From Detection to Decision: An Inside Look at Next-Generation ADAS

Texas Instruments

Taylor Gage
Systems Engineer
Automotive ADAS

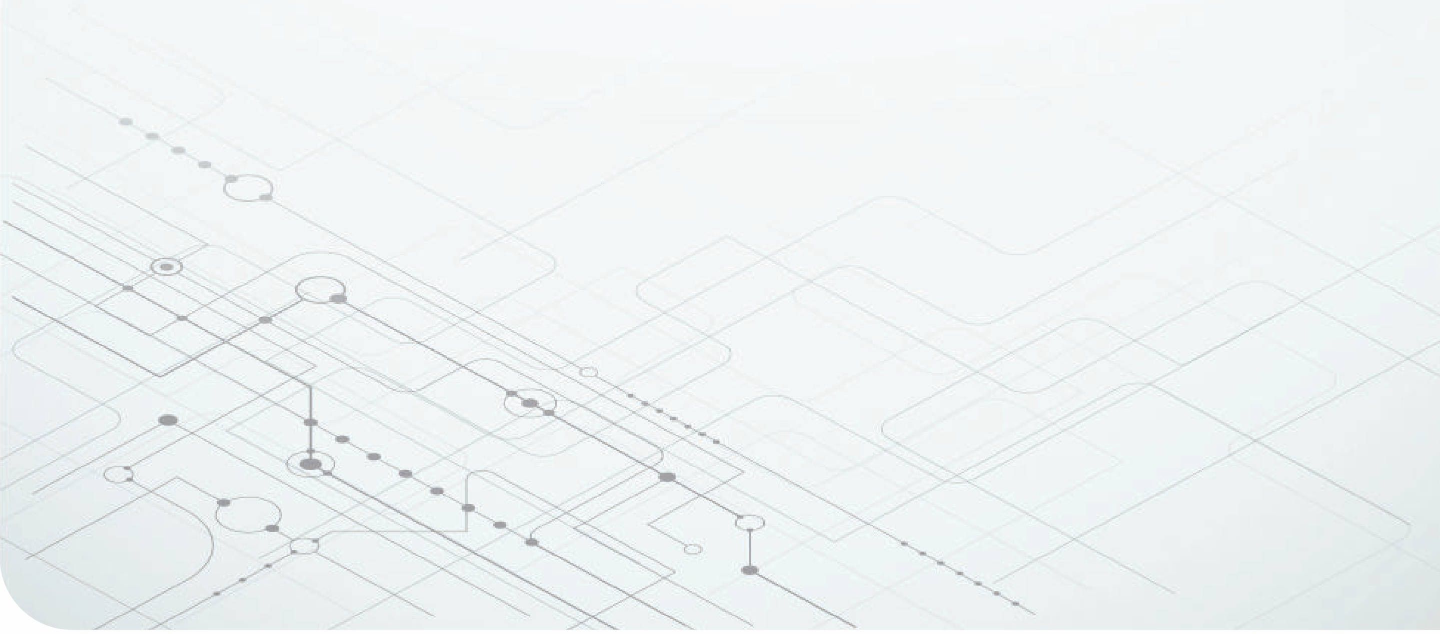

At a glance

1

Sensor fusion

Radar, lidar, and camera integration deliver the redundancy and cross-validation needed for safe, real-world ADAS operation

2

End-to-end ADAS

Autonomous vehicle success depends on the synergy of advanced driver assistance systems (ADAS) perception, computation, communication, and actuation

3

Enabling infrastructure

Advancements in semiconductors determine whether systems perform reliably and scale toward autonomy levels beyond Level 3

Introduction

Vehicle autonomy has advanced faster than most have predicted, but the engineering challenges required to deliver the underlying technology improvements have grown with it. Adding more sensors or increasing computing power alone won't get vehicles to the next SAE level. A vehicle must have continuous and accurate sensing data about the surrounding environment, enabling it to process complex situations and take action in milliseconds despite the existence of adverse environmental conditions such as fog, heavy rain, glare, or road debris.

