batching, content filtering and querying, compression and bandwidth reduction, time-sensitive networking (TSN), and more.

Each DDS implementation can choose where to focus its efforts in these QoS categories, from simple on/off controls to fully tunable parameters, or to omit the category entirely. These will have a direct impact on whether your system can run in a particular environment or not. Simply put, more QoS means more adaptability. A capable DDS implementation will be able to operate in countless environments where ROS 2 cannot.

Interoperating with Non-ROS Systems

As DDS is used in thousands of critical systems and in standards such as AUTOSAR and FACE, a ROS 2 system also may need to interoperate with these non-ROS environments based on DDS.

While DDS provides standards-based interoperability, the implementation within ROS 2 imposes a set of rules and restrictions on the names of data types and topics that it will accept. Any data that falls outside of these rules is ignored or unusable. This has a direct impact on ROS 2's ability to interoperate with non-ROS systems built on DDS, so a separate bridge application must be written to translate these non-conforming data types into something that ROS 2 can accept.

However, an application written to directly use the DDS API can easily interoperate with any DDS-based system, such as ROS 2, AUTOWARE, FACE, and many others. In this role, DDS can be thought of as something like a "Hypervisor for Distributed Systems," capable of directly integrating many different types of DDS-based applications (ROS 2, AUTOSAR, etc.) into a high-performance system.

Getting the Best of DDS and ROS 2

DDS maintains a very stable, standards-

FOR THE NEXT STAGE IN POWER TESTING

- lUltra-compact:20kW in a single 3U size frame.
- Maximum voltage:Supports high voltages of 1500V
- Bidirectional output: Source and Sink in a single unit! Increased capacity suitable for testing batteries, inverters, etc.
- 100% full output, whether the AC input voltage is 200 V or 400 V.
- Operating modes* : CV (constant voltage)/CC (constant current)/ CP(constant power)
  *Priority setting available
- Parallel operation up to 10 units(200kW)
- Equipped with an intuitive touch panel.
- Digital I/O: Customizable control items (5 inputs, 6 outputs)

Bidirectional DC Power Supplies

PXB series

https://kikusuiamerica.com kikusui@kikusuiamerica.com

Kikusui America, Inc. 3625 Del Amo Blvd, Ste 160 Torrance, CA 90503 Tel: 310-214-0000

based software API, and ROS 2 has a very stable data model and is itself based on DDS. However, ROS has a history of introducing changes to its software API with each major release, which creates an ongoing burden for developers who must update and version-stamp their applications to keep pace with ROS.

There's a way to get the full benefits of DDS while retaining the benefits of ROS 2 and eliminating its drawbacks: Implementing critical components directly in DDS using the ROS 2 data types.

This hybrid combination has some compelling benefits:

- Significantly improved performance. Eliminating the ROS 2 software layers can eliminate more than 90% of the latency of data communications. This approach also enables full access to advanced capabilities of DDS implementations that aren't supported in ROS 2, such as zero-copy transfers, reduced encoding, compression, etc.
- Scalability unbounded. By using keyed topics and DDS-enabled hierarchical design, your system can scale to massive proportions, as was done with the Constellation control system at the Kennedy Space Center (the largest SCADA system in the world).
- Access to all DDS QoS. Quality of service enables your systems to work in the most demanding environments, where ROS 2 cannot venture.
- Fully interoperable with ROS 2, and with non-ROS DDS systems.
- Supported on production-grade hardware and operating systems, and available in safety-certifiable versions.

While this might seem like a risky introduction of new technology to a ROS 2 development, it's actually the opposite. Developers using ROS 2 are already working with DDS; this change merely gives them a more direct connection to DDS. Let's now look at how that affects the software.

Examining the Software API

ROS and DDS are both data-centric, publish/subscribe technologies with very similar design patterns. Taking a closer look at the ROS 2 API shows that the bulk of the API functions cover data communications, which are provided under the hood by the DDS API. The remainder are mostly common system-level functions (files, timers, callbacks, etc.) that can be found in standard libraries.

Therefore, the API patterns of creating applications directly in DDS should feel very familiar to developers accustomed to ROS 2. But what about the source code? *Figure 2* illustrates a comparison of a simple "Hello World!" application written using the ROS 2 and DDS APIs.

Implementing systems directly in DDS can be a natural progression for developers using ROS 2. It follows similar design patterns but provides far greater control over system communications, while eliminating many layers of software that hinder debugging.

Implementing the Improvements

A typical pattern for this hybrid use case is to replace critical system components written in ROS 2 with their native DDS equivalents. Because of the standards-based interoperability of DDS, the replacements can be dropped in without disruption to the remainder of the ROS 2 system. An example might be a signal-processing module for a camera or LiDAR that needs to operate at minimum latency, or a gateway application to communicate with many peer units over a radio transport.

The system developer can implement these improvements strategically, replacing only critical components while leaving the remainder of the system in ROS 2, or as part of an overall system transition to native DDS. Interoperability with the ROS 2 ecosystem can be maintained continuously.

Wrapping It All Up

ROS 2 is a popular way to quickly create robotics systems using distributed applications, but it may struggle with scalability, performance, and operation in challenging environments—all places where DDS excels.

Fortunately, ROS 2 is implemented on top of DDS, meaning that system developers can freely intermix ROS 2 and DDS applications to solve specific ROS 2 shortcomings. Or they can migrate their entire system to a higher performing, open-standard DDS option without losing interoperability with ROS 2 and its excellent ecosystem of tools and packages. Truly, the best of both ROS 2 and DDS is available now for builders of robotic and autonomous systems. ed

To learn more, visit https://community.rti.com/ros.

ROS 2 / C++

```
#include <chrono>
#include <functional>
#include <memory>
#include <string>
#include "rclcpp/rclcpp.hpp"
#include "std_msgs/msg/string.hpp"

using namespace std::chrono_literals;

class MinimalPublisher : public rclcpp::Node
{
public:
    MinimalPublisher()
        : Node("minimal_publisher"), count_(0)
    {
        publisher_ = this->create_publisher<std_msgs::msg::String>("topic", 10);
        timer_ = this->create_wall_timer(
            500ms, std::bind(&MinimalPublisher::timer_callback, this));
    }

private:
    void timer_callback()
    {
        auto message = std_msgs::msg::String();
        message.data = "Hello, world! " + std::to_string(count_++);
        RCLCPP_INFO(this->get_logger(), "Publishing: '%s'", message.data.c_str());
        publisher_->publish(message);
    }
    rclcpp::TimerBase::SharedPtr timer_;
    rclcpp::Publisher<std_msgs::msg::String>::SharedPtr publisher_;
    size_t count_;
};

int main(int argc, char* argv[])
{
    rclcpp::init(argc, argv);
    rclcpp::spin(std::make_shared<MinimalPublisher>());
    rclcpp::shutdown();
    return 0;
}
```

DDS / C++

```
#include <iostream>
#include <dds/pub/ddspub.hpp>
#include <rti/util/util.hpp> // for sleep()
#include "string.hpp"

void publisher_main(int domain_id)
{
    dds::domain::DomainParticipant participant (domain_id);
    dds::topic::Topic<std_msgs::msg::dds_::String_> topic (participant, "rt/topic");
    dds::pub::DataWriter<std_msgs::msg::dds_::String_> writer(dds::pub::Publisher(participant), topic);
    std_msgs::msg::dds_::String_ sample;
    for (int count = 0; ; count++) {
        sample.data_(std::string("Hello, world! " + std::to_string(count)));
        std::cout << "Publishing " << sample.data_() << std::endl;
        writer.write(sample);
        rti::util::sleep(dds::core::Duration(0, 500000000));
    }
}

int main(int argc, char *argv[])
{
    try {
        publisher_main(0);
    } catch (const std::exception& ex) {
        std::cerr << "Exception in publisher_main(): " << ex.what() << std::endl;
        return -1;
    }
    return 0;
}
```

2. Here's a comparison of C++ source code for a "Hello, World!" application written in ROS 2 vs. code written directly in the underlying DDS (RTI Connext) used by ROS 2. *RTI*

Transient Immunity Testers

The Avtech AVRQ series of high-voltage, high-speed pulsers is ideal for testing the common-mode transient immunity (CMTI) of next-gen optocouplers, isolated gate drivers, and other semiconductors. GPIB, RS-232, Ethernet ports are standard.

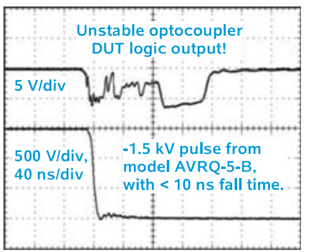

- Kilovolt amplitudes (±1 kV, ±1.5 kV, ±2 kV)
- Transition times down to < 10 ns, dV/dt rates up to 120 kV/us
- Daughterboards to handle a variety of DUT package styles

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 ±30V, ±50, and ±100 Volt general-purpose lab pulsers. In many applications, these can replace the discontinued Agilent 8114A or HP214. Higher-voltage models are also available.

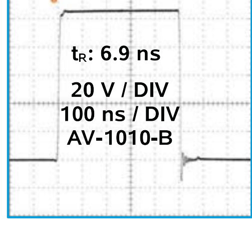

AV-1015-B: ±50V, 10 MHz, 20 ns - 10 ms, 10 ns rise
AV-1010-B: ±100V, 1 MHz, 20 ns - 10 ms, 10 ns rise
AV-1011B1-B: ±100V, 100 kHz, 100 ns - 1 ms, 2 ns rise
AV-1011B3-B: ±30V, 100 kHz, 100 ns - 10 ms, 0.5 ns rise

Includes switchable output impedance and double-pulse modes. Flexible triggering. GPIB / RS-232 / Ethernet.

Nanosecond Electronics
Since 1975

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

Nanosecond Laser Diode Drivers With Butterfly Diode Sockets

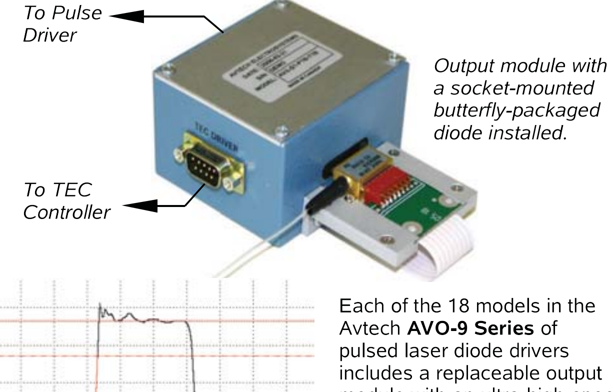

Output module with a socket-mounted butterfly-packaged diode installed.

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

Each of the 18 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 pulses.

Models with maximum pulse currents of 0.2A 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

High Output, Low Risetime Pulsers

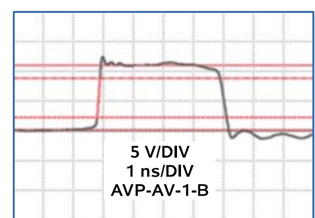

Nanosecond and faster pulsers covering a wide range of amplitudes!

GPIB / RS-232 / Ethernet standard on “-B” models.

Standard and customized models available.

Ampl	tRISE	Max. PRF	Model
5 V	80 ps	1 MHz	AVP-AV-1S-C
10 V	120 ps	1 MHz	AVP-AV-1-B
20 V	120 ps	1 MHz	AVP-AV-HV2-B
20 V	200 ps	10 MHz	AVMR-2D-B
40 V	150 ps	1 MHz	AVP-AV-HV3-B
50 V	500 ps	1 MHz	AVR-E5-B
100 V	500 ps	100 kHz	AVR-E3-B
100 V	300 ps	20 kHz	AVI-V-HV2A-B
200 V	1 ns	50 kHz	AVIR-1-B
200 V	2 ns	20 kHz	AVIR-4D-B
400 V	2.5 ns	2 kHz	AVL-5-B-TR

Nanosecond Electronics
Since 1975

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

STEPHEN OXLEY |
Business Development Engineer, Fixed & Variable Resistors, Sensors & Specialist Components Division, TT Electronics

