

Developing Smarter, Safer Cars with ADAS IP

Optimizing performance, safety, and reliability in ADAS applications via design, verification, and processor IP and a holistic design approach

By Charles Qi, Sr. Design Engineering Architect, and Neil Robinson, Product Marketing Director, IP Group, Cadence

Today’s cars are a full-fledged electronic system on wheels, where every part is interrelated and must be designed, optimized, and verified simultaneously. As a result, it’s important to apply a holistic approach when developing automotive systems, taking into consideration the chip, the package, the board, the subsystem, and, ultimately, the whole vehicle. In this paper, we’ll examine the role of Cadence® advanced driver-assistance system (ADAS) intellectual property (IP) and the Cadence System Development Suite’s holistic design approach in creating dependable, high-quality automotive systems.

Contents

- Introduction..... 1
- Delivering a Smarter, Safer Driving Experience with ADAS2
- Subsystems Within the ADAS Environment3
- The Role of CNNs in ADAS Applications3
- ISO 26262 and Functional Safety..4
- Processing All of the Data4
- Why Only a Holistic Design Approach Can Meet Standards for Automotive Quality and Safety5
- Cadence’s System Design Enablement Strategy6
- How IP Helps Optimize ADAS Design6
- Designing and Verifying Automotive Systems..... 7
- Summary8
- Further Information.....8
- Sources.....8

Introduction

As the automotive industry continues to develop more autonomous driving features and self-driving cars, automotive electronics will maintain an integral role in enabling these capabilities. Future vehicles will boast more integrated infotainment functions, sensor clusters, computer power, car-to-object (Car2X) communication technology, high-bandwidth Ethernet networks, and high-definition (HD) displays. While electronic components are enhancing the drive, government regulations are increasingly calling for automotive manufacturers to integrate redundant sensing and control systems with more cameras, radar, and other ADAS features into vehicles for increased safety and reliability. Figure 1 depicts how technology is transforming transportation.

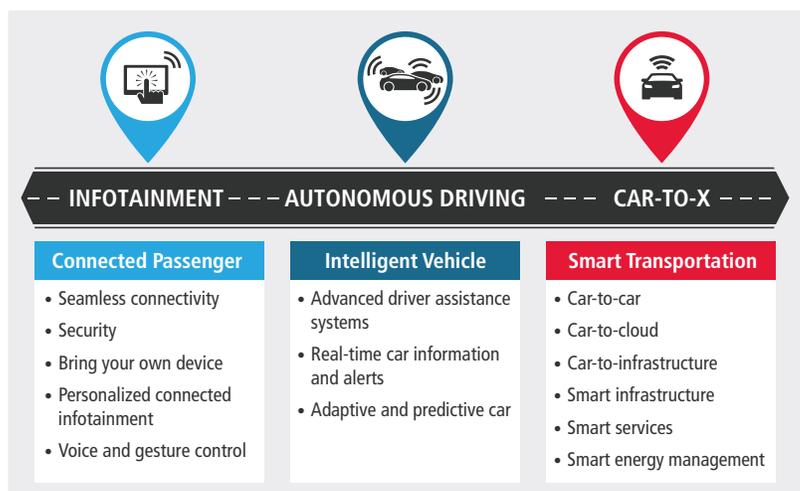


Figure 1: Technology is making vehicles increasingly smarter

Ensuring that these electronically sophisticated vehicles operate as intended in an array of road and environmental conditions requires an array of design and verification tools, new algorithms and runtime software, and services aimed at optimizing the whole automotive system. Each vehicle contains multiple in-vehicle networks, from the infotainment network to the safety, guidance, engine management, and several other networks. As requirements for each subsystem within the vehicle are defined, it's important to consider how they might impact each other as well as overall vehicle operations. That's why all of the electronic circuits need to be verified to ensure correct hardware and software functionality, expected component tolerances, temperature variations, stress-induced failure mechanisms like electromigration and electrostatic discharge, and a host of other parameters to safeguard against failures in the field. Additionally, automotive electronic subsystems must be simulated with each other, to ensure that each interconnected subsystem will continue to operate as intended over a long period of wear and tear. Indeed, exhaustive simulations, fault analysis, and yield and reliability analyses are critical to avoiding catastrophic safety problems and potential costly recalls.