Emerging government and industry regulations with stricter safety certification, reliability, and cost constraints continue to shape design decisions. The real challenge for automakers is to architect a complete end-to-end ADAS platform that safely and reliably translates detections into decisions in milliseconds. The success of these smarter, self-aware vehicles depends not on how advanced any individual system may be, but on how seamlessly all of these systems work together.

The forces at play

Driver assistance features once reserved for luxury vehicles now come standard for today's drivers. Consumer expectations have pushed the market firmly into Society of Automotive Engineers (SAE) J3016 Level 2 autonomy, requiring lane-keeping assistance and adaptive cruise control as standard features, with automated parking and highway piloting making their way into mid-range and entry-level vehicles. For most automakers, Level 2 autonomy and beyond have become the starting point.

Government regulations are also driving the development and adoption of advanced technologies to improve vehicle safety. Requirements now include driver monitoring to keep attention on the road, automatic emergency braking for collision avoidance, and blind-spot detection and occupancy monitoring to prevent children from being left behind in an unattended vehicle. Specific requirements vary by region, but governing bodies worldwide – from the National Highway Traffic Safety Administration in the U.S. to the European Commission in the European Union and the Ministry of Industry and Information Technology in China – share the same goal: keeping drivers, passengers, and vulnerable road users safe.

Addressing these challenges and delivering safe and reliable autonomous vehicles worldwide requires automakers to drive innovation across four areas:

- Perception, where systems capture and interpret the environment.
- Computation, where processed data drives real-time driving decisions.
- Communication, where large volumes of raw and processed data transfer across the vehicle.
- Response, where intelligent chassis systems such as control-by-wire translate decisions into action in milliseconds.