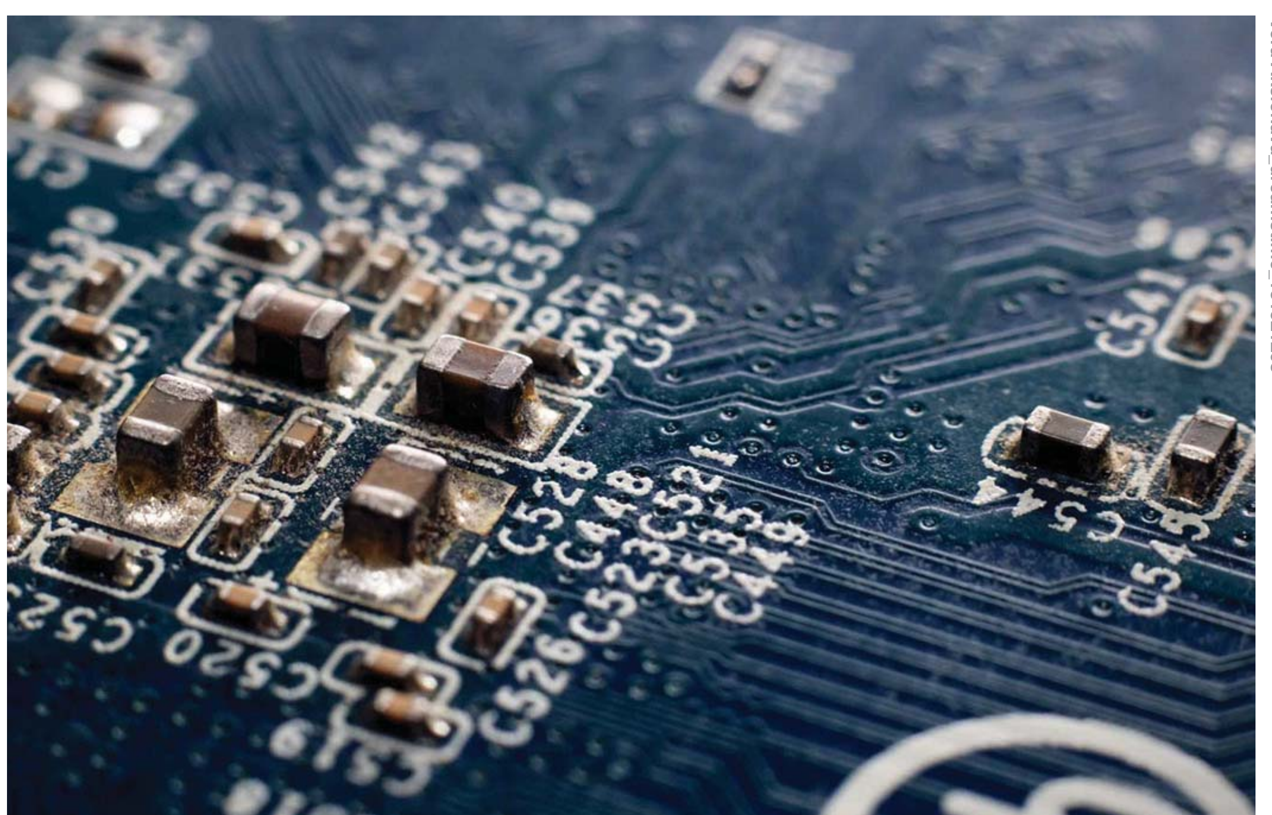

Vera Aksionava_dreamstime_131521268

Overcome the Challenges of Using **Sub-Milliohm SMD Chip Resistors** (Part 1)

Today's low-ohmic-value current-sense resistors have driven downward from milliohms to hundreds of microhms. Working with such low ohmic value, designers face challenges as well as opportunities throughout their design and manufacturing processes.

The simplest and most cost-effective way of converting a measured current to a voltage signal is to use a low-ohmic-value current-sense resistor. The upswing in products containing batteries, motors, or actuators that call for current monitoring or control has led to huge market growth for current-sense chip resistors with values below one ohm over the last two decades.

More recently, though, driven by power-efficiency demands and enabled by low-noise voltage-sense amplifiers, the value range has extended downward from milliohms to hundreds of microhms. Such low ohmic values present challenges to the user at many stages in their design and manufacturing processes.

This two-part series considers the nature of these challenges and suggests strategies to overcome them at various stages, including component selection, PCB layout design, verification of the ohmic value of unmounted components, and critical assembly processes. Each stage features potential pitfalls but

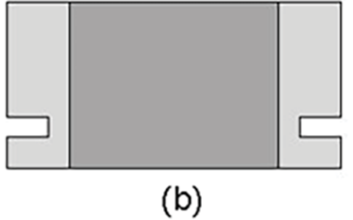

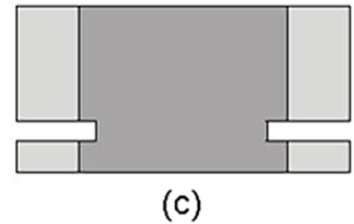

1. A Kelvin connection, or separation between the current-carrying and voltage-sense circuit, becomes increasingly important as nominal ohmic value is reduced. The most common option is a Kelvin connectable two-terminal style. (a) There's also an intermediate option, a semi-Kelvin four-terminal style (b), and a true Kelvin four-terminal style.

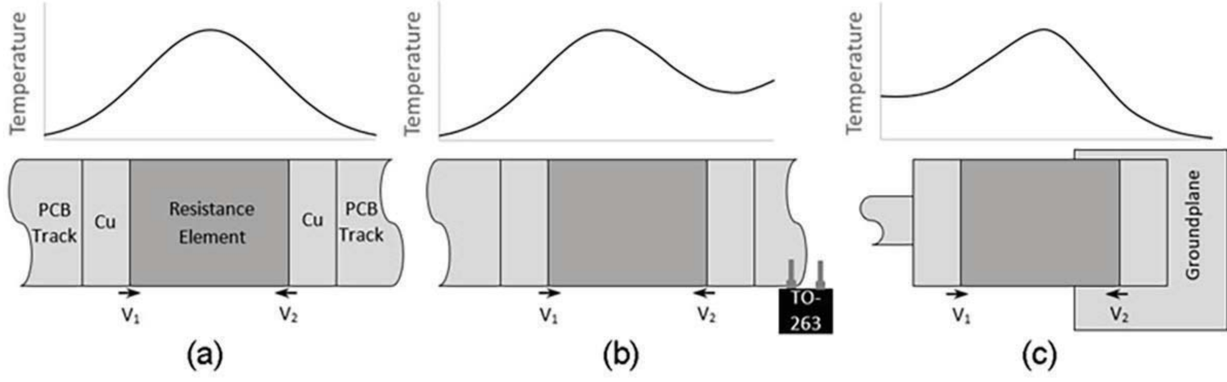

2. Shown is a balanced state in which the thermal voltages V1 and V2 are equal (a), and examples of imbalance due to the external influence of a heat source (b) and a heatsink (c).

also opportunities to quantify and minimize error and variation.

Although sub-milliohm chip resistors are still just chip resistors, ideal design strategies treat them as a separate class of component, embracing considerations and techniques that enable their successful use.

Component Selection

Termination Styles

When restricting the temperature sensitivity of a resistor-based current-sensing circuit, it's well-known that the total terminal resistance from copper, common to both the current-carrying circuit and the voltage-sense circuit, must be restricted as well. That's because this element of the total measured resistance has a temperature coefficient of resistance (TCR) of +3,900 ppm/°C, which contributes to the total TCR in proportion to the resistance ratio.

For example, a total terminal resistance of 100 μΩ, i.e., 50 μΩ at each end, for a 1-mΩ resistor contributes 100 μΩ/1000 μΩ × 3,900 ppm/°C = 390 ppm/°C to the TCR. This contrasts with the TCR of the resistance element itself, which is typically better than ±30 ppm/°C. Such separation between the current-carrying and voltage-sense circuits is referred to as a Kelvin connection, an issue that becomes clearly more important as the nominal ohmic value is reduced.

There's a spectrum of termination styles to address this problem. The most common is the Kelvin connectable two-terminal resistor *(Fig. 1a)*. While this is often the lowest-cost option, it places the onus on the PCB designer to realize a Kelvin connection in the PCB track layout (later we'll look in detail at how this may best be achieved). Such a component must have a low termination resistance since the Kelvin connection strategy necessarily ends at the surface of the termination.

Figure 1b shows an intermediate component type where four solder terminals are provided, but the separation of circuits doesn't extend all the way into the resistor element. *Figure 1c* shows a true Kelvin format where there's no current-carrying termination whatsoever within the voltage-sense circuit. The latter two types offer error-proof PCB layout design and the lowest achievable magnitude of TCR, but generally it comes at a higher cost.

Element Materials

Low-value resistors can be made from both thick- and thin-film materials, but the lowest values available in these technologies are in the multiple-milliohm range. Both types are relatively susceptible to damage from high-current surges. In the case of thick-film technology, the

lowest values are associated with high TCR values of several hundred ppm/°C and thus are suited only to low precision use.

For these reasons, most current-sense chip resistors are based on a bulk metal element. This may be either a foil supported on a substrate or a self-supporting metal element. Though the former option allows for the use of thin metal layers to achieve higher values, the latter lends itself to sub-milliohm values.

A range of alloys, each with differing resistivities, are selected by device designers to provide the required ohmic value within the dimensional constraints of the product. From the point of view of the user, the material choice is often unimportant, but there are two exceptions: one is the control of thermally generated errors, and the other is application for non-dc circuits.

A copper-terminated metal-element chip resistor contains at least two boundaries between dissimilar metals. These act as thermocouples and generate a thermoelectric voltage in the presence of a temperature gradient. Furthermore, they're connected in series, and because of the component's symmetry, they're of opposite polarity when the resistor element itself is the main heat source. As a result, if the temperature distribution across the chip resistor is symmetrical, any generated thermoelectric voltages will be cancelled out.

Figure 2a illustrates this balanced state in which the thermal voltages V_1 and V_2 are equal. *Figures 2b and 2c* show an example of imbalance due to the external influence of a heat source and a heatsink, respectively. This would lead to a finite value of $V_1 - V_2$, which would sum with the measured sense voltage and create a source of error.

In many designs, it's simply not possible to guarantee thermal symmetry under all operating conditions. In such a case, a chosen part should employ a resistance alloy with a low thermoelectric voltage against copper. These alloys contain manganese in a copper nickel alloy in which the proportion of copper exceeds 80%. The thermoelectric voltages generated across a junction with copper can be as low as 3 µV/°C, which is an order of magnitude lower than for a copper nickel alloy.

The second application-specific driver for resistance alloy selection is the need to avoid iron-bearing alloys in circuits where ac or rapid step changes in dc need to be tracked accurately, as this is not possible with ferromagnetic alloys.

Thermal Design Format

An inherent problem with resistive current sensing is heat, generated at a rate proportional to the square of the current. This may need to be restricted for one of two reasons. First, the design must reduce the effect of temperature increase on the linearity of the component, which can stem from TCR or from thermoelectric voltage errors. Second, it's necessary to avoid overheating the resistance alloy, which can lead to irreversible ohmic-value change.

3. Illustrated are an open-air format metal-alloy resistor (a), and a flat-chip format metal-alloy resistor (b).

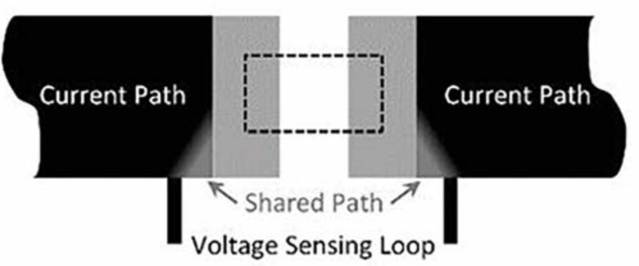

4. Minimizing the conductive path shared between the current path and the voltage-sensing loop increases both the effective ohmic value and the TCR of the mounted part.

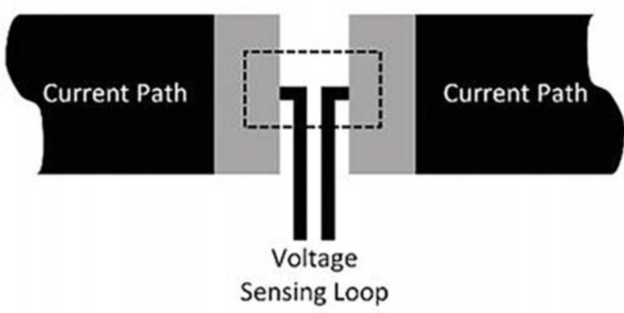

5. To achieve what's illustrated in Figure 4, connect the voltage-sense tracks to the inner edges of the solder pads.

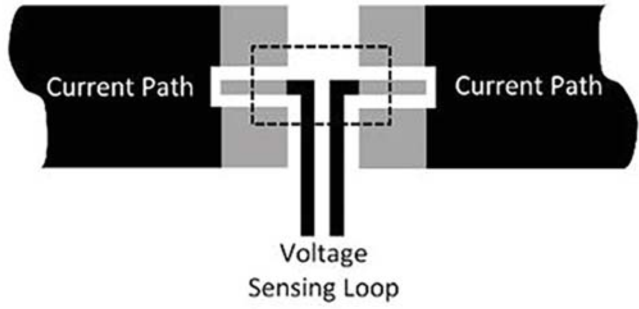

6. By splitting the voltage-sense pads from the current-path pads, the solder joints themselves also are removed from the shared path.

This consideration calls for careful thermal design of the assembly. It begins at the component selection stage in response to the basic decision as to where the heat generated should be dissipated; for example, in the air or in the copper PCB tracks. The answer to this depends on the overall thermal-management strategy.

A well-ventilated assembly—with either a high thermal loading already on the PCB or with temperature-sensitive components—would benefit from a resistor that dissipates heat into the air. Conversely, a PCB that's heatsinked, or has no excess of heat generation and no temperature-sensitive parts, can employ a resistor that dissipates heat mainly to the PCB tracks.

An example of a primarily air-dissipating open-air format is shown in *Figure 3a.* This can sustain a temperature rise of the hotspot above the solder joints well in excess of 100°C. Its flexible nature makes it virtually immune to temperature cycling or board flex stresses on the solder joints. An example of the primarily PCB dissipating flat-chip format is shown in *Figure 3b.* The design benefits from its low profile and is generally the lower-cost option.