IP plays an important role in all of this, supporting essential communications protocols like Ethernet, as well as functions including real-time data/audio/voice processing, sensor fusion, pattern and voice recognition, sound enhancement, and connectivity. The ADAS segment is one of the fastest growing of the automotive semiconductor space—essential for enhancing the driver experience and overall safety. We'll center our discussion in this paper on the role of IP in ADAS.

Delivering a Smarter, Safer Driving Experience with ADAS

ADAS technology (Figure 2) enables vehicles to become more aware of their surroundings and, generally, safer to drive. Adaptive cruise control, driver monitoring systems, automatic parking, collision avoidance, lane departure warning systems, and traffic sign recognition are just a handful of the many functions covered by ADAS. Data is at the heart of many of these functions—various in-vehicle sensors collect huge amounts of data in real time. Various subsystems, in turn, must process this data accurately in real time, using it to make informed emergency braking, steering, and communications decisions that impact driver safety and the overall driving experience. Over time, vehicles will increasingly be able to exchange data with other vehicles and their environment for an even safer driving experience. Market research resource Research and Markets projects that the ADAS market will reach a value of almost \$6B in 2016, growing to about \$10B by 2022. From a unit standpoint, roughly 50 million vehicles equipped with ADAS are expected to ship in 2016, and this number is anticipated to grow to over 60 million units by 2022¹.

ADAS applications are unique in that purpose-built, low-power, high-performance system-on-chip (SoC) devices and software must react and interact reliably, and be ready for continuous enhancements to align with market requirements and safety upgrades. While key technologies for vision, radar, ultrasound, real-time networking, and embedded control can be adapted from other applications, the special requirements of ADAS limit the off-the-shelf chip choices for designers. Because of this, it's a challenge for OEMs (the automakers) and Tier 1 suppliers (such as automotive application designers) as well as a big opportunity for vendors who provide the tools, software, and services to create these SoCs.

Making ADAS capabilities possible requires some unique design specifications:

- Higher compute performance, at >1000GMAC/s, with support from a digital signal processing (DSP) architecture tuned to process compute-intensive algorithms while delivering an optimal SoC power, performance, and area (PPA) ratio
- Higher network/memory bandwidth, up to 1Gb, to support increased video/image resolution, increased frame rates, increased number of video streams and cameras, and the need to store and access intermediate results generated by highly complex algorithms
- Greater integration

