TECH BRIEFS

SAE INTERNATIONAL

SPECIAL REPORT:
AUTONOMOUS VEHICLE TESTING

NOVEMBER 2021

CONTENTS

FEATURES

RESOURCES

ON THE COVER

Sensors are a vehicle's eyes and ears but their inputs must be analyzed and processed. Some sensor modules include an integrated microprocessor that reduces bandwidth requirements. Some lidar packages also include processors to analyze light that bounces back to the sensor, sending the main controller more pertinent data. Learn more on page 3.

(Image: Velodyne)

Peering into the **Distance**

by Terry Costlow

New sensors of all types look out longer distances — and provide higher resolutions — for engineers pushing ADAS capabilities and higher-level vehicle automation.

Sensors are the frontline technology for advanced driver-assistance systems (ADAS) and future vehicles with high-level (SAE Levels 4-5) automation. Designers at all levels are working to find the optimal and number of sensors and their ideal performance levels, capturing data that's farther away while increasing field of view and resolution.

Tier 1s and OEMs have long leveraged different sensors — radar, cameras and lidar — to provide redundancy and gain insight into what's around the vehicle. The first high-level automated vehicles for public use will be heavily laden with sensors as engineers and developers strive to prove that driverless vehicles can navigate safely. Automated taxis and shuttles are being loaded with input devices.

"When we look at Level 4-5 robotaxis, ultimately they'll be at 30 or more," said Andy Whydell, VP of systems product planning for ZF. "Some of the challenging situations are slow movement when you're maneuvering the vehicle."

The latest generation of ADAS, mostly considered to be at the high end of SAE's Level 2 classification, generally use several radars and cameras. Both the number of sensors and the technologies being leveraged will grow as automakers move to higher automation. Even suppliers of cameras and radar generally agree that most vehicles will move from one sensor modality to multiple versions of varied sensor types.

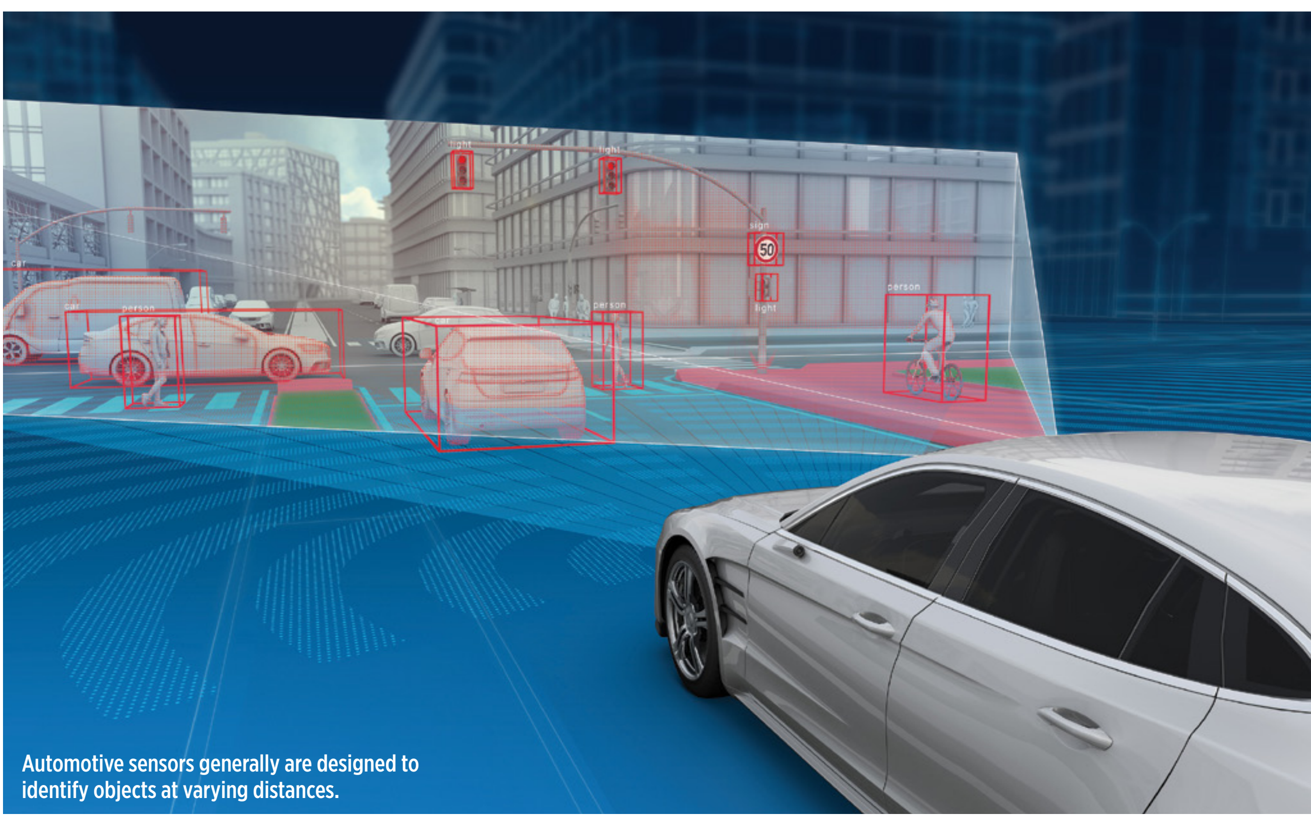

Automotive sensors generally are designed to identify objects at varying distances.

ZF

Sensor developers are working to stretch sensing systems' distance capabilities.