Most Common Type

After considering the many options of termination style, element material, and thermal design format, the most common type of sub-milliohm resistor is a two-terminal metal-element chip resistor. This resistor type will be considered hereafter in this article.

PCB Layout Design

Near-Kelvin Connection

The PCB layout design around a very-low-value resistor is critical to its performance. The most important aspect of this design is the fact that four rather than two tracks must be provided to form a Kelvin connection, even where the component itself has only two terminals.

The aim is to minimize the conductive path shared between the current path and the voltage-sensing loop *(Fig. 4),* which would increase both the effective ohmic value and the TCR of the mounted part. This may be achieved by connecting the voltage-sense tracks to the inner edges of the solder pads *(Fig. 5).*

Designers also can take this a step further and split the voltage-sense pads from the current-path pads, so that the solder joints themselves are removed from the shared path, too *(Fig. 6)*. By using this method, it's possible to approach the accuracy obtained from a true four-terminal resistor.

Furthermore, a study[1] by Analog Devices based on TT Electronics' ULR3 0.5-mΩ mounting-pad options has shown that a mounted value tolerance close to 1% may be achieved on a 1% tolerance component. This indicates low additional error due to mounting effects, using a centralized, isolated sense-pad design similar to that of *Figure 4c.*

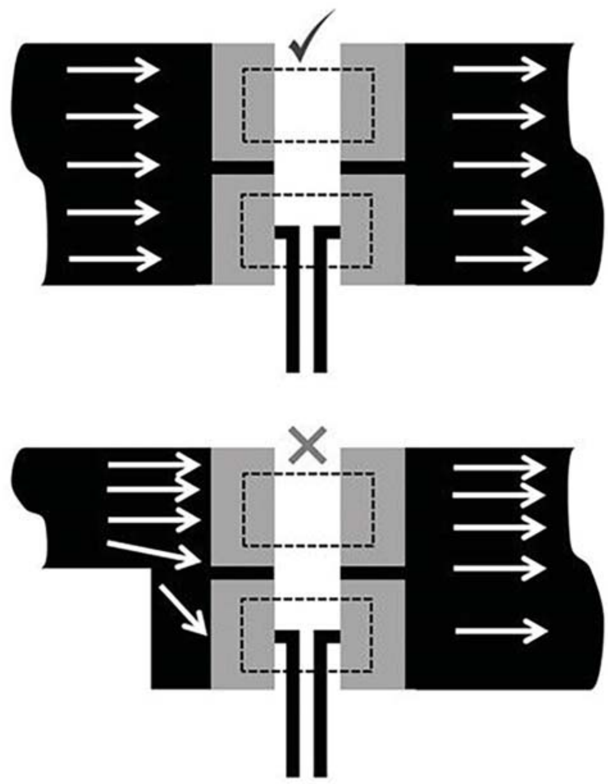

7. Parallel connection requires equal distribution of current.

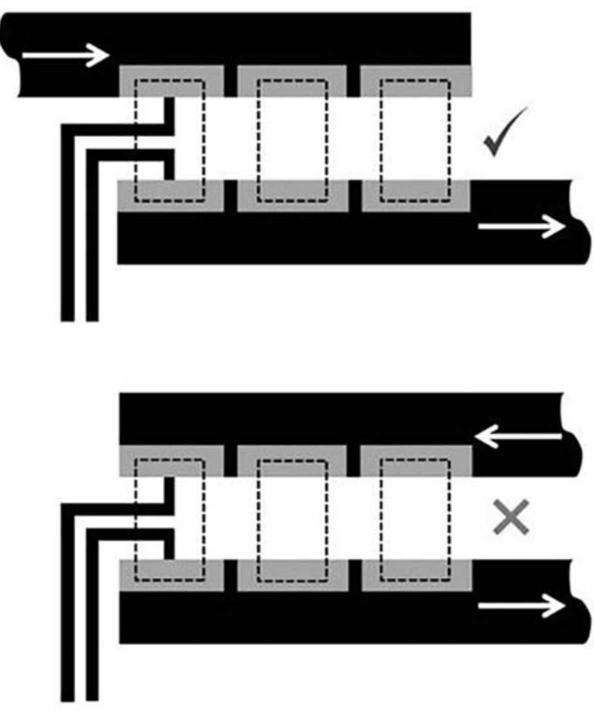

8. The total series track resistance should match for all resistors.

Minimization of Sense-Loop Area

A source of error involving high currents that are ac or changing dc is due to the voltage-sensing loop linking with changing magnetic fields. This can induce a noise signal superimposed on the desired voltage-sense signal.

To reduce the noise, the loop area contained within the sense resistor, the two voltage-sense tracks, and the sense circuit input should be minimized. This means keeping the sense circuitry as close as possible to the sense resistor and running the voltage-sense tracks close to each other.

A good way to keep these tracks very close is to superimpose them in different PCB layers. Where long track runs are unavoidable, it's also possible to use periodic vias to cross over the tracks into alternate layers. On a PCB, this replicates the effect of a twisted-pair cable, which, by means of cancellation of induced voltages, allows the circuit to withstand the effect of any changing magnetic fields that have small variations across the spatial periodicity of the twisting.

Connecting Multiple Resistors in Parallel

Sometimes, designers are forced to use more than one current-sense resistor connected in parallel, either to meet a high power or surge rating, or to achieve an ohmic value lower than the minimum available. This is problematic but possible. Resistors may be connected in parallel with voltage-sense connections made to just one of the resistors, provided the track layout ensures equal distribution of current between all resistors.

For example, the position in the current trace where the resistors are placed should be well clear of bends or constrictions that could affect the distribution of current density *(Fig. 7)*. The goal is to ensure that the total track resistance in series with each resistor should be the same (*Fig. 8),* so that the sensed resistor carries the required fraction of the total current.

Moreover, this ensures that the proportion of the total current carried by the sensed resistor doesn't vary with temperature. This would otherwise be the case with unequal series track resistances due to the high TCR of the copper PCB tracks.

Design for Heatsinking

A flat-chip resistor dissipates more than 80% of its heat by conduction into PCB tracks, making it important to provide sufficient copper area to act as a heatsink. Copper area is, for this purpose, defined as the total area directly surrounding the solder pads, including the first two squares of connected tracks. *Figure 9* indicates the general relationship between effective power rating and PCB copper area.

In region (A), there's relatively low thermal conduction through copper connected to the pads, and conduction by substrate and convection to air predominate. In region (B), the copper connected to the pads acts as a heatsink to raise the effective power rating. In region (C), further increase in copper area gives diminishing returns, as the internal thermal impedance of the chip restricts the rating.

This limiting factor of the internal thermal impedance can be lowered significantly by changing the resistor's orientation. If the terminations are formed on the longer edges of the chip rather than the shorter edges, the solder joint width is approximately doubled and the maximum distance from film center to termination is approximately halved *(Fig. 10a)*. TT Electronics' ULR2N and ULR3N *(Fig. 10b)* are examples of products that use this enhanced cooling method.

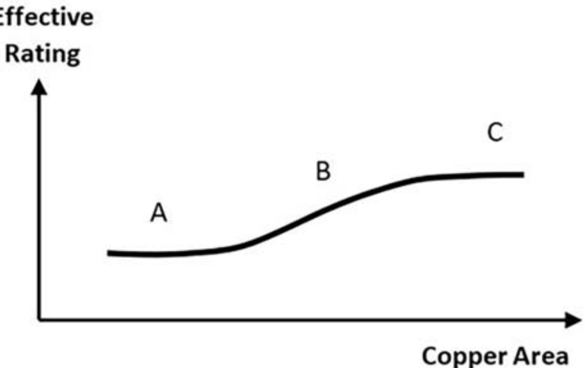

9. The general relationship between effective power rating and PCB copper area can be divided into three distinct regions.

Product Safety Test Equipment

ED&D, a world leader in Product Safety Test Equipment manufacturing, offers a full line of equipment for meeting various UL, IEC, CSA, CE, ASTM, MIL, and other standards. Product line covers categories such as hipot, leakage current, ground, force, impact, burn, temperature, access, ingress (IP code), cord flex, voltage, power, plastics, and others.

ED&D

1-800-806-6236 ▪ www.ProductSafeT.com

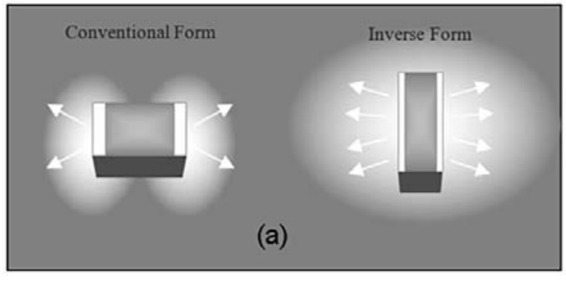

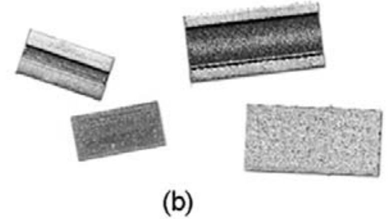

10. The limiting factor of the internal thermal impedance may be lowered significantly by changing the orientation of the resistor. If the terminations are formed on the longer edges of the chip rather than the shorter edges, the solder joint width is approximately doubled and the maximum distance from film center to termination is approximately halved (a). TT Electronics' ULR2N and ULR3N are examples of products that make use of this enhanced cooling method (b).

The resistor datasheet should contain information on the mounting conditions used to obtain the rated power. This indicates the minimum copper area that should be provided by the designer.

Summary

The growing use of sub-milliohm chip resistors for current sensing creates a spectrum of challenges for the designer and the process engineer. The component format should first be selected to support the chosen thermal-management approach, with metal-element flat-chip resistors having two terminals being the most cost-effective solution. It's then essential to design the PCB tracks and pads to meet the needs of the Kelvin connection, heat dissipation, and avoidance of induced noise.

For Part 2 of this series, which addresses additional stages such as critical assembly processes, go to *https://electronicdesign.com/21242473.*

References

1. O'Sullivan, Marcus. "Optimize High-Current Sensing Accuracy by Improving Pad Layout of Low-Value Shunt Resistors," *Analog Dialogue 46-06 Back Burner*, June 2012.

2. Folmar, Patrick. "The Truth about Placement Accuracy," *SST Semiconductor Digest.*

Attracting Tomorrow TDK

Engineered to make life easier!

CUS-M SERIES

High Performance Medical and Industrial AC/DC power supplies with low EMI

Tired of spending hours resolving radiated EMI issues? Use a power supply that actually meets class B, with margin. From 30 to 1500W, we have a product to suit your application!

- Industrial and medical safety certifications
- Industry standard footprints
- Efficiencies up to 96%
- 4kVAC isolation and low leakage currents
- Suitable for B and BF rated equipment
- Class B conducted and radiated EMI

For more information on how TDK-Lambda can help you power your unique applications, visit our web site at www.us.lambda.tdk.com or call 1-800-LAMBDA-4.

BILL SCHWEBER | Contributing Editor

Enhanced 3D-Printing Technique Yields Custom MEMS Sensors

Researchers have used a specialized form of additive manufacturing to create MEMS sensors that can be fabricated as custom, single-quantity units.

MEMS technology in its various implementations has revolutionized sensors over the past few decades, with tiny, highly accurate, low-power/cost devices for sensing acceleration, motion, pressure, and other challenging variables. There's no need to expound on that success story. Still, there's an aspect of MEMS devices that can't be overlooked: While its virtue is that it is inherently a high-volume, (mostly) silicon-based fabrication technology, these same attributes make it a poor fit for unique, single- or low-quantity, custom sensors.

Along a similar timeline, the very different technology of additive manufacturing (AM, commonly called 3D printing) has radically changed the way that custom-designed components can be built using a range of metals, polymers, and other specialty materials. With 3D printing, a design that exists on a computer screen via a CAD package can be fabricated "on the spot" and is especially well-suited for one-off or low-volume production runs.

The problem is that these two are at opposite ends of the volume and customization spectrum. But what if you could use 3D printing to create and build a one-off custom-designed MEMS sensor as needed, thus giving you the best of both worlds?

Two-Photon Polymerization

Yet that's what a team at the KTH Royal Institute of Technology in Stockholm (Sweden) has done *(Fig. 1)*. The researchers built on a process called two-photon polymerization, which can produce high-resolution objects as small as few hundreds of nanometers in size, but not capable of sensing functionality.

To form the transducing elements, the method uses a technique called shadow-masking, which works something like a stencil. On the 3D-printed structure, they fabricated features with a T-shaped cross-section that work like umbrellas. They then deposited metal from above, and as a result, the sides of the T-shaped features aren't coated with the metal. This means the metal on the top of the T is electrically isolated from the rest of the structure.

Prof. Frank Niklaus, who led the research team, noted, "This is something that has not been possible until now, because the startup costs for manufacturing a MEMS product using conventional semiconductor technology are on the order of hundreds of thousands of dollars and the lead times are several months or more. The new capabilities offered by 3D-printed MEMS could result in a new paradigm in MEMS and sensor manufacturing."