Perception: Understanding the environment

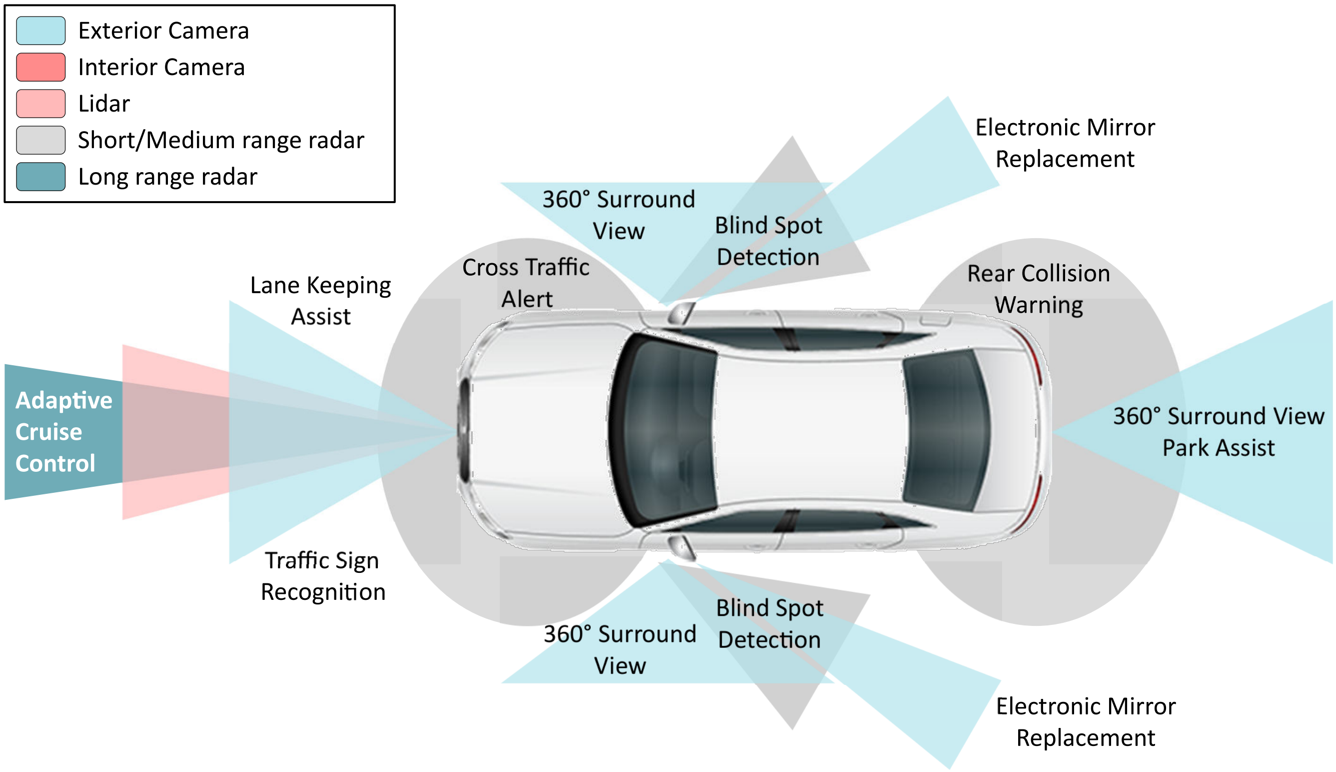

Figure 1. *Vehicle coverage by sensor type*

Fully understanding the vehicle's environment at all times is one of the biggest engineering challenges in ADAS design. What may seem straightforward can quickly become a major bottleneck. A picture of how each sensor typically provides coverage on a vehicle is shown in **Figure 1**. Each type of sensor has its own strengths and weaknesses, and no single sensor type alone can provide a comprehensive view of a vehicle's surroundings.

Sensor considerations for ADAS design

Radar excels at range and velocity sensing through frequency modulated continuous wave (FMCW) measurement, delivering this information independent of environmental conditions despite fog, rain, or low light. High-resolution millimeter wave (mmWave) radar sensors continue to improve accuracy at ranges >300m, enabling vehicles to detect objects sooner.

Where radar has traditionally fallen short is in object classification in lower-resolution and high-dynamic-range scenarios. However, **advanced 4D imaging radar** is closing the gap, enabling more accurate object classification and height detection previously unachievable with standard mmWave radar. Even with these advancements, radar alone cannot match the color, text, and object classification richness that cameras provide.

Lidar provides complementary advantages, particularly in long-range resolution. Traditional lidar systems use time-of-flight measurement to construct detailed 3D maps of a vehicle's surroundings through near-infrared light, capturing object distance, shape and relative location – important data when navigating complex situations such as sensing and driving around obstacles on the road or making tight maneuvers around sensitive terrain.

In long-range scenarios, scanning lidar provides lower power consumption and eye safety benefits, although its reliance on motor-driven mirrors introduces mechanical stress and reduces operational lifetime. For short-range scenarios, flash lidar has gained rapid adoption by eliminating moving parts, instead triggering all lasers in an array simultaneously to capture the full scene in one frame.

Flash lidar systems are smaller and easier to implement than their scanning counterparts but cannot support long-range applications given the laser power output levels necessary to comply with eye safety requirements. While inclement weather such as fog and snow remains a challenge for lidar, advancements in **laser-pulse lidar receiver technology** are improving performance in these conditions. Engineers are working to bring FMCW laser modulation to automotive lidar to increase reliability, reduce the needed aperture size and add object speed data.

Cameras complete the three sensing modalities, offering capabilities that radar and lidar cannot match: color, text and rich object classification data. Cameras accurately identify road signs, traffic lights and pedestrians. Features such as lane-keeping assistance rely entirely on camera-based systems because lane markings on the road are two-dimensional and undetectable using radar or lidar. Cameras do carry the greatest environmental vulnerability of the three modalities, however. Low-light conditions, glare from bright lights and debris such as mud on the lens can render camera systems effectively blind.

Sensor fusion. Each sensing modality has scenarios where it excels and environments where it falls short. Thus, automakers have pursued higher levels of autonomy through sensor fusion, combining two or more modalities rather than solely relying on one type. As an example, pairing radar for object detection and camera for object classification improves sensing accuracy through cross-sensor validation and consensus-based decision-making. Multiple sensor types also provide redundancy for enhanced safety. If a camera cannot see because of poor lighting conditions, radar or lidar systems can supplement the necessary field of view to keep the vehicle operating safely. Combining sensor data from the available sensor modalities combines the individual strengths of color perception, high resolution, object speed data, and robustness against harsh weather conditions.