Continental

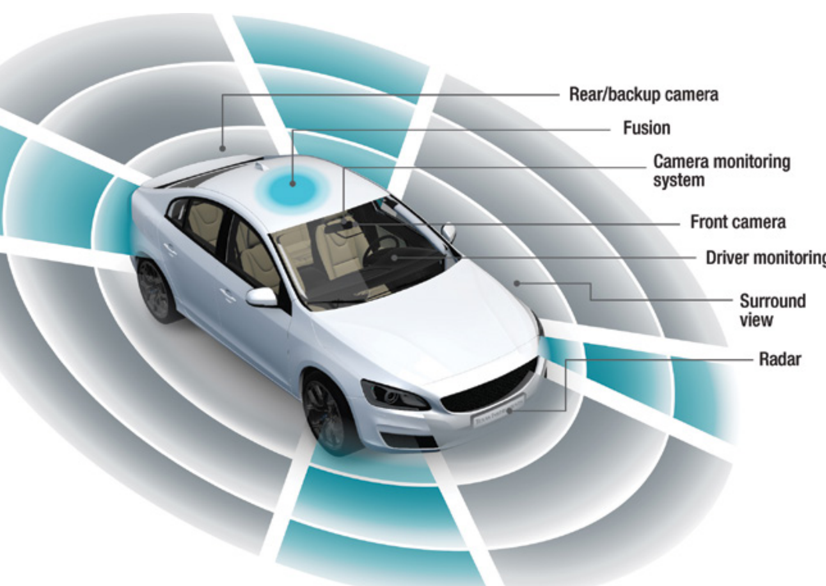

Texas Instruments

Radar can help cameras monitor front, sides, and rear.

"For [the so-called] Level 2-plus ADAS, we are seeing five radar sensors being required to achieve the ADAS functions — two in the rear corner, two in the front corner, and one in the front," said Prajakta Desai, marketing manager for Texas Instruments mmWave automotive radar. "For Level 3, additional sensors on the side would be needed for 360-degree coverage. For Level 4 and beyond, we believe that all the sensing modalities (vision, radar, and lidar) might be required to achieve fully autonomous driving."

Sensors Galore

Most front-facing sensing systems rely on a combination of sensor types. Though it's possible to provide full autonomy with a single technology, combining cameras, radar, and lidar provides some redundancy while also adding complimentary sensing capabilities.

"None of the available sensor technologies — be it camera, lidar, radar, or ultrasonic — will be able to realize automated driving functionalities on their own," asserted Arnaud

Processing Capability TOPs-out

Sensors are the vehicle's eyes and ears but their inputs must be analyzed, processed, and acted upon. A hefty amount of computing power is needed to fuse complex data from several sensors and determine how those inputs should be used.

Many, but not all, sensor modules possess little intelligence, feeding information to zonal processing modules that do a bit of processing before the data goes to a centralized controller. These powerful central controllers fuse data from cameras, radar, and lidar, then decide how to assist the driver — or automatically brake, steer, or accelerate. Meanwhile, Tier 1 suppliers are taking advantage of rapid increases in processing power.

"We've gone from 30 TOPS (trillion operations per second) to 1,000 TOPS in four to five years," said Andy Whydell, VP of systems product planning for ZF. "The industry is ramping up processing to handle all the inputs but that needs to be balanced with power consumption."

The performance gains in processors can be combined with sensors that also are leveraging semiconductor production advances. Design teams are using multiple sensor types to blend the strengths of different technologies, then they're exploiting powerful controllers to run multiple analysis tools. It's a potent combination.

"By putting cameras and radar together to check each other, it's amazing what the processors can do with the extra information, especially if you have enough processing power to run different algorithms in parallel," said Martin Duncan, ADAS division general manager at STMicroelectronics.

Stripping intelligence from sensors generates significant benefits. Sensor counts are rising, with plans for 30 or so

Velodyne

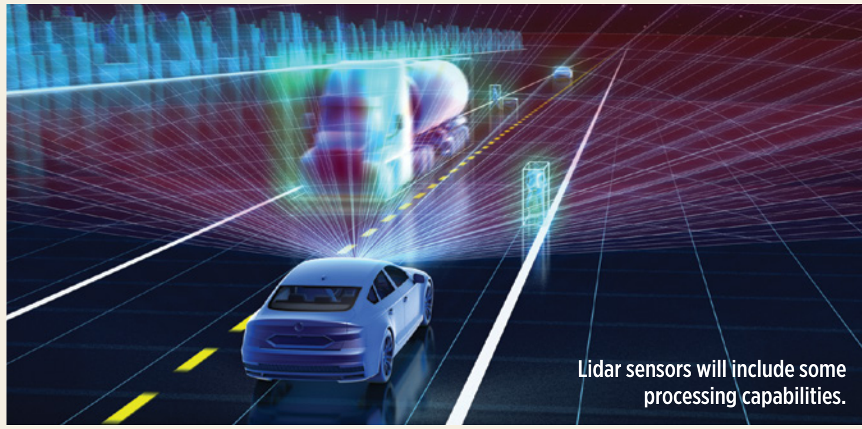

Lidar sensors will include some processing capabilities.

ZF

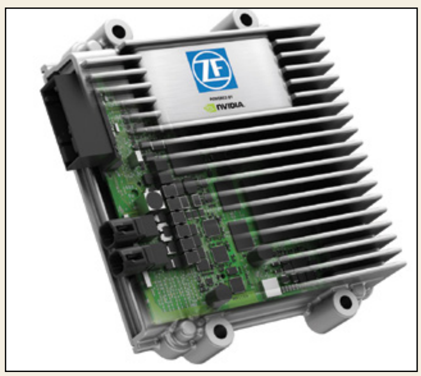

Processing power has leapt 30-fold in the past five years.

on a highly automated vehicle. Eliminating power demand and shrinking size are important factors for hiding sensors.

"Removing some processing power means sensors are smaller and power consumption is less," Whydell said. "Smaller packages make it easier to integrate sensors into lighting structures or the A- and B-pillars."

However, some sensor modules will include an integrated microprocessor. This "distributed processing" model reduces bandwidth requirements while also reducing the workload of the vehicle's central controller. Some lidar packages include processors to analyze light that bounces back to the sensor, sending the main controller more pertinent data.

"We look to subsume more functions, doing some first-level computing," said Anand Gopalan, Velodyne's CEO. "The camera guys talk about using GPUs (graphics processing units), lidar does more at the 'edge,' so systems can get away with using low-cost FPGAs (field programmable gate arrays) or low-cost CPUs."