1. Shown is the finished 3D-printed MEMS accelerometer next to a 2-cent Euro coin (coin diameter is 18.75 mm/0.75 in.).
Images courtesy KTH Royal Institute of Technology

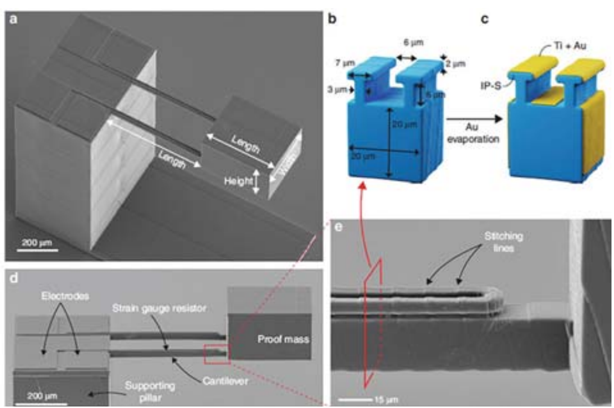

2. The 3D-printed accelerometer: (a) SEM image of the 3D-printed accelerometer structure. (b) 3D schematic view of the cantilever cross-section before metal evaporation. (c) 3D schematic view of the cantilever cross-section after metal evaporation showing the shadow-masking mechanism that enables the electrical isolation of the resistors. (d) SEM image of a lateral view of the top part of the device. (e) Close-up view of the T-shaped resistors on top of the cantilevers. The structures shown in the SEM images were coated with a thin sputtered Au-Pt layer to improve the SEM image quality.

He added that "scalability [in volume] isn't just an advantage in MEMS production, it's a necessity. This method would enable fabrication of many kinds of new, customized devices."

They printed the accelerometer structure on a glass substrate, with a supporting pillar having two single-sided, clamped horizontal cantilevers and a proof mass attached at the end of the two cantilevers *(Fig 2)*. The design freedom offered by the 3D-printing process allowed them to pattern shadow-masking structures with T-shaped cross-sections on top of the cantilevers and the supporting pillar to define the areas of the strain-gauge transducers, the electrical interconnects, and the probing electrodes.

Test Setup and Results

To assess performance, the team conducted a range of tests using a fairly standard arrangement that included a laser Doppler vibrometer *(Fig. 3)*. The piezoshaker driving frequency was swept between 1.4 and 2 kHz, with voltage amplitudes ranging from 1 to 7 V_{RMS}.

They conducted both mechanical and electrical tests, although the two aspects are related. The mechanical characterization assessed factors such as the material and stress/strain performance, which was compared to their COMSOL models (there was close agreement).

The electrical characterization assessed parameters such as relative resistance change (ΔR/R) as measured at different frequencies and driving voltages of the piezoshaker on the accelerometer *(Fig. 4)*. They also tested stability over a 10-hour period, as even a "perfect" sensor is of little use if it drifts over time (in contrast, temperature-based variations can often be accommodated or compensated, but time-related drift is more difficult to accept).

The work is detailed in their readable 13-page paper "Micro 3D printing of a functional MEMS accelerometer" published in *Nature Microsystems & Nanoengineering*. You can also read the 100-page doctoral thesis "Additive Manufacturing and Integration of 3D MEMS using Ultrafast Lasers and Magnetic Assembly" by one of the authors, on which much of the work is based.

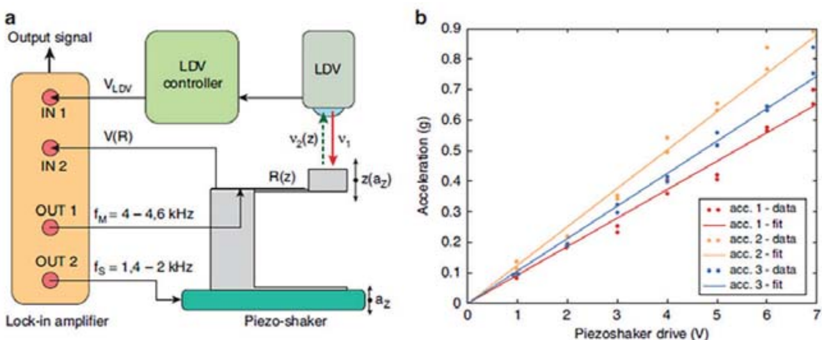

3. Measurement setup and piezoshaker calibration: (a) Schematic of the setup used to measure the responsivity of the 3D-printed MEMS accelerometer. The lock-in amplifier drove the piezoshaker, extracted the signal from the resistor through downmixing, and demodulated the laser Doppler vibrometer signal. (b) Acceleration applied by the piezoshaker was measured with the laser Doppler vibrometer on top of the supporting pillars of the three devices under different driving voltages at a frequency of 1.775 kHz. The corresponding linear fits are shown. A variation of about ±10% between the different devices can be seen.

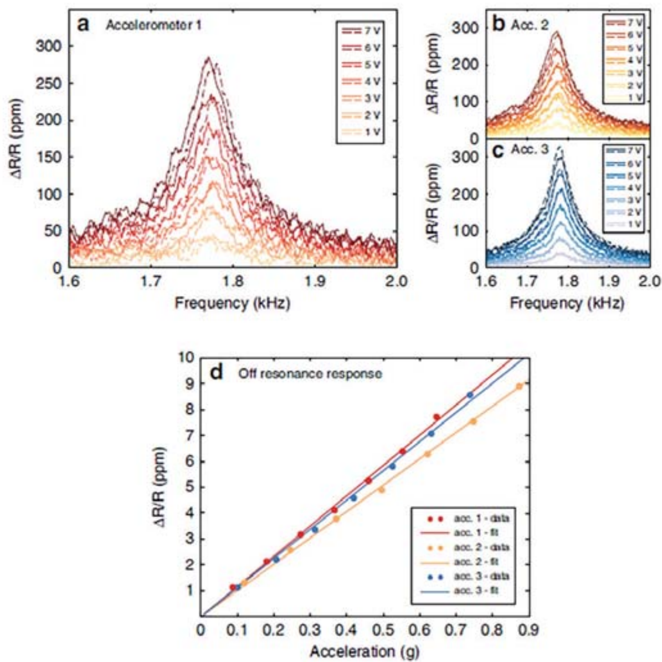

4. Electrical characterization: Relative resistance change (ΔR/R) measured at different frequencies and driving voltages of the piezoshaker on accelerometer #1 in (a), accelerometer #2 in (b), and accelerometer #3 in (c). The noticeable shift in resonance frequency is attributed to the increase in temperature in the polymer at large oscillation amplitudes. (d) Relative resistance changes of the strain-gauge transducers as a function of the acceleration applied to the MEMS accelerometers computed off-resonance.

MICHAEL STEFFES | Analog Signal Path Consultant

Three Major Design Pitfalls Plaguing **New Analog Signal-Path Designers**

Wouldn't it be great to not repeat the same amplifier application errors many new designers fall into? Read on to head off these common confusions and oversights.

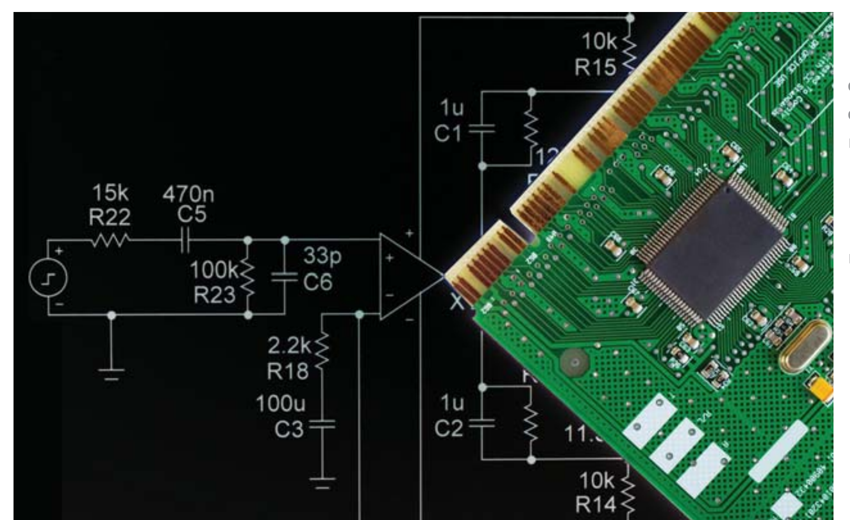

Traveling-light_dreamstime_25319773

Having been on the receiving end of designer queries from 1985 forward, there are some common oversights and misunderstandings that show up regularly. These essentially fall into three areas:

- Not considering the actual operating voltage range on the I/O or internal pins of the device being used.
- Misunderstanding the elements that contribute to an output dc offset or drift error.
- Accidentally building an oscillator (or even worse, a nominally stable design that slips over into oscillation over production and/or temperature ranges).

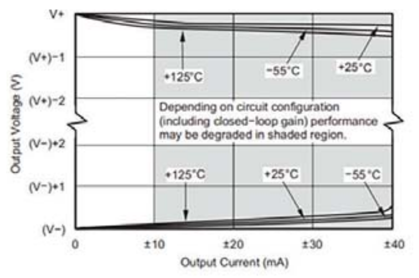

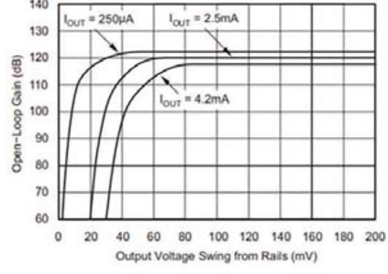

1. This example output swing "Claw" curve shows the added headroom with output current demand and the open-loop gain reduction near the supply's warning of a loss of linearity.

Texas Instruments

1. Running Into I/O Range Limits with Op Amps, FDAs, and INAs

The evolution of op amps began with very simple designs requiring considerable headroom to the supply voltages in both the input pins and the output pin. This carried over into the early fully differential amplifier (FDA) and instrumentation amplifier (INA) developments.

Over time, the need to provide more of the available supply voltage range on the I/O pins first led to rail-to-rail output (RRO) designs, then added negative rail input (NRI) designs, and more recently rail-to-rail input/output (RRIO) designs. These each come with compromises in the internals to the device.

Many single-supply designs will select at least a RRO and NRI device and then expect the device to operate with no input signal with the V+ and output pins at 0 V. All RRO devices require some small headroom to the supplies to operate linearly.

While that may be as low as 10 mV, it's still not zero. Asking the op amp to perform as expected with 0-V input will usually cause performance problems. Most NRI devices can actually operate slightly below the negative supply; therefore, a 0-V input on a single-supply design usually will not hit an "input" limit.

Another source of confusion has been the long-term evolution of how this "headroom" is specified. Early devices (still available) talk about a ground-centered maximum $\pm V_{OUT}$ swing on a bipolar supply. While accurate, it's much more useful to

think in terms of required headroom at the output (and input) to each supply voltage being used. Most early devices specified no load or a specified load for this swing.

The required headroom always increases as the demand escalates for more output current. This, and the open-loop gain reduction near the supplies, are captured in more recent op amps as shown in the curves of *Figure 1* from the OPA350 datasheet. While a true swing to ground is required in a single-supply design, some designs apply a fixed −0.23-V bias generator like the LM7705.

The first commercial FDA—the AD8138—emerged in 1999. Subsequent developments have pushed toward extreme dc precision (and speed with low power) in mainly NRI and RRO designs like the THS4551.

One common confusion when applying these modern FDAs is that a single-supply design can, in fact, take a dc-coupled bipolar input and operate all I/O pins with enough headroom on a single supply. The key here is that the common-mode (CM) control loop will force a dc level-shifting current back through the input resistors to level-shift the input CM voltage across the two inputs to operate above ground, even with a bipolar input signal. This effect is illustrated in *Figure 2.*

Any FDA circuit can reduce the input networks to Thevenin equivalents as shown in *Figure 2.* A good design will have equal feedback resistors and equal Thevenin impedances looking back from the two inputs to a source and ground (or low-impedance reference voltage).

The easiest way to see that the input CM voltages are above ground is to consider the lower output side of the FDA in *Figure 2* dividing back to ground. If the output is correctly swinging ±0.5 V on each side around the stated 0.95-V CM voltage, the 0.45- to 1.45-V absolute swing on that lower output will divide back to the lower input pin as 0.177 × 0.45 V to 0.177 × 1.45 V equal to 80 to 256 mV.

Yes, the input CM voltage moves with the full-scale swing of the input signal but never goes below ground. Actually, since the outputs can't go below ground, that feedback signal to the lower summing junction can't swing below ground. The FDA differential loop forces the error voltage across the inputs to zero. Thus, the input pins move together for a single-ended input to differential output application.

A very popular solution in precision industrial applications is the instrumentation amplifier (INA). These typically present two high-impedance inputs with a settable differential gain to a single-ended output stage. Such an output often includes a reference voltage input that independently sets an output dc level separate from the input-signal-induced output swing.

Those all have specified input and output headrooms much like an op amp or FDA device. They also often have internal swing limits not directly observable in application or simulation. These hidden limits have tripped up many a design engineer.

Several INA suppliers have developed tools to expose these limits in application. One is the instrumentation-amplifier diamond plot tool. This tool allows designers to enter their intended input conditions and a desired gain and reference voltage, along with a candidate device, and immediately expose internal and external clipping issues.

Figure 3 shows an example drawn from an actual thermocouple design where the input CM voltage is fixed at 2.048 using a single 5-V supply on the LT1789 INA. If the red line in the diamond plot is completely within the white area, unclipped operation is assured.