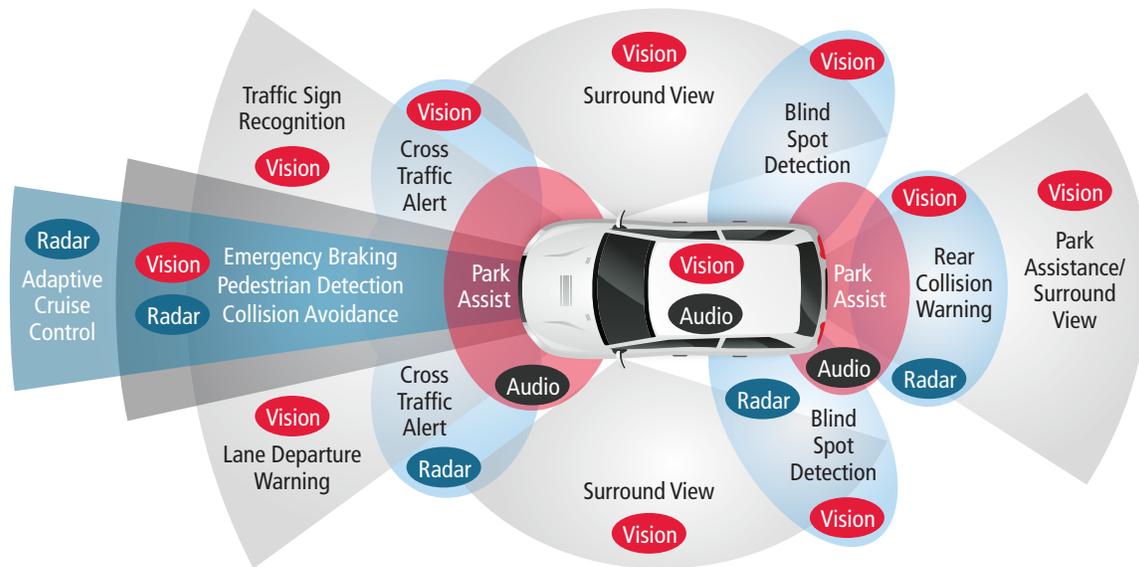


Figure 2: ADAS provides many of the capabilities that help make vehicles safer to drive

Subsystems Within the ADAS Environment

Creating these sophisticated ADAS functions requires a mix of sensors (monocular and stereo cameras, LiDAR, radar, ultrasound, etc.) and communications (both in-vehicle networks and Car2X wireless). These subsystems often use specialized IP building blocks to execute key algorithms for sensor processing, including computer vision, voice recognition, radar analysis, and reliable communications. Essentially, ADAS connects a variety of subsystems within a vehicle.

Imaging and Video Assistance

Imaging and video assistance encompasses everything including emergency braking, lane and vehicle tracking, traffic sign recognition, parking assistance, driver alertness monitoring, and low-distraction human-machine interfaces (HMIs). These applications must accurately process large volumes of data in real time and use this information to make decisions quickly and reliably.

Voice Control and Gesture Recognition

Advanced HMI support, voice control, and gesture recognition can enhance safety by providing the driver with hands-free control of a variety of functions, from the car radio to the air conditioning unit to navigation functions.

Vehicular Communications Systems

Car2X communication systems will make it possible for vehicles and roadside units to exchange information, such as traffic and safety-relevant data. In addition, vehicle-to-vehicle (V2V) communication has an advantage compared to vision, LiDAR, and radar systems in further enhancing safety by “looking around the corner.” Of course, all of this data can be shared with other cars via the cloud.

Automotive Ethernet

Delivering high-speed in-vehicle communication, automotive Ethernet is the key enabler for ADAS applications. It’s the data highway that allows, for example, video streams from side- and rear-view cameras to be processed and transferred to the dashboard display with the high bandwidth and low latency required. In addition, it can also be used as the data backbone that directly connects all domain controllers in a car.

The Role of CNNs in ADAS Applications

Underlying many of these functions are deep-learning technologies, such as convolutional neural networks (CNNs), that support applications like vehicle and pedestrian detection, road surface tracking, sign and signal recognition, and voice command interpretation. CNNs represent a special case of a neural network, a system of interconnected artificial neurons that exchange messages. The connections have numeric weights that are tuned during a training process; a properly trained network responds correctly when, for instance, presented with an image or pattern to recognize. The network has many layers of feature-detecting neurons and each layer itself has many neurons that respond to different combinations of inputs from the previous layers. A CNN has one or more convolutional layers, often with a subsampling layer, followed by one or more fully connected layers.

Compared with traditional pattern-detection methods, CNNs offer advantages in their ability to deal with distortions, their reduced memory requirement, and in their efficient training process. Tapping into the German Traffic Sign Recognition Benchmark (GTSRB), and using a proprietary hierarchical CNN methodology, Cadence has developed traffic sign recognition algorithms that have yielded a better correct detection rate (99.80%) versus a previously established baseline³. There's potential in the future for CNNs to be trained for more complex tasks, such as judgment and strategy.

ISO 26262 and Functional Safety

Unforeseen or unexpected errors in the electrical components of a car can lead to minor inconveniences, brand-damaging product recalls, or, worse, injuries or fatalities. Meeting functional safety standards helps ensure that an overall system will remain dependable and function as intended even when there is an unplanned or unexpected occurrence. To be functionally safe, a system must have redundancy to limit the risk that any single error will upset the entire system, as well as checkers to monitor the systems and trigger error-response and recovery features.