Applying sensor fusion to accurately discern objects

Sensor fusion addresses challenges that go beyond poor weather and lighting. For example, municipalities across British Columbia implemented a "Pavement Patty" optical illusion, as shown in **Figure 2**.

Designed to surprise and slow down speeding drivers, the pavement art appears to show a child chasing a ball in the middle of the road. In reality, an 8-foot vinyl sticker applied at an angle creates a 3D illusion to the human eye. The problem is that this illusion fools camera systems too. An autonomous vehicle relying purely on image data would determine that a child is running in front of the vehicle and take drastic action, such as sudden braking or entering into oncoming lanes of traffic. Because object classification is a strength of image sensors, but object detection is not, this confusion can be avoided by adding a radar or lidar sensor to the vehicle. Now when the camera signals that it sees an object or person in the road, this information can be checked against the sensors whose strength is object detection. Once this happens, the radar or lidar system would correctly indicate a hazard-free road ahead, and the overall ADAS system can determine that the camera is hallucinating the hazard and avoid making dangerous maneuvers.

Illusions such as these are intended to encourage human drivers stay attentive and drive safely, but pose challenges for autonomous vehicles, which already follow posted speed limits and monitor the environment at all times by forcing them to address hazards that are not actually present. Pavement Patty is an interesting case that exposes the fact that every sensing modality has blind spots. Fusing radar, lidar, and camera sensor data provides the redundancy and cross-validation needed to achieve higher levels of autonomy safely and avoid ambiguous situations such as these.

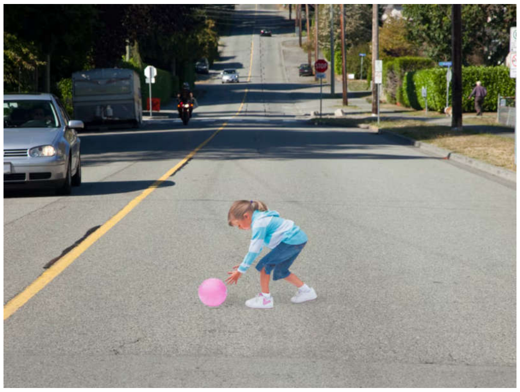

Figure 2. *The Pavement Patty 3D optical illusion in Vancouver, Canada*

Computation: Turning perception data into driving decisions

Perceiving environmental data is only the first step to autonomous vehicle operation. Vehicles must convert that data into comprehensive images of their surroundings and make real-time decisions within milliseconds. In many of today's vehicle architectures, data processing for a specific application happens at the sensor itself. For example, long-range front radar could be the only input to adaptive cruise control. As self-driving advances and sensor fusion increases, processing is shifting from the edge to a centralized architecture.

Centralized computing units deliver several benefits: increased processing efficiency, simplified software update rollouts, and easier scalability of features, since changing sensors at the edge no longer requires a complete electronic control unit or software overhauls.

Sensor fusion applications and automated decision making require specialized computation acceleration such as neural processing units, vision digital signal processors (DSPs), and radar fast Fourier transform accelerators, all of which help create machine learning (ML) and predictive artificial intelligence (AI) networks driving complex, real-time decisions on the road. Texas Instruments' **TDA5** driver assistance system-on-a-chip integrates multiple computing domains to address sensor processing, real-time control and ML all in one package. Reference designs and software also help engineers shorten their development timelines and reach production faster.

Communication: The backbone of the vehicle

As important as reliable sensing and high-performance computation are, without a network of high-bandwidth, low-latency and long-distance communication systems, these functions remain isolated. As shown in **Figure 3**, a range of communication protocols, each tailored to specific data input types and support mission requirements, connect systems across the vehicle.

High-speed, data-dense applications such as cameras require cable lengths of >10m to reach the computing unit. **Serializers and deserializers (SerDes)** provide an asymmetric pipeline for data, with forward channel speeds reaching >7Gbps and employing low-speed, back-channel protocols such as serial peripheral interface (SPI) or inter-integrated circuit (I^2C) for bidirectional communications.