2. Designer Oversights in Assessing Output DC Precision and Drift

The calculations for output dc error and drift are well-trodden trails in academic and vendor material. Several detail issues continue to trip up new (as well as

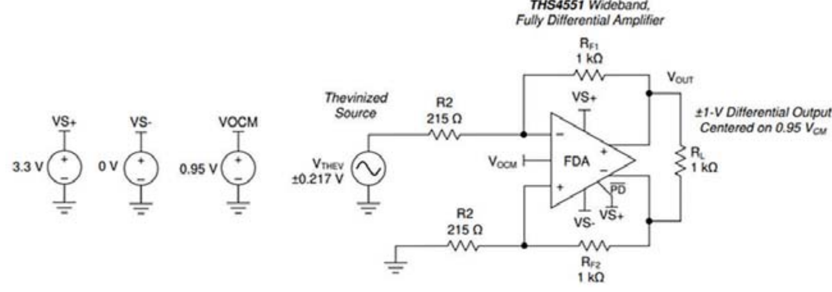

2. Input common-mode swing analysis for a single-supply dc-coupled application example.
Texas Instruments

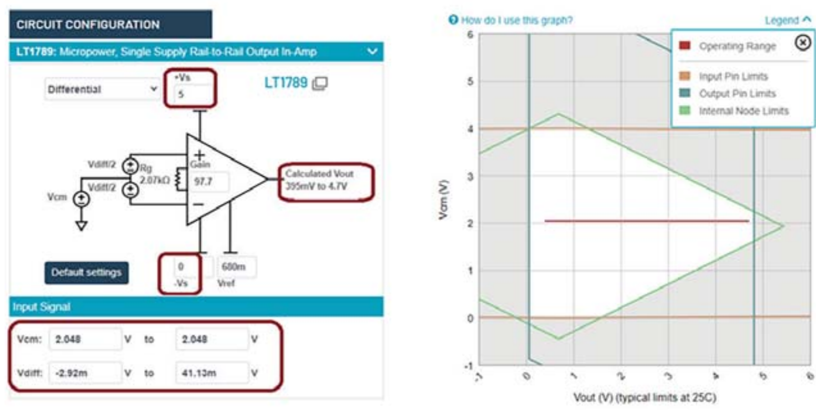

3. Diamond plot tool setup for a thermocouple application showing valid operating range.
Analog Devices

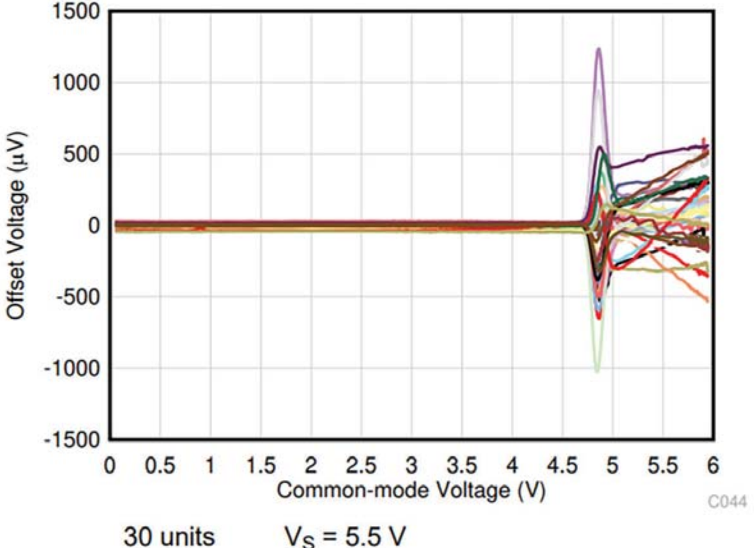

4. The input crossover network shows a large offset voltage step near the positive supply for the OPA396 precision CMOS RRIO op amp. *Texas Instruments*

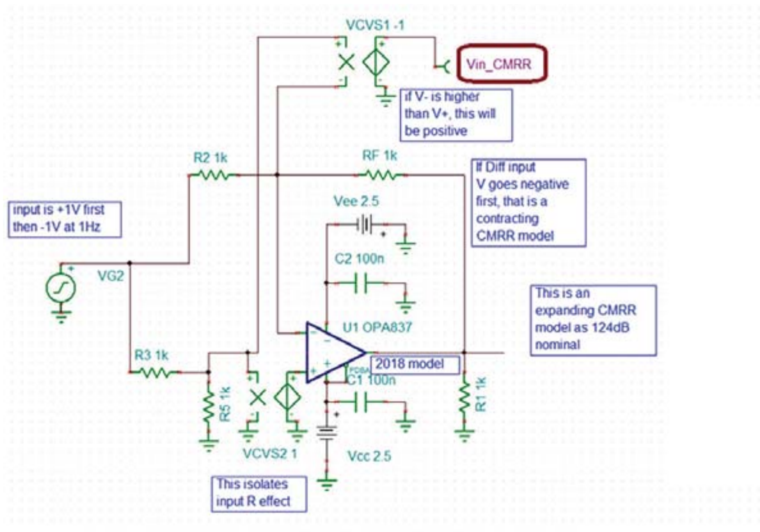

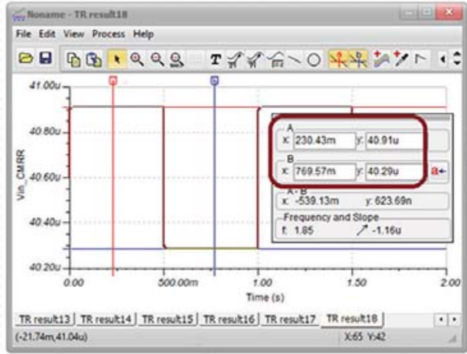

5. Simulation to test the effect of input CM swing on input offset voltage in the OPA837 model. *Texas Instruments*

experienced) designers with the vast proliferation of op-amp and FDA solutions.

Classic bipolar input op amps and FDAs usually offer a well-matched input bias current error if it's a voltage feedback amplifier (VFA). It's effect on an output dc error can be reduced using a "bias current cancellation" resistor solution to reduce the output dc error to the offset current at the inputs (mismatch specification) times the feedback resistor value. For voltage-feedback op amps, this simply requires you to match the dc impedance looking out the V+ pin to the parallel combination of the feedback and gain resistors on the inverting side.

But where does this actually work—and not work? It will always work with simple NPN or PNP input stages having matched bias currents. Some very-low-bias current bipolar input devices use cancellation currents into the input pins. If so, the offset currents aren't as low as for the simpler input-stage designs. Bias currents are never matched for CMOS or JFET input devices, so designing for bias current cancellation is a waste of time; lower R's on the V+ pin are usually desirable to reduce added noise from those resistors.

Some of the very lowest input offset voltage and drift VFAs emerged with the chopper-input types of devices. These chopper-, and trimmed non-chopper-input, CMOS devices provide sub-10-µV input offsets with very low drift. Later developments added rail-to-rail input (RRI) options using crossover networks at the input to pass control between the CMOS device types. Zero-crossover RRI types include an on-chip boost regulator to provide enough supply voltage for the input stage to get RRI without a crossover network, like the OPA328.

Those types of RRI devices with a crossover region will show a discontinuity in the input offset voltage as control is passed across the CM input range of the op amp. Many designers have been tripped up by this, where simply avoiding that area of the input CM might have been possible.

Figure 4 shows a good example from the recent OPA396 RRIO CMOS precision op amp, a non-chopper device quoting a

maximum input offset of 100 µV. This gain-of-1 plot clearly shows the large step in input offset voltage near the positive supply. This is easily avoided by operating with a small non-inverting gain or running inverting mode with fixed bias on the V+ pin well below this crossover.

The very best input-drift VFA op amps use a chopper-input structure. Those intrinsically need an internal switching clock that then shows up in the input current noise spectrum. Though this often isn't shown, it's usually there. Whether this affects the accuracy in the application depends on many things, but at minimum it's prudent to plan on at least a post-RC filter well below that chopping frequency to filter that off.

It's also prudent for chopper-input op amps to design for source matching as in dc bias-current cancellation. This will reduce the higher-frequency output noise due to the chopper-input current spikes. Some, but not all, chopper-input op amps report that chopping frequency.

CMRR and PSRR

The earliest op-amp literature spent quite some time discussing the common-mode rejection ratio (CMRR) and power-supply rejection ratio (PSRR) effects on output error terms. Those usually end up showing a plot over frequency that's almost always a designer simulation as the measurement is nearly impossible.

Here, only the dc values are of interest for output dc error concerns. The PSRR gets confused in the datasheets, sometimes showing the supplies moving together—but that's the same thing as a CMRR test. ATE flows move only one supply at a time to extract out an apparent shift in the input V_{os} voltage. These are often assumed to have a bipolar distribution in adding to the other dc error terms to get full output dc error band.

For typical single-supply designs with say a ±10% supply tolerance for a +5-V design, such an error term for modern devices is very small. Typical PSRR numbers are 110 dB or greater, so ±0.5-V supply shift in production maps to a ±1.6-µV expansion in the input offset span using a 110-dB specification.

CMRR has been presented as a shift in the input offset voltage as the CM input voltage travels across the available input span. In fact, all models and ATE data show this as a gain error term. Since the error depends on the input CM level, why would it be a static dc error when in fact it's more like the LG/(LG+1) gain error (where LG is the loop gain, the Aol/(noise gain)). Often, this CMRR gain error is on same order or smaller than the Aol at a gain of 1, and it becomes even less significant at higher noise gains as that LG term becomes the dominant gain error.

A simple simulation *(Fig. 5)* can easily illustrate what the model is producing.

Here, the precision OPA837 is set up with four equal resistors in a classic CMRR test. The output should be very close to zero swing, but here the input offset voltage is probed showing a very small 0.62-µV p-p amplitude square wave (around the nominal 40.6-µV offset voltage in the model) for a 1-V p-p CM input swing at the V+ input pin. Here a dependent unity-gain voltage buffer was inserted to isolate the V+ pin input resistance from the four-equal-resistor test setup.

The polarity indicates this model is showing a very small expansion in the gain due to CMRR effects. This 124-dB CMRR level will in practice combine with any gain error introduced by the LG/(LG+1) term. It's not clear that this expanding gain effect in the OPA837 model is matching the physical device.

Both expanding and contracting CMRR effects can be found using different op-amp models in the test circuit of *Figure 5*. Holding a fixed (non-zero) input CM voltage (as in an inverting op-amp design) will add a fixed CMRR error contribution to the total input V_{os} calculation.

3. Ignore Nominal Design Phase Margin at Your Peril

Once we have the I/O ranges satisfied and the output dc error band estimated for particular candidate solution device, the actual functional design can proceed. Many different implementations and applications can call upon the vast range of op amps and FDAs for numerous end applications.

With a schematic and maybe even a layout developed, do you know your phase margin? Perhaps you should. Op amps have always had the risk of instability. It's been exacerbated by the more aggressive designs in recent years trying to deliver the most performance at the lowest power.

For instance, the common RRO stage designs come with a very reactive open-loop output impedance *(Fig. 6).* Hopefully this is in the simulation model. Often, it's a little uncertain if this critical feature to loop phase margin is correctly captured by the model.

A circuit that's already oscillating has one set of bench tools to isolate down to the suspect device. It's far more prudent to attempt a phase-margin simulation prior to board build to head off any problems. That does, of course, depend on good simulation models, and those have been improving. Still, they come with a variety of pitfalls across the industry.

There are several easy techniques to extract the loop phase margin from an amplifier schematic. If possible, any layout and source impedance parasitics should be added to the simulation, and by all means the intended load has to be there—even if it's only a parasitic RC of the next device. Essentially, the simple techniques break the loop in some way, inject a small

Keep Your Cool!

PFH SERIES

500 Watt AC-DC Power Module

Keeping your system cool just got easier. TDK-Lambda's next generation AC-DC power modules offer efficiencies of up to 92%, helping you to reduce both power consumption and heat sink size.

The PFH500F offers more intelligence too! A PMBus™ interface can program the output voltage and fault management settings, with read back capability on actual voltage, current and operating temperatures. Don't sweat the details, contact TDK-Lambda for an evaluation board or check our website for distribution inventory.

https://product.tdk.com/en/power/pfh

- Compact Metal 4" x 2.4" Case
- 85 to 265Vac Input
- Up to 92% Efficient
- Parallel Operation Optional
- Operation from -40 to 100°C
- Read/Write PMBus™ Communication

For more information on how TDK-Lambda can help you power your unique applications, visit our web site at www.us.lambda.tdk.com or call 1-800-LAMBDA-4.

signal test signal in the loop, and trace the gain and phase around the loop to assess phase margin where the loop gain magnitude goes to 0 dB (or 1 V/V).

If we think about the transfer function having that LG/(LG+1) term, it can be rewritten as 1/(1+(1/LG)). The LG has a gain and a phase-shift component. If it drops close to a 1 magnitude (0 dB crossover) near where the phase shift is approaching 180 degrees around the loop, it becomes a 1/(1-1) term, which may result in sustained or intermittent oscillation. This can cost lots of man-hours and re-spin dollars when a little bit of simulation time could have headed off this pain and suffering.