Complying with functional safety standards requires:

- Implementing requirements tracing from the system level down to the components
- Performing functional verification and safety verification at all levels of abstraction for all parts of the system
- Performing error (fault) injection in the functional verification environment

In the automotive world, ISO 26262 is the standard for functional safety of electronic systems in vehicles. The standard outlines automotive safety lifecycle phases, covers functional safety across the whole development process, and provides requirements to ensure that a sufficient and acceptable level of safety is being achieved². Compliance with ISO 26262 calls for:

- Detection and correction of hardware errors
- Detection and resolving of systematic faults
- Ability to prevent software tasks from affecting each other
- Reduced use of variable-latency system components
- Fast processing
- Documentation, including guides to safety features, tool qualification support, and verification reports

One of the biggest challenges of complying with ISO 26262 involves collecting and analyzing the massive amount of data involved in achieving the accepted safety-integrity level. Traditionally, this has been a manual, time-consuming process.

Processing All of the Data

Automotive engineers must determine whether processing the data from all of these ADAS subsystems should happen in a distributed or centralized manner. In a distributed manner, data processing takes place close to each respective sensor or camera and requires high-speed interface IP and DSPs, with automotive Ethernet IP enabling in-vehicle communication. In these scenarios, sensor fusion DSPs can integrate the output of multiple sensors, reducing data traffic to the central head unit. A centralized data-processing arrangement requires an automotive head unit connected to each of the subsystems (via automotive Ethernet IP, for instance). Interface IP, DSPs, and memory subsystems are essential here to ensure low-latency response of the system. Every piece of each of these scenarios must be verified to ensure that it will work reliably in the vehicle. Figure 3 shows an example of a generic ADAS SoC architecture.