Artificial intelligence will play a role when processors analyze sensor inputs. AI takes a fair amount of computing power but its benefits far outweigh the processing demand. Determining what's being seen by an array of sensors can be confusing for machines that only know what's been written into software.

"AI can help cameras deal with something they're not trained to recognize," Whydell said. "A European OEM [during on-road testing] came across a kangaroo that was sometimes on the ground, sometimes in the air, so the system couldn't tell whether it was a bird or not."

Ryan Gehm

Lagandré, VP of the ADAS Business Unit, Continental North America. "Furthermore, we need to understand redundancy as sensors use different physical principles. Adding more forward-looking cameras will not help when you are driving directly towards the sunset — they will all be equally blinded. Considering the top three sensor technologies currently in use (camera, radar, and lidar), you would always need to have at least two different physical principles operational to ensure safe sensing of the environment in any complex driving scenarios."

Sensor suppliers are continuing to expand their capabilities, extending distances and improving resolution. Automakers utilize CMOS imagers like those in phones, but most automotive-grade cameras remain in the 1-2 megapixel

> *"There's a lot of growth for long-range, 4D, high-resolution imaging radar — high-performance radars with more signal channels than you find in most mass-market passenger cars today."*

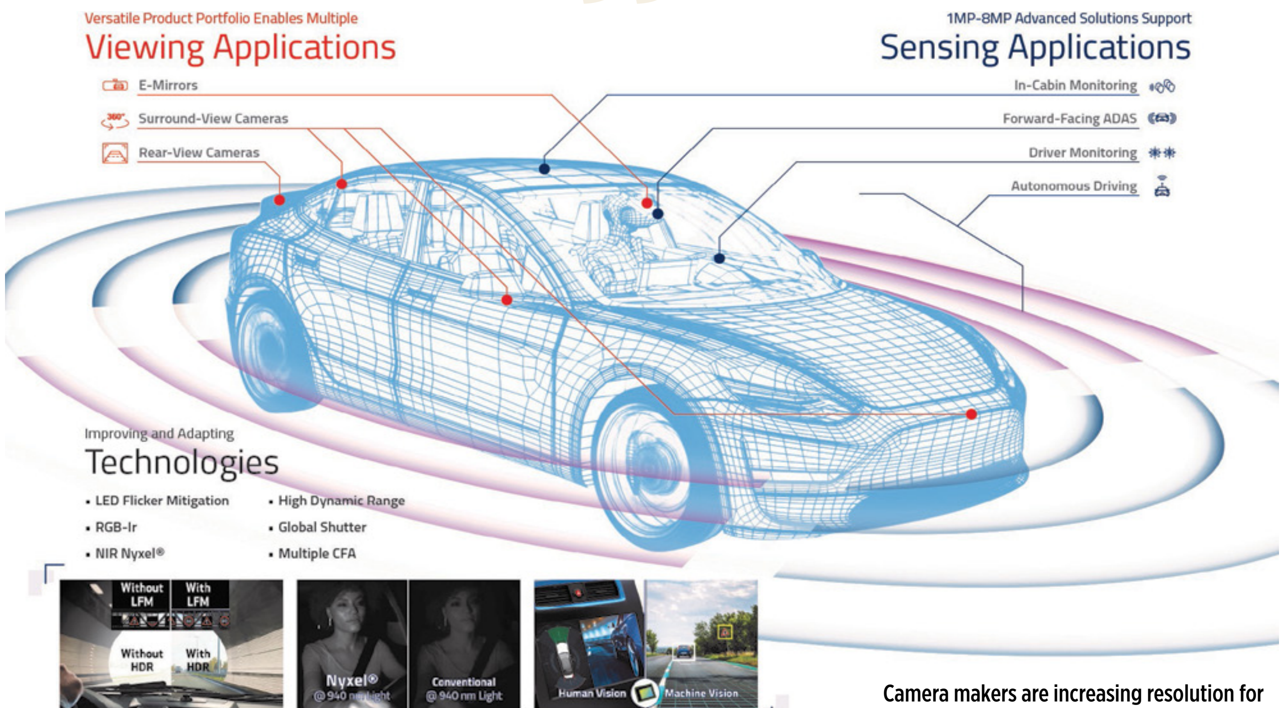

Omnivision

Camera makers are increasing resolution for automotive-grade imagers.

range. However, higher-resolution automotive cameras are beginning to ship.

"The most widely discussed challenge within CMOS imaging sensors is to make the sensor with smaller pixel and higher performance targets, such as higher resolution, and to make sure they operate smoothly in extreme temperatures," said Andy Hanvey, director of automotive marketing, OmniVision Technologies Inc. "As the level increases, the resolution increases to 8 megapixels in order to see farther distances."

Changing Channels

Radar developers are adding more channels to boost resolution and increase sensing distance. When channel counts double, the number of virtual channels soars significantly, giving systems greater ability to determine distance and identify objects.

"Radars have gone from a 2x2 channel format to 4x4, so there's four times higher resolution," said Martin Duncan, ADAS division general manager at STMicroelectronics. "The four transmitter and four receivers communicate independently and can communicate at different times, creating massive numbers of virtual channels."

These higher-resolution sensors are critical for advanced ADAS and high-level automated vehicles, experts note. At highway speeds, they offer more time to determine what's ahead, while at city speeds, the same radar devices can provide wider fields of view.

"There's a lot of growth for long-range, 4D, high-resolution imaging radar — high-performance radars with more signal channels than you find in most mass-market passenger cars today," Whydell said. "Our long-range radars have a distance of up to 350 meters. At high speeds, the beam is narrow; at slower speeds, energy shifts to a "flood," more with a wider field of view."

Tiny is Terrific

Engineers want the benefits that come from using multiple sensors but design stylists don't want sensors to mar their sleek exterior lines. Those competing concerns put pressure on sensor designers, making package size even more critical than it's been in the past. That concern ripples out to wiring harnesses that connect sensors to controllers.

Radar size has reduced dramatically over the past several years, driven in part by declining prices and a shift to higher frequencies. A few years ago, 24-GHz modules that cost well over $100 were fairly common. But 77-GHz devices that cost in the realm of $50 now are mainstream. That

Commercial Vehicle Sensor Challenges

Deere

Sensors combine with satellite positioning inputs for the specific guidance requirements of most agricultural vehicles.