Even simple circuits can run into phase-margin problems *(Fig. 6, again)*. Here, a simple inverting gain of –1-V/V design added a feedback capacitor to bandlimit the signal channel to 1 MHz. A small parasitic capacitive load, along with that feedback C_f, interacts with the reactive open-loop output impedance to

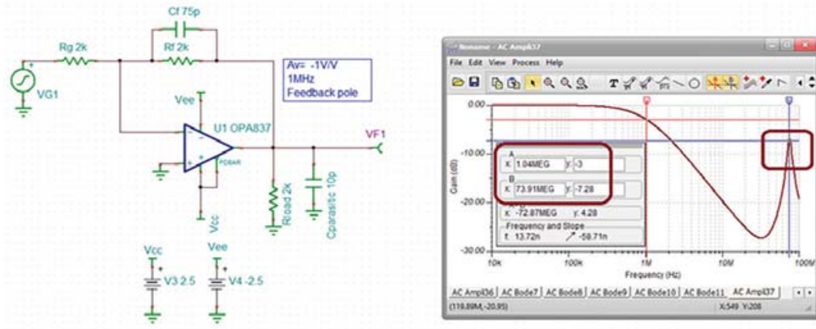

6. Inverting gain of –1 V/V with the OPA837 shows marginal stability with the feedback bandlimiting capacitor. *Texas Instruments*

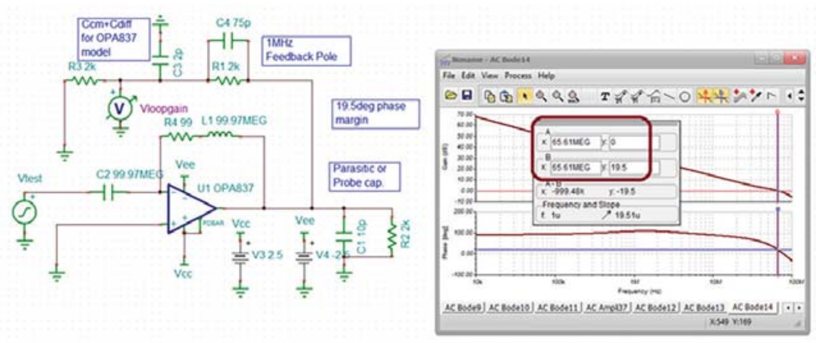

7. One possible loop-gain phase margin simulation setup shows only 19.5-degree phase margin for the circuit of Figure 6. *Texas Instruments*

HRi RANGE
EZi RANGE
BBi RANGE
HARWiN
CONNECT TECHNOLOGY **WITH CONFIDENCE**
// WWW.HARWIN.COM

cause the peaking at 74 MHz. This is a warning that the circuit might go unstable in production.

To run a LG phase-margin simulation, it's necessary to first establish a good dc operating point for all nodes in the circuit. Older approaches found the exact input offset voltage to add to an open-loop circuit to zero the output-pin voltage. That works, but it's much easier to use simulation tricks of impossibly high L and C elements to do this job for us, as shown in *Figure 7* using the OPA837 model again.

The large feedback inductor closes the feedback loop at dc, then immediately opens up on the first frequency test step. The large input capacitor is open at dc, then immediately shorts out to apply the test signal on the first frequency step.

This approach requires you to manually add the op-amp input impedance at the loop gain measurement point (2 pF here). The measurement meter is rotated to report phase margin directly for this setup. Looking for the 0-dB gain point around the loop and then the phase margin at that same 66.61-MHz frequency shows only 19.5 degrees. This would require some attention where several approaches (and this phase-margin simulation approach) are detailed in the OPA396 datasheet.

What your minimum target phase margin might be depends on your circuit and the device you're using. Many older devices (National Semiconductor in particular) targeted a nominal 45 degrees and just took the peaking that results from it. More modern devices feature a nominal phase-margin target around 60 degrees to get close to a Butterworth closed-loop response. As a rough guideline for most circuits:

- Phase margin >30 degrees is probably okay if the intended circuit operation is acceptable.
- Phase margin between 20 and 30 degrees, if easy to do, should be improved to >30 degrees.
- Phase margin <20 degrees probably should be raised to at least the mid 20s.
- Phase margin <10 degrees—you should never go to production like this; it absolutely needs attention.

How sensitive a design is part to part and over temperature variation really depends on the circuit and devices chosen. Older op amps and FDAs have a wider spread on their open-loop gain and phase where more modern devices (especially those with supply current trim) are much better and will have far lower risk of large dips in phase margin over production.

Keep these three hazardous areas in mind as you set out to apply the vast range of op amps, FDAs, and INAs to your design. ed

DC-DC Converters

0.25W~150W
1" x 1"/ 2" x 1"/Half-Brick
SIP/SMD/DIP Package
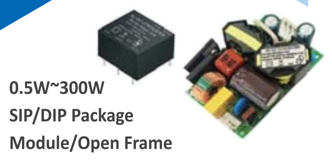
0.5W~300W
SIP/DIP Package
Module/Open Frame
AC-DC Converters
NEW PRODUCT
Ultra Compact Size
YTB Series
10~15W
24.0 x 14.0 x 9.5 mm

1xN,2xN Port
1G/2.5G/5G/10G Base-T with PoE+
LAN Filters
RJ45 Connectors
10:1 Wide Input Range
27DX-6W Series
SIP8 Package
1500 Vdc Isolation

High Power Density
YKB Series
15~25W
DIL24 Package

Inductor
Pulse Transformer
Telecom Transformer
High Frequency Transformer
Transformers
8:1 Wide Input Range
YNC Series
20~40W
2" x 1" Size

No.61, Keji 3rd Rd., Annan Dist., Tainan City 70955, Taiwan
TEL:+886-6-3842899 / FAX:+886-6-3843288
E-mail : ed@yds.com.tw / Website : www.yuandean.com

Broadcom Optocouplers

Certified to IEC 60747-5-5 Safety Standard with Proven High Voltage Lifetime Performance

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

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

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

For more information, visit **broadcom.com/automotive**.

Copyright © 2021 Broadcom. All Rights Reserved. Broadcom, the pulse logo, and Connecting everything, and R²Coupler are among the trademarks of Broadcom. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. 010521

CABE ATWELL | Technology Editor

Dreamstime | Pavel-Chagochkin_151097024

What's the Difference Between **CAN and SPE in the Automotive Industry?**

The CAN bus protocol is being phased out of the automotive industry in favor of single-pair Ethernet due to its increased data bandwidth, node efficiency, security, and more.

It's been 36 years since the CAN (controller area network) bus was released by the SAE (Society of Automotive Engineers). Lying at the heart of vehicle communications for decades, it supports a wide variety of automotive innovations.

The CAN bus is described as a vehicle bus standard that allows microcontrollers and devices to communicate with each other's applications without a host computer. It's a message-based protocol, designed originally for multiplex electrical wiring within automobiles to save on copper, but it also can be used in many other applications.

Automotive technology has come a long way since the first CAN buses were deployed, and the platform has since been tasked with more and more functions beyond what was envisioned in the 1980s.

CAN remains a favorite of auto manufacturers even into its fourth decade. That said, the automotive industry has been undergoing a paradigm shift in response to cutting-edge technologies and fast-evolving consumer demands. CAN's long-running reign is set to face new challenges.

To that end, the automotive industry is looking toward single-pair Ethernet (SPE) to function as the automotive network's backbone, an alternative that brings higher performance, increased security, and increased efficiency over CAN buses.

According to a 2020 market report, the global connected-car market is expected to reach $225.16 billion by 2027, up from

TEST EQUIPMENT THAT SIMPLIFIES THE ADVANCEMENT OF HUMANITY.

www.dewesoft.com | sales.us@dewesoft.com | +1-855-339-3669

FLEXIBLE

MEASURE ANYTHING
ANYWHERE QUICKER AND
FASTER THAN BEFORE

POWERFUL

NANO VOLTS TO KILO VOLTS IN
THE SAME ACQUISITION

RELIABLE

GUARANTEED QUALITY
GUARANTEED ACCURACY

$63.03 billion in 2019. This shift toward increased connectivity will play a decisive and accelerating role in the move to SPE networks, even as CAN buses continue to provide an important communication medium (primarily for legacy components).

CAN Bus Advantages and Limitations

CAN's technical features have made it effective in the vehicular environment, starting with its high tolerance for noise that's supported by CAN's physical layer and protocol.

Moreover, CAN supports native multicast and broadcast, provides built-in frame priorities, offers non-destructive collision resolution, is 100% distributed in operation, and easily supports long single buses measuring in the tens of feet. With only a pair of wires and a cheap physical-layer interface, CAN still remains an attractive option in traditional automotive networking.

The CAN bus is limited to a maximum data-transmission rate of 8 to 10 Mb/s, making it challenging to keep up with the demands of current and next-gen connected vehicles. This means that as more devices become connected to the same bus, its performance is diminished.

Manufacturers have mitigated that issue by incorporating more CAN buses into vehicles, with some packing 10 or more. The ad hoc connections between these buses have been sufficient to facilitate network communications, but security is typically an afterthought, despite connected vehicles' increasing vulnerability to cyberattacks.

Upgraded CAN versions, including CAN FD and CAN XL, have improved on the basic design by supporting more throughput (from 1 Mb/s in traditional CAN to up to 10 Mb/s in CAN-FD/XL), faster line rates, and larger frame sizes. CAN FD is currently in production, while CAN XL remains in the design phase.

Nonetheless, CAN offers some features that manufacturers still find attractive, including its ability to be implemented into a single microcontroller, significantly reducing the amount of software needed to establish communication. It also can detect and recover from an error condition within 23 µs at 1 Mbit/s, while TCP/IP at 100 Mbits/s is limited to a response time of about 1 ms (a factor of roughly 44).

CAN also provides collision-free and predictable arbitration to manage network access between competing nodes. Furthermore, the bus can be implemented in industries outside the automotive field, including as a fieldbus in general automation environments, primarily due to the low cost of some CAN controllers and processors.

While these improvements will certainly enhance CAN's robustness and reliability as cars become more connected and autonomous, SPE offers another pathway to building a vehicle communications network that meets modern needs.

The next generation of standard, semi-autonomous, and fully autonomous vehicles are switching to SPE from CAN bus systems for increased data throughput, security, and efficiency. *Pexels*

SPE's plug-and-play capabilities are suited to the high-performance, service-oriented environments that will define next-gen vehicles.

Single-Pair Ethernet and Its Advantages

Meeting the demands of next-generation connected and autonomous vehicles will likely require code that's hundreds of millions of lines long, in contrast to the 50-100 million lines for current models. Thus, the industry is transitioning to SPE. It uses a pair of copper wires that can transmit data at speeds of up to 1 Gb/s over short distances while simultaneously delivering Power over Dataline (PoDL).

SPE's plug-and-play capabilities are suited to the high-performance, service-oriented environments that will define next-gen vehicles. Devices can be connected and disconnected in real-time, with zero downtime, providing a significant advantage over CAN buses.

SPE also offers distinct advantages that necessitate integrating into a mixed-criticality automotive architecture. By supporting quality of service, security,

time-critical extensions (time-sensitive networking, or TSN), and increased performance, SPE enables the high-performance communication required in zonal architecture.

Furthermore, it facilitates identifying and troubleshooting problems within in-vehicle systems. On top of that, SPE provides smaller wiring harnesses, resulting in reduced investments of time and resources than the installation of often-ungainly CAN systems.

SPE operates at full duplex and has a maximum reach of about 50 ft. (100BASE-T1, 1000BASE-T1 link segment type A) or up to 130 ft. (1000BASE-T1 link segment type B) with up to four inline connectors. Both physical layers require a balanced twisted pair with an impedance of 100 Ω. The cable must be able to transmit 600 MHz for 1000BASE-T1 and 66 MHz for 100BASE-T1. In addition, 2.5 Gb/s, 5 Gb/s, and 10 Gb/s over a 1-m single pair is standardized in 802.3ch-2020.

Transitioning to service-based communication and modular design are supported by Ethernet-based architectures rather than CAN buses, allowing for greater flexibility. This will be ideal as those hundreds of millions more lines of code are added to future car models, along with new infotainment capabilities, safety-critical features, over-the-air (OTA) updates, and autonomous and semi-autonomous features.

Moreover, SPE enables consistent data transmission up to the field level. This means that only one pair is required to transmit the signals instead of the previous two or four pairs, which benefits the automotive industry and its requirement profile. Additional advantages include the use of thinner cables, lowered costs for assemblies, less space and weight requirements, smaller bending radii, and fewer materials needed for manufacturing.

Conclusion

According to a report from Ixia, there are more than 400 million automotive Ethernet ports than for all other interfaces combined. Though SPE is still in its infancy, it's expected to command a market share of over $4.3 billion by 2024 due to increased demand for next-gen automotive features such as infotainment systems, advanced driver-assistance systems (ADAS), sensors, cameras, and more.

As technological innovation continues at its current pace and consumer demands grow in tandem, a wide range of new applications and connected features will be integrated into vehicles. These include everything from more sophisticated ADAS to in-car gaming consoles. Integrating those new capabilities will require more bandwidth, smaller hardware, and increased data throughput, which is why more automakers are currently looking at SPE to handle those loads and more.

SCOTT BEEKMAN | Vice President, Memory Business Unit, Kioxia

11 MYTHS About 3D Flash Memory

Think you know all about 3D flash memory and what the future has in store for this technology? Keep reading...