While SerDes technology has traditionally only been used in camera applications, the increasing evolution of radar and lidar system technologies has led to streaming architectures that send raw data directly to the central computing unit, which also requires the high-speed performance of SerDes. As these applications expand, time synchronicity becomes essential to prevent radar or lidar interference and to enable higher-resolution cascaded radar configurations that rely on very tight time coupling between sensors. For a specific example, see the **Single-Chip mmWave Streaming Radar Reference Design for Automotive Front Radar**.

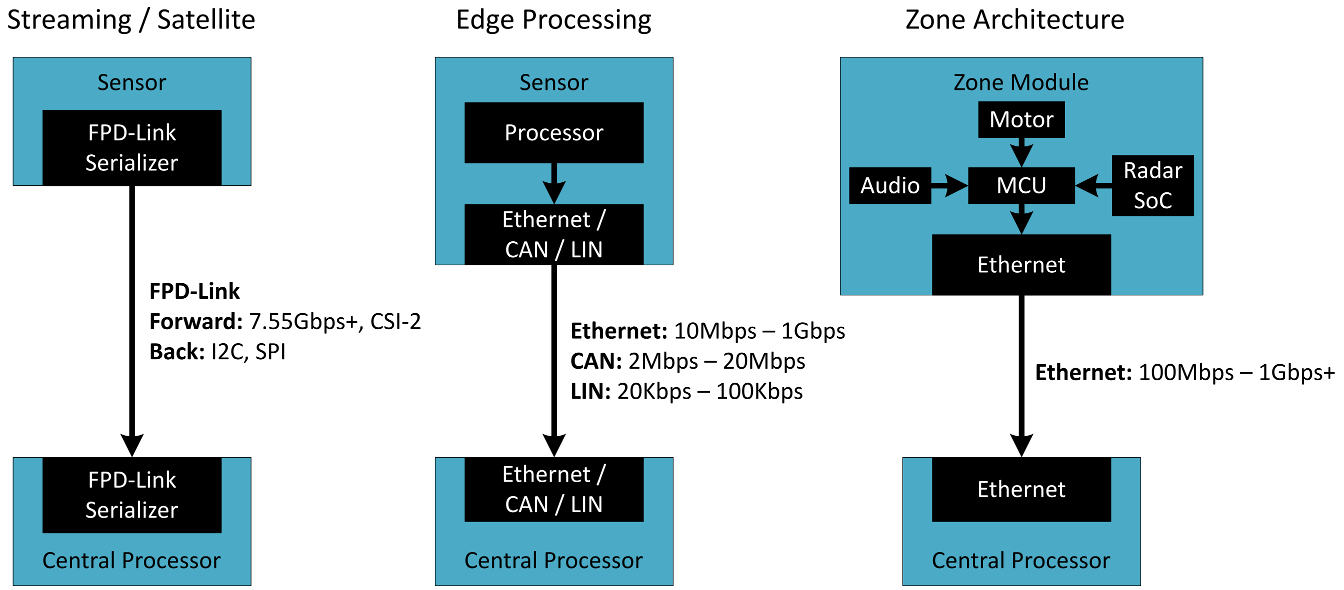

Figure 3. *Diagram of various networking protocols and bandwidths to support different vehicle architectures*

In addition to SerDes, **Automotive Ethernet** has emerged as a reliable protocol for in-vehicle networking. Ethernet connects domain controllers and enables distributed processing and zone architectures. For ADAS and infotainment applications requiring separate compute units, Ethernet provides a long-distance bidirectional signal chain with integrated security features such as media access control security (MACsec) to meet growing cybersecurity requirements.

Ethernet also supports tight time synchronization, keeping data and decision-making consistent throughout the vehicle. In zone architectures, Ethernet aggregates multiple data sources such as corner radar, seat motor controls and audio signals, along with legacy protocols such as controller area network (CAN) and local interconnect network (LIN), into a single output to the centralized compute unit, reducing cabling costs and complexity. Standards such as 10BASE-T1S extend Ethernet connectivity to lower-bandwidth end nodes within the zone, further simplifying wiring and reducing cost. This zone networking architecture directly enables software-defined vehicles where over-the-air (OTA) updates push new features and capabilities to the vehicle post-production, reducing the need for costly hardware revisions. To learn more about this topic, see the article, "**How Ethernet Accelerates the Move to Software-Defined Vehicles**."