Commercial vehicles share many of the same challenges as passenger vehicles but they also have additional requirements. Reliability and longevity demands increase and operating conditions in agriculture and construction are vastly different.

Autonomous vehicles have operated in mining and some other closed environments for some time. In other off-highway fields, automated systems aid operators but don't replace them. In agriculture and construction, sensing ranges often are relatively short.

"We expect most of our challenges to come in close range and involve signal-blocking structures like trees, silos, and storage sheds," said Nancy Post, director, Intelligent Solutions Group at John Deere. "As a result, our next-generation receivers will have technology that mitigates most, if not all, of the scintillation and interference issues that are common today."

Driving-related sensors on these vehicles often will be relegated to secondary status, because companies like Caterpillar and Deere have their own satellite-positioning technologies. "Our primary sensor is the satellite navigation system, which typically doesn't have the same issues with inclement weather," Post said. "During short periods where it may be affected, we can fuse the positioning information for the additional sensors like inertial navigation systems, camera, and imaging radar to bridge these gaps."

On highways, reliability over the long haul is an essential factor. Commercial vehicles have demanding environmental prerequisites that are challenging suppliers' capabilities. Getting sensors that can "see" far enough away to safeguard vehicles that haul heavy cargo payloads is not easy.

"Sensors' range is quite important for some applications, and we are always looking for sensors that are certified for safety purposes," said Luca Delgrossi, head of technology, Volvo Autonomous Solutions. "Today, there are a few certified sensors and the distance they can cover is still relatively short. Lidar's reliability is one of the factors that determines how fast we can drive under safe conditions. Different sensors can handle adverse weather conditions in different ways. It is important to understand specific sensors limitations and avoid running operations when the system cannot operate safely."

Ryan Gehm

Volvo

Certified sensors with longer sensing distances will hasten AV development for on- and off-road environments.

helps increase distance performance while trimming both size and cost.

"Moving to 77 GHz gives you smaller packages so more can be squeezed in — people are also adding the antenna in the basic package," Duncan said. "The circuit boards used for radar are expensive, so making them small reduces the price significantly."

When packaging engineers devise ways to trim size and lower cost, engineers often clamor for more components. Small packages are easier for stylists to hide. Tiny packages also can be housed in areas like headlights and by rearview mirrors, where they're less likely to get dirty or be covered by snow.

"We see a trend to integrate sensors in small spaces on vehicles, such as door handles or headlights, for better coverage," Desai said. "Antenna-on-package technology enables extremely small form factor that removes small space constraints and enables sensor integration into new places enabling newer applications and functions inside and outside the vehicles."

Wiring is another critical issue that arises in the proliferation of sensors that generate large data streams. CAN doesn't have the necessary bandwidth, while Ethernet requires bulkier and more costly cables. That's prompted many engineering teams to deploy a standard used in many phone cameras: the MIPI-A-PHY standard, developed by the MIPI Alliance, is gaining acceptance as a sensor link.

"Most sensor transmissions are largely mono-directional, so you don't need to go to Ethernet," Duncan said. "A new protocol, MIPI-A-PHY, can be quite efficient. It uses cheap cables — they're about a third the cost of Ethernet cables." ■

10 Best Practices for ADAS and AV Testing

by Jeremy Dahan

A veteran tester who has datalogged many thousands of miles in the U.S. and Japan offers suggestions for rapidly acquiring good test and validation data.

Testing for advanced driver-assist systems (ADAS) has required a completely new approach to testing. The most obvious reason for this is the sheer number of sensors and actuators involved in any given feature. Whereas engineers used to test single sensors at a time, today there are upward of 20 sensors to specify. These sensors — short- and long-range radar, mono/stereo cameras, sonar, and lidar — need to act in concert. This means that several layers of sensor fusion must be implemented.

For example, when it comes to the actuators for the braking system, there are up to seven distinct systems able to apply the brakes. As a result, a failure is not easily tracked to a single root cause but can be a complex combination of several factors. Imagine expanding this to ADAS and autonomous-vehicle (AV) functionalities such as automated cruise control (ACC), lane-keeping assist (LKA), automated emergency braking (AEB), and valet parking as

The simplicity and ease of use of the system in the hands of the driver will positively impact the quality of the collected data.

Elektrobit

Especially for AV testing, identifying unusual real-world scenarios — such as this bison the author encountered in Montana — can open new possibilities.

Jeremy Dahan

well as convenience features like blind-spot monitoring, night vision, and steering beams, to name just a few.

We have a task of gargantuan proportions in front of us. With such complexity, best practices are mandatory. The following is a "Top-10" checklist for testing today's increasingly complex ADAS and AV systems. It is based on my company's decades of experience testing systems for leading Tier 1s and automakers as well as my own recent experience logging thousands of miles in support of a new cloud-based testing and validation solution, EB Assist Test Lab, across the U.S. and around Tokyo.

1 Filter or annotate your data as soon as and wherever you can as part of a multi-step approach to annotation

Doing so will ensure that each step takes advantage of the actual capabilities of your system. This is the single best piece of advice I can share with anyone embarking on a test project. By tagging a recording in real- or near real-time around a new, unique, or unusual use-case, you can dramatically reduce the amount of data you ultimately need. This logic should be applied throughout the entire process of data acquisition.

2 Make your driver and/or passenger responsible for manual annotation

Traditionally the driver drives and that's it. Introducing ways for the driver to easily tag data (by pushing a button on the steering wheel or via voice control, for example) as it comes in makes it possible to quickly know which part of the recording might contain something relevant.

3 Run algorithms in the vehicle data logger to identify the most relevant data

Many checks can be done inside the vehicle data logger, saving time and frustration later on. These can be as simple as checking that the recording is consistent and not corrupted. Another example is if the values of signals exceed pre-defined thresholds — say, for acceleration — and therefore must be tagged. Doing so will result in a dramatic reduction of the data that needs to be uploaded.