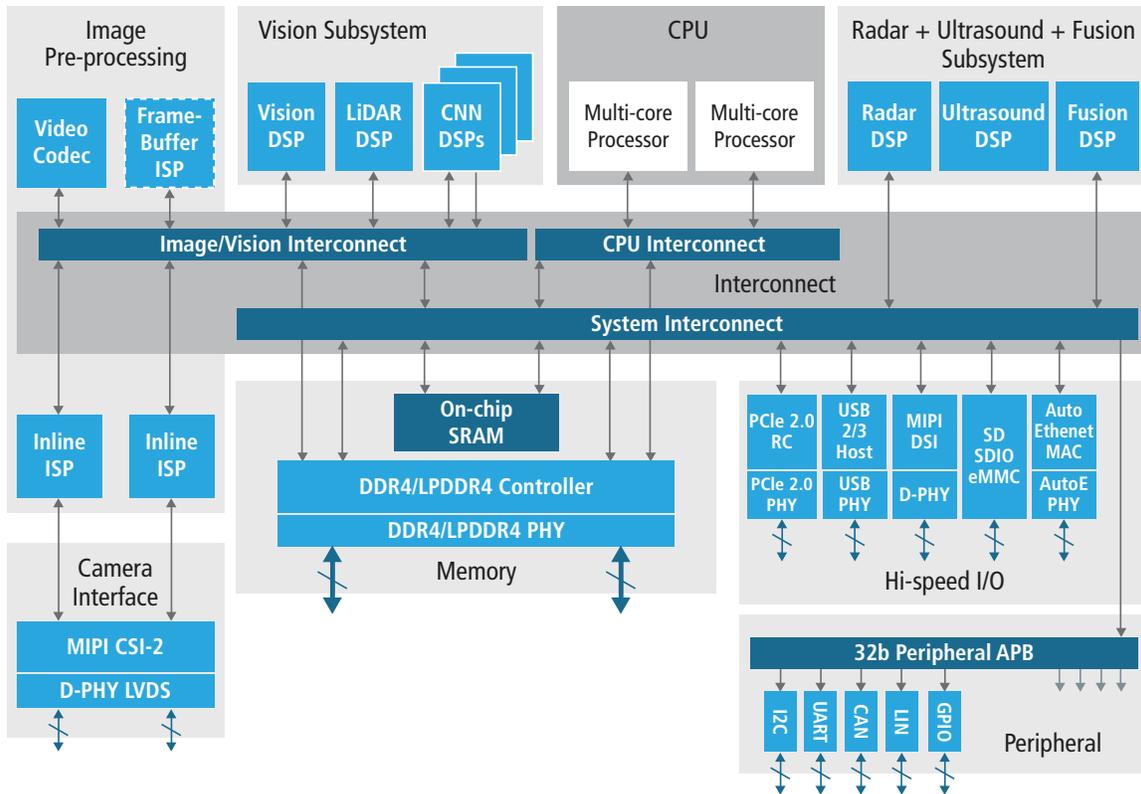


Figure 3: Generic ADAS SoC integrated architecture

Why Only a Holistic Design Approach Can Meet Standards for Automotive Quality and Safety

Clearly, a vehicle’s electronic components are all interrelated. At the same time, the non-electronic parts of the car, along with the vehicle’s operating environment, will impact performance and functionality. That said, a holistic design approach that accounts for all of these internal and external impacts is critical for ensuring that the vehicle meets automotive quality and safety standards.

One of the most valuable aspects of a holistic design approach is that it supports “test driving” of a chip and the entire system early on, well before bugs become too costly to fix and/or lead to time-consuming and expensive respins. It focuses on throughput/processing performance for compute-intensive algorithms (voice, vision, communications, etc.), with an eye for optimizing these algorithms early in the development cycle. The approach is also about thorough verification of hardware and software for safety and reliability.

Just as what was once done manually to develop a chip has been automated, so, too, has the process for designing and verifying a system. The way in which systems are designed is increasingly being driven by the end applications. As a result, semiconductor companies are delivering more than just the chip—they are also providing some layers of the software stack. For example, their responsibility includes OS ports, hardware/software verification in a system context, system-level analysis, packaging- and board-level co-development, and analysis of the chip, package, and board.

Cadence’s System Design Enablement Strategy

Companies who have built their foundation on electronic design automation (EDA) tools now have a rich opportunity to expand their scope into optimizing electronic system design. Their customer base is no longer only traditional chipmakers but also includes system designers, such as automotive OEMs and Tier 1 automotive application designers.

Cadence has built its foundation on a broad EDA portfolio of sophisticated SoC, packaging, and board design tools. Over the years, we have extended our expertise and portfolio into the system design space, delivering design, verification, and processor IP, hardware/software convergence platforms, software content, and services along with the foundational EDA tools. Taking a holistic design approach from an end-product viewpoint forms the basis of our system design enablement strategy. In the automotive space in particular, Cadence has been amassing a depth of knowledge and experience through years of working closely with customers to help them meet the challenges of designing and verifying automotive components, subsystems, and whole systems.

We work not just with automotive semiconductor suppliers, but also with Tier 1 vendors and OEMs to deepen our understanding of their requirements, obstacles, and opportunities. We continually enhance our products, ensuring that our solutions are relevant for each level of the supply chain. Our technologies and methodologies, along with offerings from our ecosystem partners, make Cadence ideally suited to enable automotive designers to overcome system integration, package and board, and chip design and verification challenges. Figure 4 depicts our strategy for optimizing automotive systems.

How IP Helps Optimize ADAS Design

ADAS designs continue to increase in complexity as automakers strive to meet functional safety requirements:

- Multiple sensors and sensor types are deployed in a single system to ensure reliable driving decisions
- High-performance sensing, DSP computation, and high-bandwidth memory and networks contribute to the vehicle’s ability to handle complex road and weather conditions
- Advanced process nodes, high-speed analog/mixed-signal circuits and I/Os, and multi-chip solutions with advanced packaging are all part of the mix

IP can help in each of these areas. For example, Cadence Tensilica® processor IP can support many of the core algorithms in ADAS. Tensilica Vision DSPs, based on a VLIW and SIMD architecture and operating at 70mW to less than 300mW of power, provide the scalable high performance and low power that ADAS computation demands. The latest processor in this family, the Vision P6 DSP, was designed for CNN applications. The Tensilica Fusion and ConnX DSP families are ideal for always-on, sensor-fusion applications, supporting the real-time data processing for sensors and cameras that enables automated parking and lane departure, as well as radar and communications.