Although zone architecture and streaming ADAS sensors drive the need for Ethernet and SerDes, CAN and LIN remain essential for control systems throughout the vehicle, including seat motors, post-processed radar data, and lighting. For high-performance applications such as ADAS sensing or powertrains, CAN Flexible Data-Rate (CAN FD) offers a multicommander topology for bidirectional commands. And for simpler applications such as heating, ventilation and air conditioning, and seat motors, CAN FD Light provides a fixed bit rate and a single commander topology, reducing costs and complexity at the end nodes and zone controllers.

As vehicle autonomy increases and architectures evolve, designers need flexible networking options that support precise time synchronization, functional safety and cybersecurity protections.

Response: Turning data into actions

Sensing and computing are important, but mean nothing without action. More comprehensive data, advanced computing, and faster communication are converging at the final step in autonomous vehicle operation: making the vehicle act. Intelligent chassis systems add the critical actuation layer that transforms driver assistance decisions into precise physical responses, with control-by-wire technology delivering superior control precision and performance within millisecond timeframes.

As these systems grow more sophisticated, predictive capabilities will draw on proven real-world data to reduce reaction times to instantly avoid hazards, protecting passengers and pedestrians before a driver realizes the threat. Figure 4 shows an example of such a process.

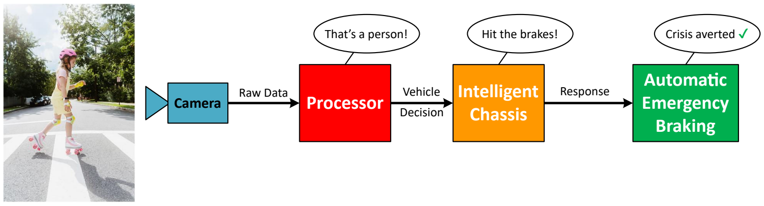

Figure 4. *Detection to decision flowchart*

The hidden enablers: IC technologies that make it all possible

Even the most advanced ADAS platform depends on analog semiconductors and embedded processors supporting every system across the vehicle. Bulk acoustic wave and power-management integrated circuits (ICs), often overlooked in broader autonomy discussions, enhance system performance across perception, computation and networking, ensuring that these systems operate reliably under real-world conditions. TI designs these devices with the full vehicle architecture in mind, not just the individual application.

Conclusion: The road to higher autonomy

The next decade will bring significant advances across every layer of autonomous vehicle technology. FMCW lidar will deliver more accurate velocity and direction data from already precise ranging sensors. More sophisticated AI models will address complex driving environments and the operational edge cases that today's systems cannot handle. And as the bottleneck shifts from hardware to software, OTA updates will offer automakers a path to SAE J3016 Level 4 autonomy built on Level 3 hardware. Intelligent chassis systems will use real-time information to turn data into action and enable a smoother and more responsive ride.

These technologies have moved well beyond the concept phase. Adoption continues to grow, driving costs down and introducing higher levels of safety to previously cost-prohibitive designs. Full autonomy depends on seamless integration across every system in the vehicle. By driving innovation across sensing, processing, connectivity, and actuation, TI is helping automakers create safer, more intelligent vehicles for every driver, on every road, around the world.

About the author

Taylor Gage is a systems engineer on TI's automotive ADAS engineering and marketing team, specializing in external sensing, including camera, radar and lidar system design. He leverages expertise in FPD-Link networking and TI's radar portfolio to help OEMs and Tier-1 suppliers advance autonomous-driving technology and develop safer vehicles. Taylor earned a B.S. in electrical engineering from Texas A&M and is pursuing an MBA at Southern Methodist University.

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, regulatory or other requirements.

These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you fully indemnify TI and its representatives against any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI’s products are provided subject to TI’s Terms of Sale, TI’s General Quality Guidelines, or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI’s provision of these resources does not expand or otherwise alter TI’s applicable warranties or warranty disclaimers for TI products. Unless TI explicitly designates a product as custom or customer-specified, TI products are standard, catalog, general purpose devices.

TI objects to and rejects any additional or different terms you may propose.

Copyright © 2026, Texas Instruments Incorporated

Last updated 10/2025