4 Ensure your upload station takes advantage of the pre-annotated data and can properly prioritize uploads to the cloud or your on-premise data center

It will invariably take a lot of time to upload all the recorded data, whether it's to on-premise storage, a cloud provider, or some hybrid solution. You've already annotated your data, so you should take one extra step and determine which data shall go where. Data that is not tagged or annotated should not be uploaded to a high-cost storage location as it will invariably be deleted or archived later. Delete it now and save money!

This best practice is especially relevant if your teams are located in the same country or region and for urgent/priority projects, as your teams can start working on the relevant data as soon as it's uploaded. Also, be sure that the upload station provides updates on its current status to the relevant individuals.

5 Build a custom testing "platform" that meets your needs

Establishing a best-of-breed partner ecosystem is the only way

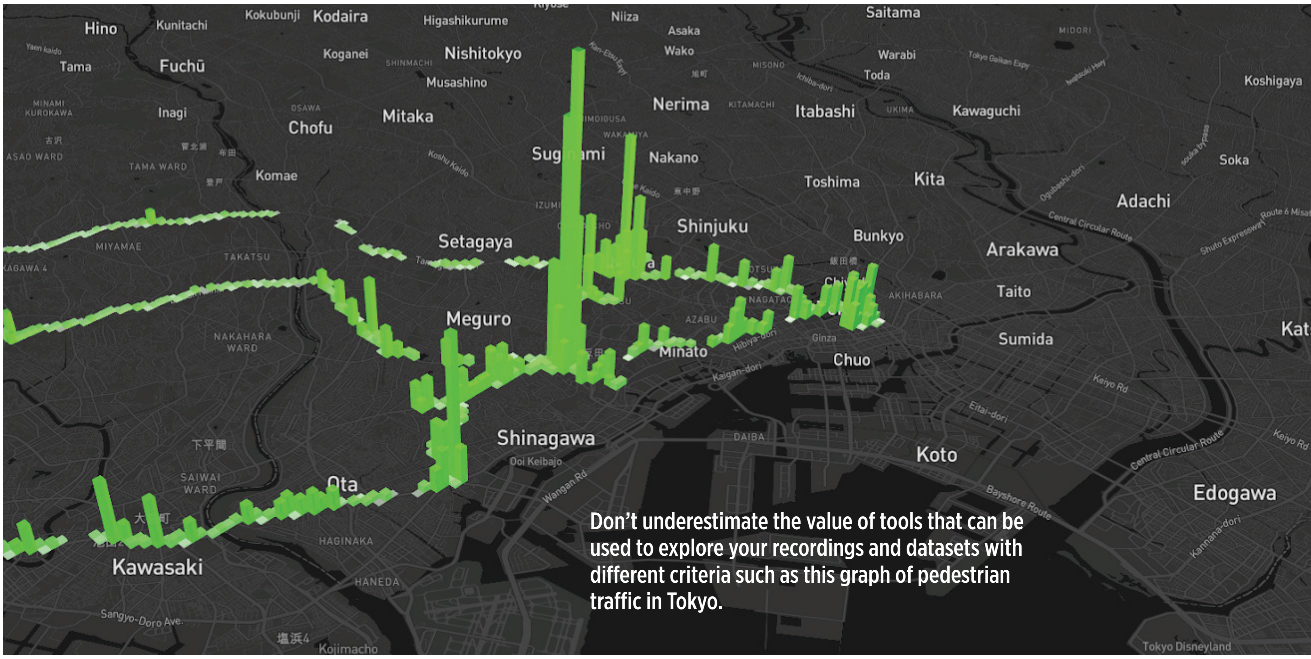

Don't underestimate the value of tools that can be used to explore your recordings and datasets with different criteria such as this graph of pedestrian traffic in Tokyo.

Elektrobit

to guarantee you can truly benefit from the different offerings on the market. Such an ecosystem should include a fast-upload datacenter, hyperscalers, and cloud providers to handle all potential use cases as well as software and hardware partners.

Partnering with companies that are experienced and understand your needs will ensure your resources are well utilized. There is no one-size-fits-all solution and each company has different needs — and needs often differ even inside a single corporation.

All organizations, however, have a few common goals when it comes to their testing approaches and these include flexibility, cost effectiveness, and speed. Depending on your use-case, each of these eventually will play a role.

The approach used by some OEMs and Tier 1s is to try to build it all themselves but this isn't efficient or cost-effective. These companies should instead focus on improving their products and bringing them to market more quickly and efficiently. In the end, this is what will differentiate them from their competitors.

6 Create a direct pipeline from your logger to the cloud that will allow the driver to share a snippet of a recording if something unusual is happening

This direct pipeline approach ensures that an engineer can investigate the issue ASAP and provide a fix or insights as well as determine whether the test drive is still relevant. This is critical to avoid dealing with unnecessary data. Examples range from a wrong version of software running on the ECU, to the updated software behaving unusually under certain conditions such as low light.

7 More data is not better — the diversity of your dataset will drive your success

Being able to identify how much coverage you have and being critical about your training dataset will ensure the best results. You must clearly define metrics to rank and evaluate your data.

Don't underestimate the value of tools that can be used to explore your recordings and datasets with different criteria such as a heatmap of the position of traffic signs, showing where they were collected and at what time of day. Several companies collect and archive as much data as they can for Hardware-in-the-Loop (HIL) tests. However, after analyzing the archived data from one vendor to determine its true value, I found that a fraction — less than 10% — was relevant. The rest was duplicate information and therefore resulted in a very expensive, time-consuming process. Another important aspect is the quality of the recorded data. You must record your data with high accuracy (25 ns) so that you can later replay it in HIL farms with a microsecond accuracy.

8 Test-drive and simulation data should be treated equally

Being able to address them equally is key. By this I mean that your data-management system should be able to search regardless of the source of the data. I strongly recommend securing relevant data from third-party companies. While the sensor setup won't be the same, the extracted scenario will be valuable.