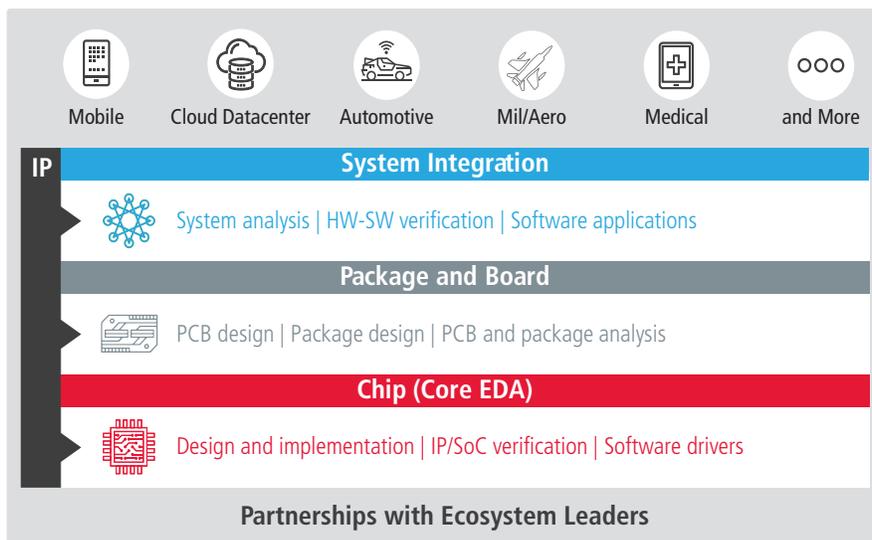


Figure 4: Cadence’s system design enablement strategy

Designing and Verifying Automotive Systems

Because software is available well before silicon, hardware/software co-verification is essential. The Cadence System Development Suite covers the entire design cycle, from early pre-silicon software development to silicon and system validation. The suite's connected platforms accelerate system design, IP and SoC verification, and bring-up, significantly reducing system integration time. Tier 1 suppliers can use these tools to prototype and test a complete system, evaluating a variety of real functions before hardware availability, for instance. OEMs can debug specific traffic situations to test and optimize their algorithms.

As an example, consider a camera system that sits behind the rearview mirror. Inside such a system is a CNN that runs on an ADAS SoC for object detection and tracking. Before silicon becomes available, the designer must develop, verify, and optimize the hardware platform running the complex CNN algorithms. The System Development Suite supports this flow, enabling the semiconductor vendor to verify the SoC to ensure that it is working as intended and to engage in early development of software drivers and firmware. To test the complete ADAS, a Tier 1 supplier can stream in video sequences of real traffic scenarios. At the OEM level, a software engineer can use the tools to validate and optimize a specific algorithm or debug a particular traffic scenario.

Addressing Reliability Issues

Automotive devices are expected to last at least 15 years⁴. To enable this degree of reliability, SoC designers must be able to account for transistor aging and interconnect electromigration, particularly for vehicle ambient temperatures and advanced SoC process nodes. Advanced technologies such as FinFETS provide power and performance advantages, but are prone to self-heating problems. Cadence's broad portfolio can facilitate total power signoff, from the transistor to cell to the entire automotive system.

Enabling ADAS Algorithms

Our scalable DSPs, along with off-the-shelf software from our partner ecosystem, support new algorithms for communication, audio, imaging, computer vision, and CNN functions that are integral to ADAS.

- Cars are being equipped with more cameras, LiDAR units, and radar sensors, collecting data from the environment around the vehicle. Our portfolio includes high-throughput DSPs that support heavy data communications for applications such as adaptive cruise control, emergency braking, sensor fusion, and V2V communications.
- Dedicated DSPs for audio, voice, and speech support always-on wake capabilities for voice triggering and automotive sound systems, including active noise control equipment
- Computer vision and imaging DSPs process data from a vehicle's many cameras, filling visual displays with meaningful information for the driver. The newest offering in this line was designed to provide the multiply-accumulate (MAC) performance that's critical in CNN applications, along with low-power consumption and on-the-fly data compression.