BiCS 3D flash technology was introduced by Kioxia in 2007. *Thinkstock and Kioxia*

Thirty-five years ago, KIOXIA invented flash memory. Did you know the company conceptualized 3D flash memory as well? It happened at a 2007 IEEE VLSI symposium. Today, 3D flash memory, which utilizes a process that vertically stacks flash-memory cells, has come to dominate the industry.

From smartphones to automobiles, data centers, and more, the demand for more memory storage, higher capacities, and improved performance will continue to grow. That means advances in 3D flash memory will need to keep up.

This article explores 3D flash-memory technology, with a focus on the challenges related to its continued evolution. It also illuminates some of the misconceptions surrounding this popular technology, some of which may surprise you.

1. Adding more 3D layers improves performance.

From a performance perspective, there's no inherent benefit to adding more layers to 3D flash memory. In fact, adding more layers to 3D flash memory increases design complexity.

As additional layers are incorporated, the string of memory cells gets longer, and the challenges of etching a uniform memory hole over the entire depth of layers becomes more difficult. Any inconsistencies in the uniformity from the top layer to the bottom layer can negatively impact the electrical characteristics of the device, which in turn may negatively impact performance.

2. Storage-class memory has no future.

While storage-class memory is still in the early stages of market development, it's clear that the storage and performance demands of real-time "big data" analysis continue to increase—driven, for example, by edge computing, autonomous-driving learning networks, and artificial intelligence. This is pushing the need for more DRAM for fast-response analysis as well as for more storage in general, which flash supports at lower cost but with less demanding response.

Storage-class memory, such as KIOXIA's BiCS flash-based XL-FLASH *(see figure)*, bridges the performance gap between DRAM and flash memory. With significantly lower latency than traditional flash memory and much lower cost relative to DRAM, storage-class memory enables new cost-effective storage solutions that support increasing memory storage requirements.

3. Adding more 3D layers reduces production cycle time and improves throughput.

As the memory hole depth increases as more layers are added, it takes longer to etch these memory holes, therefore negatively impacting production cycle time and throughput. Making 3D flash memory with more layers is a longer process, everything else being equal. However, memory suppliers can adopt one of several strategies to try to reduce cycle time, including utilizing different methods to bond the memory-cell array to the CMOS circuit.

4. Adding more layers is the only way to reduce cost per gigabyte.

As layer counts continue to increase, it will diminish the reduction in cost per gigabyte that generally results from

adding more layers. Fortunately, there are other ways to reduce cost, including boosting the density per layer by increasing the memory hole density per layer; moving the peripheral circuity, or manufacturing the CMOS circuit in parallel with the memory cell array before being bonded together (CBA); adding more bits per cell, etc. There are a number of different, proprietary ways suppliers can squeeze out costs.

5. Quad-level-cell (QLC) flash memory can't meet the reliability demands of today's applications.

The original NAND flash memory was single-level cell (SLC), which stores one bit per cell. Then came the introduction of two bits per cell (multi-level cell, MLC). Though it provided greater storage density and lower cost per bit, there were concerns about whether the reliability would be enough to meet the needs of applications at that time.

As more bits are stored per cell, it typically will negatively impact the write erase endurance. However, improvements were made to the controllers used to manage the flash memory, and MLC became the mainstream solution.

Today, 3D flash memory utilizing three bits per cell (TLC) is the dominant flash-memory solution used in the market, with designers exploiting its high densities and lower cost per bit. Likewise, QLC is in its early stages of growth and, once again, applications that demand high densities and lower cost per bit will take advantage of this new solution. Different applications will weigh cost, density, reliability, power, and performance to find the ideal flash solution for their use case.

6. 3D performance and reliability isn't as good as 2D floating gate.

3D flash memory employs a different cell architecture than 2D floating gate that improves reliability and performance. 2D floating-gate memory utilizes a trench cell architecture, whereas 3D flash memories typically use a charge trap cell architecture that's better at reducing the leakage of stored electrons.

For that reason, when 3D flash memory was introduced, primarily supporting TLC, it generally had the same reliability as the floating-gate memory solutions it displaced that supported MLC. Advances in 3D flash memory continue to enable increased performance generation to generation.

7. 3D flash memory doesn't support extended temperature ranges.

3D flash memory does, in fact, support extended temperature ranges. Based on application requirements, the general categories of temperatures supported are commercial, industrial, and automotive. Commercial grade is the least stringent and most typically used temperature range, while automotive is the most stringent.

Depending on the type of flash product and application use case, the temperature supported by flash memory can range from 0 to +70°C on one end of the spectrum to –40 to +105°C in the case of automotive.

8. Adding more 3D layers enhances reliability and improves yields

For the same reasons that can negatively impact performance when additional layers are incorporated, they also can affect reliability and yields. As more layers are added, NAND flash vendors must overcome these design and manufacturing challenges at each generation, which will continue to become more challenging as layer counts increase.

9. All 3D NAND flash is similar.

A variety of 3D flash memories are well-suited for multiple uses cases and applications. While a range of flash-memory solutions support different temperature ranges, different types of 3D flash memory support different levels of reliability and performance.

Often tradeoffs occur when optimizing flash-memory designs to enhance reliability, performance, power, or cost. This is why, for example, we see different flash-memory products or SSD solutions to support client, data-center, or enterprise markets.

10. 3D flash memory will completely replace 2D floating-gate memory.

While 3D flash memory has clearly become today's dominant flash-memory architecture, certain applications still use legacy 2D floating-gate flash memory. For example, applications that only need small densities of flash might use SLC flash memory supporting densities such as 1, 2, or 4 Gb of flash. Or applications that might only need 4 or 8 GB of MLC floating-gate-based eMMC.

On the other hand, the latest generations of 3D flash memory are typically produced in die densities of 256 or 512 Gb, or more, to achieve the best cost per gigabyte based on how they're designed and manufactured. The minimum die density supported will likely increase in future generations of 3D flash. This means that a whole range of flash-memory generations continue to be in demand based on uses cases and density requirements.

11. At each generation, it's always best to have more layers.

The optimum number of layers at each generation may be different for each supplier. This is why we see some suppliers have, say, 64 layers vs. 72 layers, or 92 layers vs. 96 layers. And we'll likely see more variety as the layer counts rise.

Though each new layer adds incremental costs, it also boosts density per wafer, which has a positive impact on the bottom line. Since the capital costs for each new generation of 3D flash memory are significantly higher than that of 2D floating gate, suppliers also need to consider how to squeeze the most out of each generation with the least amount of additional capital investment. For this reason, the optimum layer count can vary between suppliers at each generation. ed

CABE ATWELL | Technology Editor

108780810 © Lequint | Dreamstime.com

What are the Differences Between **MicroSD and MicroSD Express?**

Speed, formats, and storage capacities are key factors that separate a microSD card from a microSD Express card.

MicroSD and microSD Express are both types of flash-memory cards that are used to store data in electronic devices, such as smartphones and cameras. The main difference between the two is their transfer speeds. MicroSD (Secure Digital) cards have a slower transfer speed, while microSD Express cards are much faster. MicroSD Express cards also use the PCI Express (PCIe) or NVM Express (NVMe) interface, which is faster than the standard SD interface used by microSD cards.

The microSD format, first introduced by SanDisk in 2005, was initially known as T-Flash, then TransFlash, before being named microSD when it started to be used by the SD Card Association (SDA). The first microSD card had a 128-MB capacity and was about the size of a fingernail. Since then, the capacity of microSD cards has dramatically increased, and it's now possible to purchase cards with capacities of up to 1 TB.

The small form factor and high storage capacity of microSD cards have made them popular for use in mobile devices. They're now extensively used in smartphones, tablets, and other portable devices. MicroSD cards also have been widely adopted in embedded systems, IoT, and other connected devices.

The microSD card was initially slow to gain traction, as mobile devices at the time had limited storage capacity and used embedded memory. However, as mobile devices and digital cameras became more popular, the demand for removable memory increased, and the microSD card became the go-to memory card for many consumers.

The card has gone through several iterations over that time, increasing its capacity to 2 TB in 2020. While it's widely used in mobile devices, it's also employed in many industrial and medical applications, such as sensors, GPS systems, and medical imaging.

MicroSD Formats and Features

Every device, from a camera to a drone, has different SD card requirements. For example, a drone may require a microSD card with a high-speed class rating to write data very fast. However, if the card has a lower rating that can't write fast enough, it may not be compatible with the drone. With that in mind, microSD

1. MicroSD is the most widely used storage media for mobile devices due to its low cost and high storage capacities. *Wikimedia*

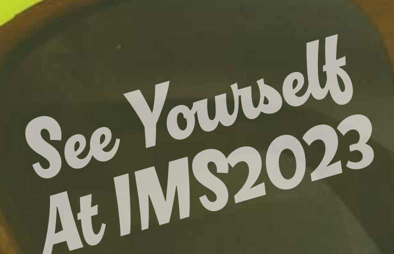

IMS

Register Now!

IEEE MTT-S International Microwave Symposium
11-16 June 2023
San Diego California

cards are available in various formats and storage capacities *(Fig. 1)*, including:

- SD (Secure Digital): The original format, with a capacity of up to 2 GB.
- SDHC (Secure Digital High Capacity): An updated format with a capacity of 4 to 32 GB.
- SDXC (Secure Digital Extended Capacity): A further updated format with a capacity of 64 GB to 2 TB.
- SDUC (Secure Digital Ultra Capacity): The latest format, with a capacity of 2 to 128 TB.

Furthermore, microSD cards may have different features, such as:

- Speed class: Indicates the minimum write speed of the card, with classes ranging from 2 to 10. The higher the class, the faster the card can write data.
- UHS (ultra-high speed): Indicates the card's maximum read and write speeds, with classes ranging from 1 to 3.
- A1 or A2: Indicates the card's performance level for running apps on a device. A1 cards are recommended for most apps, while A2 cards are better suited for high-performance apps.
- Video Speed Class: Indicates the minimum write speed of the card for recording and playback of video, with classes ranging from V6 to V90.

MicroSD card speeds are measured in terms of read and write speeds, which are typically expressed in megabytes per second (MB/s) or kilobytes per second (kB/s). The maximum read speed of a card is usually faster than its maximum write speed. Some common speed classifications for microSD cards include:

- Class 2: Minimum write speed of 2 MB/s
- Class 4: Minimum write speed of 4 MB/s
- Class 6: Minimum write speed of 6 MB/s
- Class 10: Minimum write speed of 10 MB/s
- UHS-I: Minimum write speed of 10 MB/s, maximum read speed of 104 MB/s
- UHS-II: Minimum write speed of 30 MB/s, maximum read speed of 312 MB/s
- UHS-III: Minimum write speed of 60 MB/s, maximum read speed of 624 MB/s

In addition to speed classifications, microSD cards can have different power requirements. Some cards require a higher voltage to operate, while others may have lower power requirements. A card's power requirements will be specified by the manufacturer and should be considered when choosing a card for a particular device. Certain devices may be unable to supply the necessary power to run certain cards.

It's also worth noting that a high-speed card doesn't mean it will perform well in all situations. The device that the card is being used in must support the high speed, too.

2. MicroSD Express offers increased speeds, via the NVMe and PCIe bus standards, over traditional microSD storage devices. *SD Association*

MicroSD Express Formats and Features

MicroSD Express *(Fig. 2)* is a new generation of microSD cards that uses the NVMe protocol to achieve faster read and write speeds. First announced by the SD Association in June 2018, the technology is based on the PCIe and NVMe standards. The first microSD Express cards were released in 2019 by various manufacturers such as Lexar, Samsung, and Sandisk.

These cards typically have read speeds of around 985 MB/s and write speeds of around 950 MB/s, making them much faster than traditional microSD cards and their maximum read speed of around 100 MB/s. As with microSD, microSD Express comes in several formats and storage capacities. MicroSD Express cards are available in the standard microSD form factor and are backward compatible with devices that support previous generations of microSD cards. Their capacities range from 128 GB to 1 TB.

In terms of features, microSD Express cards offer faster read and write speeds than traditional microSD cards, as they use the NVMe protocol. They also support the A2 (Application Performance Class 2) and Video Speed Class 90 standards, ensuring higher performance levels for apps and video recording. They're also built with advanced error correction and wear-leveling features to ensure data integrity and prolong the card's life.

Furthermore, some microSD Express cards have built-in encryption features, such as AES 256-bit, allowing users to protect their data with a password. This feature can be useful for storing sensitive information on the microSD card.

The differences between microSD and microSD Express can be boiled down to several key features, with the main being increased speeds and storage capacities. PCIe and NVMe protocols make the latter much faster with higher data throughput.

Overall, microSD Express cards are designed for devices requiring high-speed data transfer, e.g., smartphones, action cameras, drones, and other portable devices. They're also useful for storing large amounts of data, such as photos, videos, and music, and for running apps and games that require high performance.

As mentioned earlier, microSD Express cards come in various storage capacities, ranging from 128 GB to 1 TB. The storage capacity of a microSD Express card determines how much data can be stored on it. For example, a 128-GB card is able to store approximately 26,000 photos, eight hours of 4K video, or 130,000 songs, while a 1-TB card can store approximately 200,000 photos, 64 hours of 4K video, or 1 million songs.