9 Identify unusual real-world scenarios to open up a new world of possibilities

A few examples of scenarios I've personally encountered include

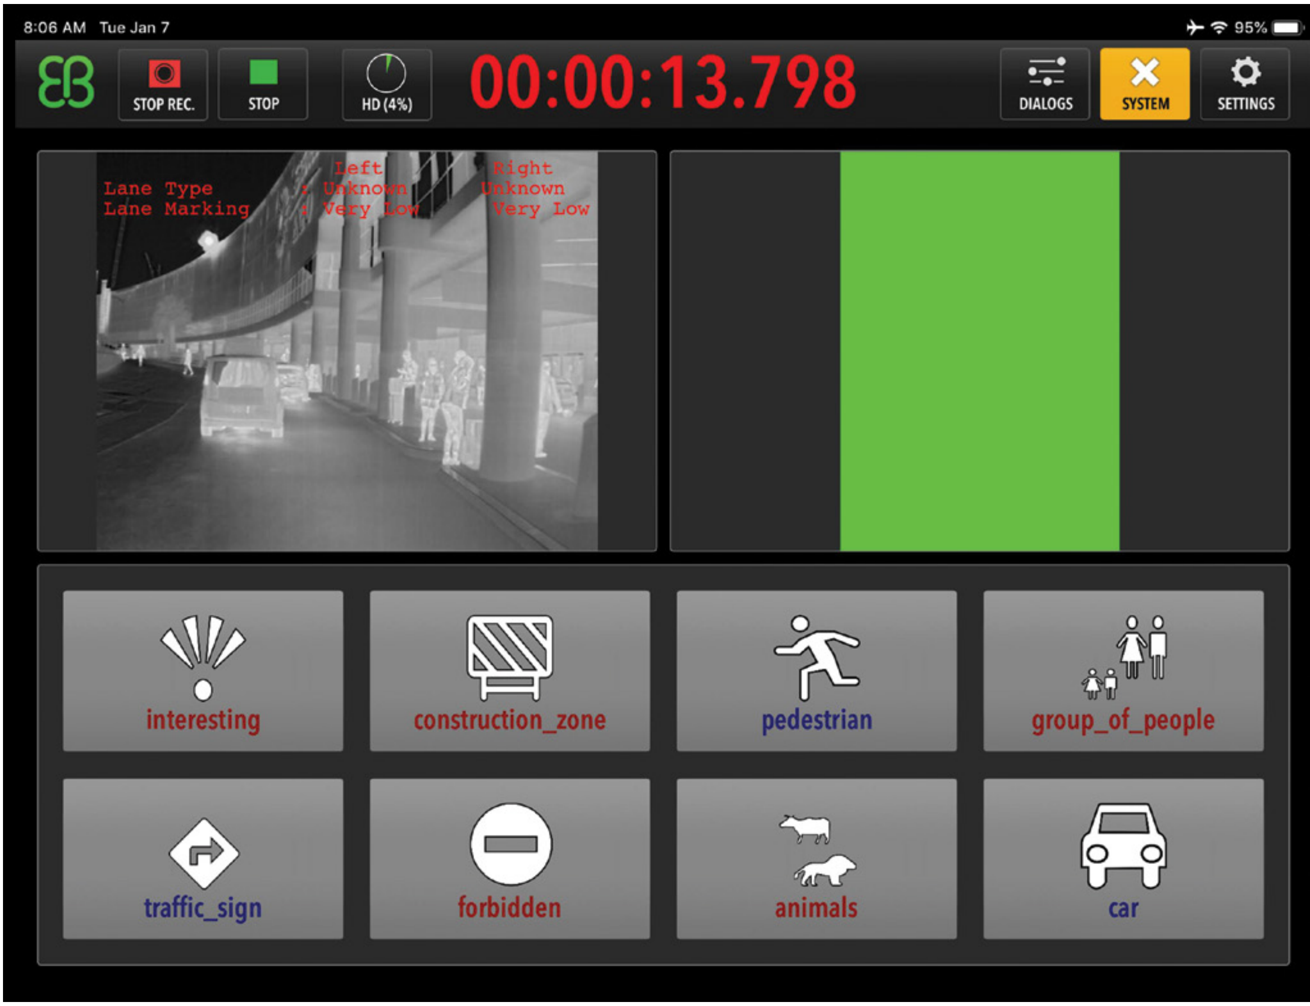

Elektrobit

Filter or annotate your data right away using a system that allows the driver or passenger to easily tag data using simple pushbutton or voice control.

life-size imagery on side of a large semi-truck, chrome that acts like a mirror, a mannequin that has been ejected from the back of pickup truck, a plow with a trailer that is drifting sideways, and a police car reducing the speed of traffic by continuously changing lanes with its sirens on.

These examples can later be reproduced in simulation and diversified in thousands of variances. The challenge is to only create plausible scenarios that can be enforced by the simulation engine and its rules and/or by the description language of a scenario used.

10 Future-proof your toolchain using industry standards

To ensure the toolchain you're creating is usable by others and future-proofed, look to industry standards such as Open Drive, Open Scenario, and Open Label. Doing so will also ensure it is interoperable with other systems you may wish to use later. Assuming you're achieving your current objectives by using clean interfaces and standards, these will also enable you to quickly onboard a new partner at any time.

The Goals: Rapid Feedback, Good Data

If there is one key lesson learned from my time as a road warrior, it's that the driver can simplify the process for all involved by playing a more active role. While often overlooked, the simplicity and ease of use of the system in the hands of the driver will positively impact the quality of the collected data.

As the driver, you're responsible for the entire chain of events. From personal experience, I can tell you that I want to find something I've never seen before. As a driver, you also want to get quick feedback on what shall be collected and understanding the why enables you to obtain better data. ■

Jeremy Dahan is head of technology research at Elektrobit (EB), where he manages business development for the company's efforts and partnerships in the Silicon Valley. He has logged thousands of miles driving across the U.S. and around Japan in EB's test vehicle in 2019 and 2020.

Mitigating **Radar-to-Radar** Interference

by Alessio Filippi, Francesco Laghezza, Feike Jansen, Jeroen Overdevest, and Dilge Terbas

An effective radar interference mitigation strategy should have the right balance between complexity and capability for dealing with the interferers.