Complying with Automotive Interface Standards

The latest iteration of the ISO 26262 functional safety standard includes a chapter that defines the requirements for IP. From the product requirement phase to final IP release, Cadence follows formal quality flows and checkpoints to assure design quality. Our portfolio includes a wide range of system, interface, and memory IP that facilitates ADAS application design, including:

- Industry-leading IP for DDR4/LPDDR4 controller and PHY
- ISO 26262-ready IP for automotive Ethernet MAC controller
- IP for MIPI® camera/display controller/PHY
- IP for PCI Express® (PCIe®) controller and PHY

Additionally, our verification IP (VIP) can validate compliance with standard interface specifications such as CAN, LIN, Ethernet, DDR4, Flash, USB, and dozens more. In fact, every vendor supplying chips to the automotive industry uses our VIP. We continue to further develop our portfolio, so our offerings will evolve and expand over time.

Meeting Functional Safety Requirements

Cadence, which has been working on fault simulation technologies for more than 25 years, offers tools that automate the process of meeting functional safety requirements. Our functional verification environment reduces the ISO 26262 certification effort by up to 50% by automating fault injection and result analysis for IP, SoC, and

system designs. The solution meets the traceability, safety verification, and tool confidence level (TCL) requirements of ISO 26262. For instance, you can use the tools to inject a fault into a memory subsystem of an ADAS SoC to determine whether there are any problems in your software stack, or whether a safety mechanism like error code correction (ECC) could detect and address the fault.

Our digital design and implementation tools ensure functional safety and high-reliability design by defining safety islands that are adhered to for placement, routing, and multi-cut via insertion. In our Tensilica processors, features such as memory/data protection, fault management and task isolation, deterministic operation, and documentation support functional safety requirements.

Summary

In August 2016, Singapore unveiled the world's first self-driving taxi. Operated by an autonomous vehicle software startup, the small fleet includes six cars that select members of the public can hail for free via their smartphones. Rides are currently limited to a 2.5-square-mile business and residential district. The startup, nuTonomy, has a goal of having a fully self-driving fleet available in the country in 2018.

The day when cities operate fleets of self-driving vehicles, with the aim of decreasing the number of accidents, reducing traffic, and even mitigating parking hassles, is quickly coming to fruition. Because lives are at stake, ensuring the safety and reliability of automotive systems is critical. Doing so calls for nothing less than a holistic design approach that accounts for the whole system and every component inside, making sure that every part functions as intended and with its counterparts.

With our system design enablement strategy, along with our 10+ years of working with automotive semiconductor designers and developing closer relationships with other parts of the automotive supply chain, Cadence is well equipped to help you deliver automotive-grade devices and systems that support a safer, more enjoyable ride.

Further Information

Explore Cadence IP solutions for automotive designs at <https://ip.cadence.com/tensilica-automotive>

Sources

1. Research and Markets, Global Advanced Driver Assistance Systems (ADAS) Market – Forecast & Analysis, 2016-2022, http://www.researchandmarkets.com/research/rz8bjp/global_advanced
2. "ISO 26262." Wikipedia. https://en.wikipedia.org/wiki/ISO_26262
3. Hijazi, S., R. Kumar, and C. Rowen. Using Convolutional Neural Networks for Image Recognition. White paper. http://ip.cadence.com/uploads/901/cnn_wp-pdf
4. Sperling, Ed (2016, January 28) "Are Chips Getting More Reliable?" [Blog post]. Retrieved from: <http://semiengineering.com/are-chips-getting-more-reliable/>



Cadence Design Systems enables global electronic design innovation and plays an essential role in the creation of today's electronics. Customers use Cadence software, hardware, IP, and expertise to design and verify today's mobile, cloud and connectivity applications. www.cadence.com