Regarding power requirements, microSD Express cards are designed to be low power and don't require additional power for their operation. They're targeted at applications that typically have built-in power sources (smartphones, tablets, cameras).

Some devices that use microSD Express cards may have specific power requirements. For example, a device that supports 4K video recording may require a higher power input to operate at the maximum recording quality. In general, users should check the specifications of their device to ensure that it's compatible with the microSD Express card they're considering and meets the device's power requirements.

Conclusion

The differences between microSD and microSD Express can be boiled down to several key features, with the main being increased speeds and storage capacities. PCIe and NVMe protocols make the latter much faster with higher data throughput. The same can be said for storage capacities, as microSD evolved over time to gain 1-TB capacities, while that same storage parameter hit microSD Express much faster. So, what does the future hold for both technologies?

Continued improvements are expected for both, especially when it comes to their storage capacities. We will likely see these cards with capacities reaching 2 TB or more. This will enable users to store even larger amounts of data, such as full-length movies, entire music libraries, and extensive collections of photos and videos.

Another anticipated advance is the speed of both card types. As technology evolves, these cards may have even faster read and write speeds. As a result, devices will be able to transfer data more quickly, improving the performance of applications and allowing for smoother playback of high-definition video and audio. There's also a chance to see the integration of newer technologies, such as 5G and Wi-Fi 6, in microSD and microSD Express cards, bringing greater data-transfer speeds and enhanced connectivity capabilities.

On top of that, as the demand for data storage ramps up, the use of these cards may expand to new types of devices and applications, further increasing their utility and popularity. Overall, the future of microSD and microSD Express looks promising, with continued advances in storage capacities, speed, and other features that will make them even more useful and versatile. ed

Ad Index

ABC TAIWAN ELECTRONICS CORPORATION 25
ABSOPULSE ELECTRONICS LTD 4
ALTECH CORP 3, 5
AVTECH ELECTROSYSTEMS LTD 13
BROADCOM 29
CERTIFIGROUP, INC. 4
COILCRAFT, INC .... BC
DEWESOFT 31
DIGI-KEY ELECTRONICS .... FC, IFC
EBM-PAPST 1
EDUCATED DESIGN & DEVELOPMENT, INC. 18
HARWIN PLC 27
IMS INTERNATIONAL MICROWAVE SYMPOSIUM 37
KIKUSUI AMERICA, INC. 11
KOA SPEER ELECTRONICS, INC. 33
PICO ELECTRONICS, INC. 6
ROCHESTER ELECTRONICS, INC. 7
TAG CONNECT .... IBC
TDK-LAMBDA AMERICAS, INC. 19, 26
YUAN DEAN SCIENTIFIC CO., LTD. 28

BILL WONG | Senior Content Director

On the Lab Bench: **Dev Kits**

Robots on the move: Check out what's on Bill's bench these days from the likes of NXP and NVIDIA.

I try to limit the number of development kits that roll by my workbench because I want to test them out and then write about them, or highlight the kits in our online Kit Close-Up video series.

One common theme on the lab bench with these kits is artificial intelligence and machine learning (AI/ML). All have support for, and target applications that can utilize, AI/ML acceleration. They also can run Linux.

Right now, the stack is pretty high and I won't get some of these up on our website for at least a couple weeks, but please check them out when they're posted. In the meantime, here are a few that I think you might find useful, especially for robotic applications.

The NavQ+ is intended for the latest HoverGames competition *(Fig. 1)*. Based on NXP's i.MX 8M Plus, it can be paired with the new HoverGames buggy or the quadcopter. The focus of the latest challenge is to help Sustainable Food Ecosystems.

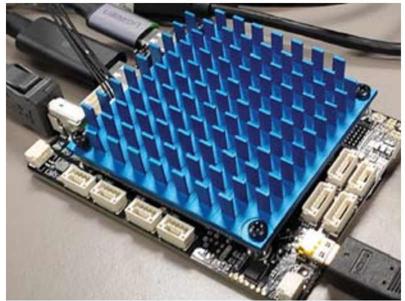

1. The HoverGames 2 platform included the NavQ+ based on an NXP i.MX 8M Plus. *NXP*

NavQ+ is designed to be the host process for the robots normally controlled by the RDDRONE-FMUK66 flight management unit (FMU), which can run Dronecode and PX4 Autopilot. The i.MX 8M Plus's multicore Arm Cortex-A53 processors are linked to a neural processing unit (NPU). The NPU, which delivers up to 2.3 TOPS, is likely dealing with data from the dual image-signal-processing (ISP) units to handle video input. There's a Cortex-M7 that can manage real-time chores with time-sensitive-networking (TSN) Ethernet support. The FMU and NavQ+ support single-pair Ethernet (SPE).

The AMD/Xilinx KR260 Robotics Starter Kit *(Fig. 2)* is built around the Kria K26 system-on-module (SOM), which contains a Zynq UltraScale+ MPSoC with an integrated FPGA. The processing sections includes multicore Arm Cortex-A53s and a dual-core Cortex-R5. The SOM can be configured to handle TSN and the carrier board exposes four Ethernet connections along with an SFP connection. There's a hardware root of trust (RoT) with secure boot and TPM2.0 support.

2. Xilinx's KR260 Robotics Starter Kit can be customized to include support for time-sensitive networking (TSN). *Xilinx*

This platform targets AI/ML applications with its FPGA support. Prototyping is supported with four Pmod connectors and a Raspberry Pi HAT header. There are four USB 3.0/2.0 interfaces as well.

NVIDIA's Jetson Orin Nano Developer Kit *(Fig. 3)* fills out the high end of the low end. The compact module consumes less space and power than the Jetson AGX Orin. The Nano supports the Robot Operating System (ROS) 2, which also could be used in other platforms mentioned here. The GPU is built around 1,024 Ampere cores and 32 Tensor cores that deliver 40 TOPS while using 15 W of power. Six 64-bit Arm Cortex-A78AE cores also are in the mix.

3. NVIDIA's Jetson Orin Nano delivers up to 275 TOPS. *NVIDIA*

The Jetson Orin Nano module is plug-compatible with other modules in the family, making it usable with a range of third-party carrier boards.

BeagleBone's BeaglePlay is more than a plaything *(Fig. 4)*. At its heart is Texas Instruments' (TI) AM6254 that contains Arm Cortex-A53, Cortex-R5, Cortex-M4, and TI PRU cores. It supports Ethernet and SPE along with a host of interfaces including GROVE, QWIIC, CSI2, USB Type-A and Type-C, plus a mikroBUS header. It can even drive an HDMI display. On the wireless side, it supports sub-1-GHz IEEE 802.15.4 that can connect to a BeagleConnect Freedom up to 1 km away.

4. The BeaglePlay from BeagleBone is built around Texas Instruments' (TI) AM6254 featuring Arm Cortex-A53, Cortex-R5, Cortex-M4, and TI PRU cores. *BeagleBone*

Again, for more insights on all sorts of kits, check out the Kit Close-Up videos at *www.electronicdesign.com* when you get a chance. ed

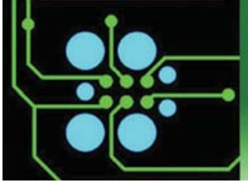

Tag-Connect

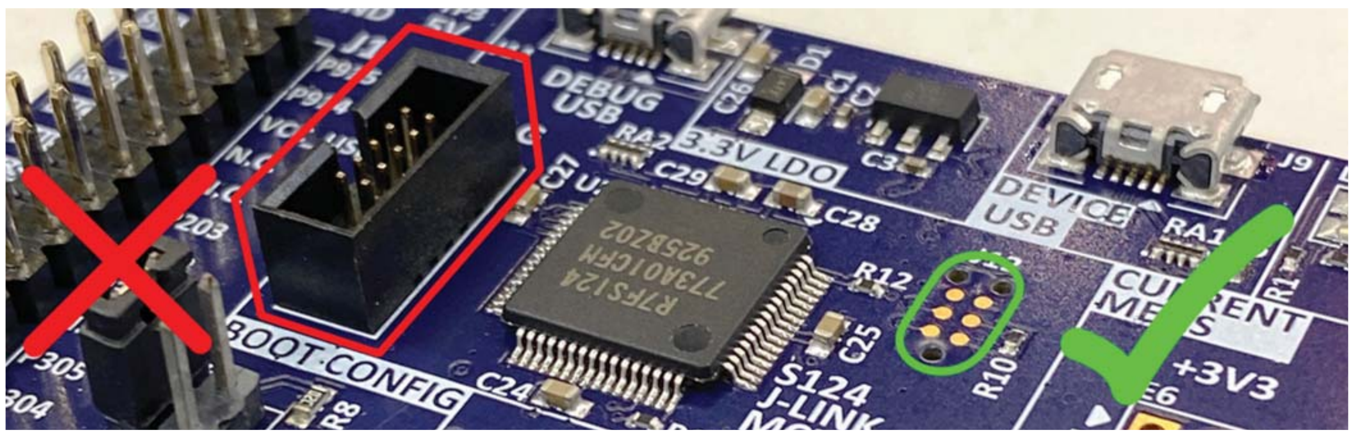

No Header! *No Brainer!*

JTAG Connector Plugs Directly into PCB!!

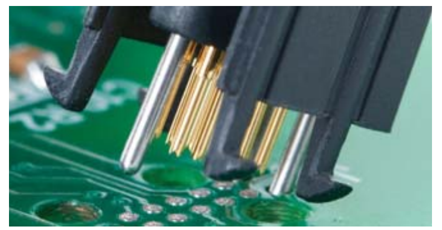

Legged version snaps into the board

Efficient board to board programming

Edge-Connect™ virtually eliminates board space!

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.

Eliminate JTAG connectors!

Tag-Connector footprints as small as an 0805 resistor.

Spring-pins tested to over 150,000 cycles!

Space and cost saving solutions for most popular programmers, debuggers, devices and protocols.

Tag-Connect - Saving the World 10,000 Connectors at a time!

www.Tag-Connect.com

STLink-V3 for illustration only

Find the Right Part Faster

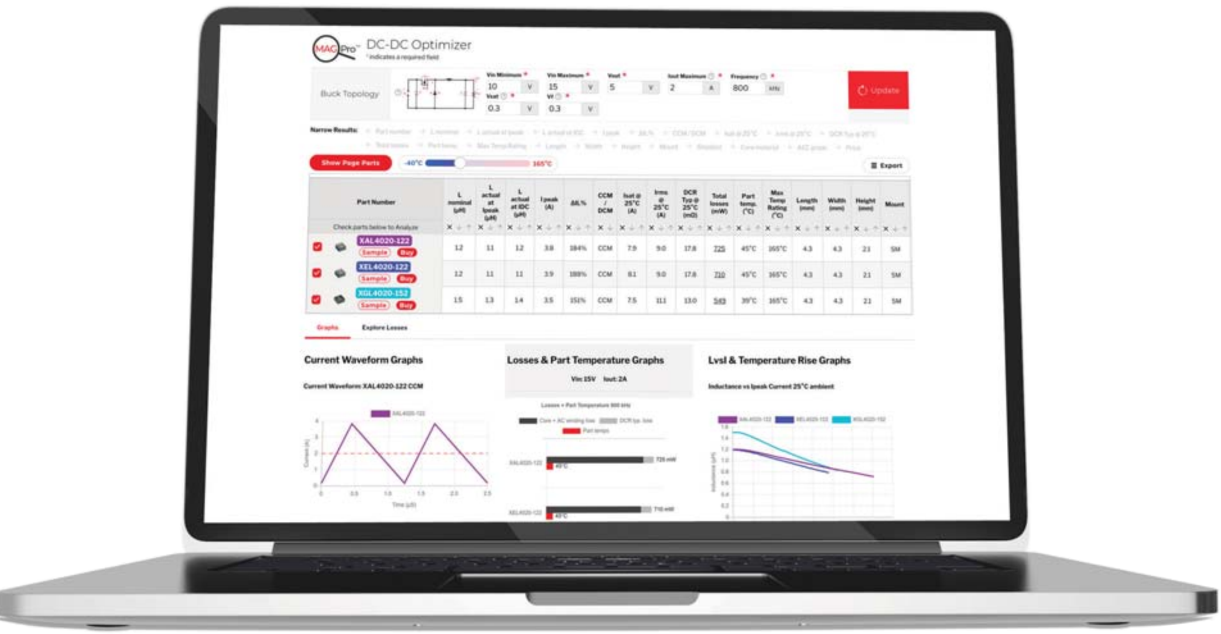

Designed by Engineers for Engineers, our MAGPro™ DC-DC Optimizer helps you find the optimal power inductors for your converter designs quickly and easily, reducing your design cycle time.

Coilcraft's MAGPro suite of online inductor analysis tools are designed to enable inductor selection and circuit optimization based on sound engineering principles and measured data.

The DC-DC Optimizer starts with your power converter parameters, calculates the needed inductor specifications, identifies off-the-shelf part numbers, and provides side-by-side performance analysis.

The tool identifies optimal inductors for buck, boost, and buck-boost converters. With just a few clicks you can go from V_{IN}/V_{OUT} converter requirements to inductor selection complete with losses and saturation analysis, all based on verified inductor data.

Reduce your design cycle time with confidence at **www.coilcraft.com/tools.**

WWW.COILCRAFT.COM