Sensors are a fundamental building block of highly automated and autonomous-driving vehicles. And each different sensor — radar, lidar, and camera systems — has strengths and weaknesses. Radars are typically more robust than cameras to adverse weather conditions. They're also more cost-effective than lidar and provide a better velocity and distance resolution. However, typical drawbacks can be lower angular or cross-range resolution and radar-to-radar interference.

There are advanced solutions that further increase the radar angular resolution, from sparse arrays to

NXP

Solutions exist to improve the radar's sensitivity to distinguish small targets from strong targets nearby such as a motorcyclist entering the FoV among passenger vehicles and large trucks.

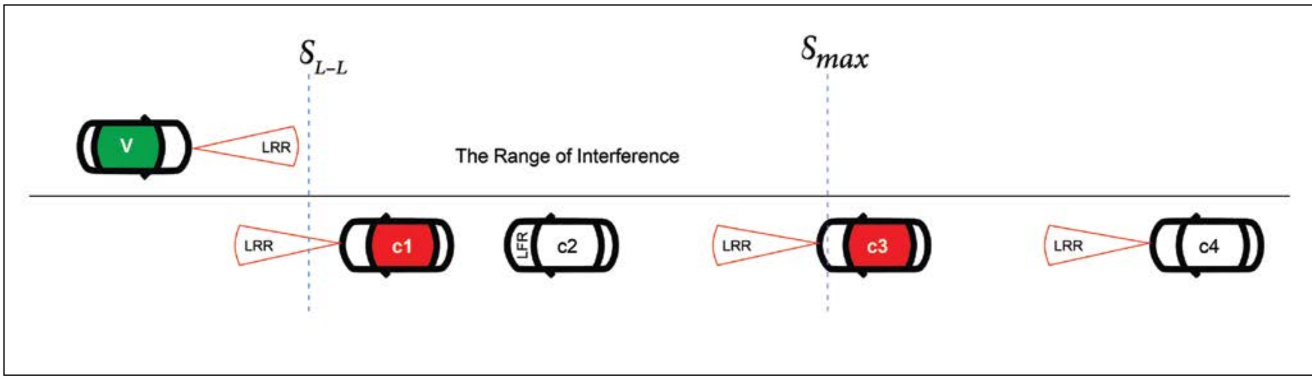

NXP

Figure 1: The simple common traffic scenario in which the interference is calculated at the victim radar (green) and is generated by the other vehicles (red).

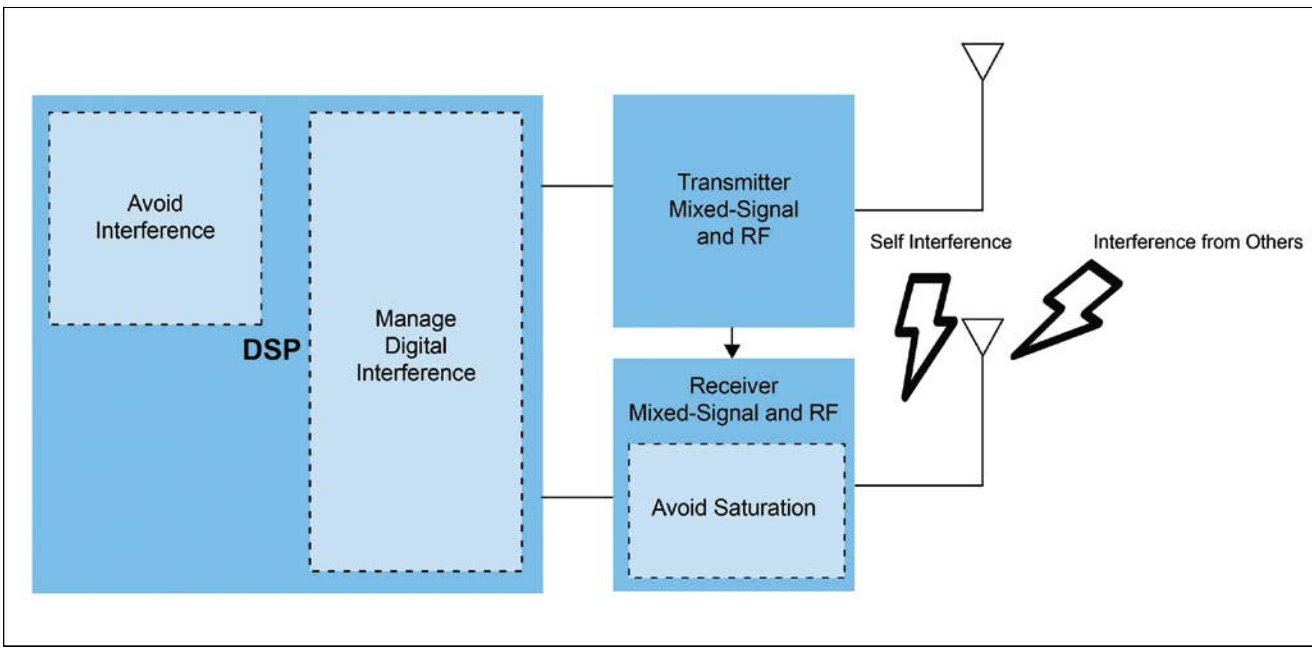

NXP

Figure 2: A radar sensor with a digital part (DSP), the transmitter, and receiver front end.

super-resolution algorithms. Solutions also exist to improve the radar's sensitivity to distinguish small targets from strong targets nearby such as a pedestrian next to a large truck or a building. As more and more vehicles are fitted with radar systems, interference is of growing concern.

Radar interference is a well-recognized challenge. As all radars use the same allocated frequency spectrum, multiple radars may transmit at the same time and frequency. If they also share a common visible path, they will interfere with each other.

The MOre Safety for All by Radar Interference Mitigation (MOSRIM) report is possibly the first project to assess the severity of the radar interference problem. A National Highway Traffic Safety Administration (NHTSA) study builds on that and adds different traffic scenarios to determine when and how often interference happens.

A common traffic scenario is shown in Figure 1: Vehicles that have long-range, front-facing radar with a relatively narrow field of view (FoV) approach each other. Interference happens as the FoV of the two sensors overlap and the radars are active simultaneously and in the same frequency spectrum.

Radar Interference Techniques

There are three classes of techniques available to mitigate radar interference without affecting overall system performance, as shown in Figure 2. The first class of

> *As more and more vehicles are fitted with radar systems, interference is of growing concern.*

techniques avoids the saturation of the front end, which happens when a radar sensor is exposed to a strong interferer. The second class manages digital interference by recognizing and removing the interference in the digital domain. The third class avoids interference by adapting the radar waveform to reduce the probability of interfering with other radars.

These techniques attempt to mitigate the adverse effects of interference before or after it happens. As they are implemented to every sensor in their respective radar system independently, they do not provide a robust strategy to avoid interference in a structured manner.

A possible approach could be to statically allocate resources to radar applications. For instance, front and rear radars could use non-overlapping parts of the spectrum to avoid the worst-case interference scenario of a front-looking radar illuminating the rear-looking radar of the car in front. Using different polarizations for different applications might mitigate this worst-case scenario as well but proper considerations on the influence of polarization on the FoV of the antenna and on the propagation should be considered.

For more advanced solutions, the radar sensor could borrow channel access rules developed by the telecommunication industries, such as TDMA, FDMA, CDMA, or OFDMA. Alternatively, the solution could take a more randomly-determined approach, using protocols such as ALOHA, CSMA, or CSMA-CA. This is usually deployed at the medium access control (MAC) layer of the communications stack.

Following the former approach would require a form of centralized coordination, which means that every radar sensor should communicate to a central unit, e.g., the telecommunication infrastructure, what it would like to do. It should receive the time and frequency slot back from the infrastructure that it can safely use to sense the environment and limit the interference.

The latter approach is distributed, whereby each radar sensor follows the same transmission protocol, which ensures fairness and a given performance. The ALOHA protocol is the most straightforward approach. In it, each unit transmits when it wants to, which is basically what the current radar sensors do.

ALOHA Simplicity

One of the first improvements of the ALOHA protocol in communications is the slotted-ALOHA. If, for instance, a collision of two transmitted signals occurs, it happens for the entire transmission frame; there are no possible partial overlaps of messages. This simple trick already doubles the efficiency of the transmission in communications.

As compared to communications, radar sensors do not need to communicate with each other to sense the environment. One could, therefore, only standardize the way they access the channel, leaving the full freedom to deploy any radar waveform. For instance, going back to the ALOHA and slotted-ALOHA example, one could think of organizing the wireless resources in time and frequency blocks; for example, 20 ms by 250 MHz. This technique would allow any radar sensor to use an integer number of time/frequency slots. The time synchronization could come from the GPS signal, which is already available in the car.

If agreed by all parties, this simple measure could allow more sensors to act in the same environment while giving the industry the full freedom to differentiate in terms of radar waveforms.

More complex and efficient multiple access schemes could also be envisioned. For instance, a form of sensing of the available time and frequency resources and a mechanism of randomization in accessing the channel, similar to the CSMA-CA protocol used in Wi-Fi networks. Further exploration is required to identify the proper resource allocation protocol for radar applications. It should be based on the MAC protocols but with significant adaptations to consider the different kinds of traffic, priority settings, and quality of service targets.

Whatever the solution will be, an agreed way to access the shared resources will surely increase the maximum number of radar sensors co-existing in the same environment.

NXP

An additional step to eliminate the interference is to look at the radar sensors in traffic as multiple units of the same sensing system. In this vision, the radar sensors work together to achieve the same shared goal.

Cooperative Sensing

An additional step can further be taken to eliminate the interference. Instead of looking at the radar sensors as individual units, one could look at the radar sensors as multiple units of the same sensing system. In this vision, the radar sensors work together to achieve the same shared goal. They will need to be equipped with a communication link to all other sensors to coordinate access to the shared channel and share information, acting as a larger ecosystem.

An example of how this would work in a scenario would be when an application needs to have a good image of the environment. It could query the radar sensing system to build that image; however, the information might come from a radar sensor mounted on a different car or road infrastructure. Such cooperative sensing would avoid the interference a priori because no unit is seen as an interferer but they all work together.

The proper radar interference mitigation strategy does not need to be the most efficient. Having the right balance between complexity and capability of dealing with the interferers strongly depends on the foreseen scenarios. Combinations of randomization, detection, and avoidance of interference might be powerful enough to support a significant market penetration of radar sensor systems.

There will be more elaborate techniques to deal with radar-to-radar interference. Eventually, there will be some form of agreement within the radar sensor community to share the sensing resources effectively. Ultimately, there will be a standardized way to access the communications channel while, at the same time, maintain the possibility to have differentiating sensing performance. ■

For further information, see: https://www.nxp.com/applications/automotive/adas-and-highly-automated-driving/automotive-radar-systems:RADAR-SYSTEMS

Data Drives **Driverless Truck Launch**

by Ryan Gehm

Smart diagnostics and advanced validation help support the reliability metrics required to gain confidence that autonomous trucks are ready for the road.

Transforming raw data into high-quality structured data is a critical path to properly fueling machine-learning (ML) models and deploying artificial-intelligence (AI) applications across autonomous fleets. Companies are working to overcome data challenges to ensure their ML algorithms can produce the AI required to achieve widespread SAE Level 4 and 5 automated-driving operations.

"When our trucks drive on the road, they're collecting terabytes upon terabytes of data and we need to get that up into the cloud and into the hands of our engineers, ultimately," said Brandon Moak, co-founder and CTO of autonomous technology developer Embark. The startup uses "active learning" techniques to identify the most relevant detections and provide the most useful insights into critical edge cases.

"You can think of active learning as a way for us to understand the ways in which our machine-learning models are failing," Moak explained. "We can actually sample our data using this technology to build high-quality datasets that are lower in volume but higher in quality to get more performance out of our systems."

The level of preprocessing required to make sure the raw data is useful for machine learning is a key challenge, said Tom Tasky, director of Intelligent Mobility at FEV North America. The supplier has a patented preprocessing and analytics

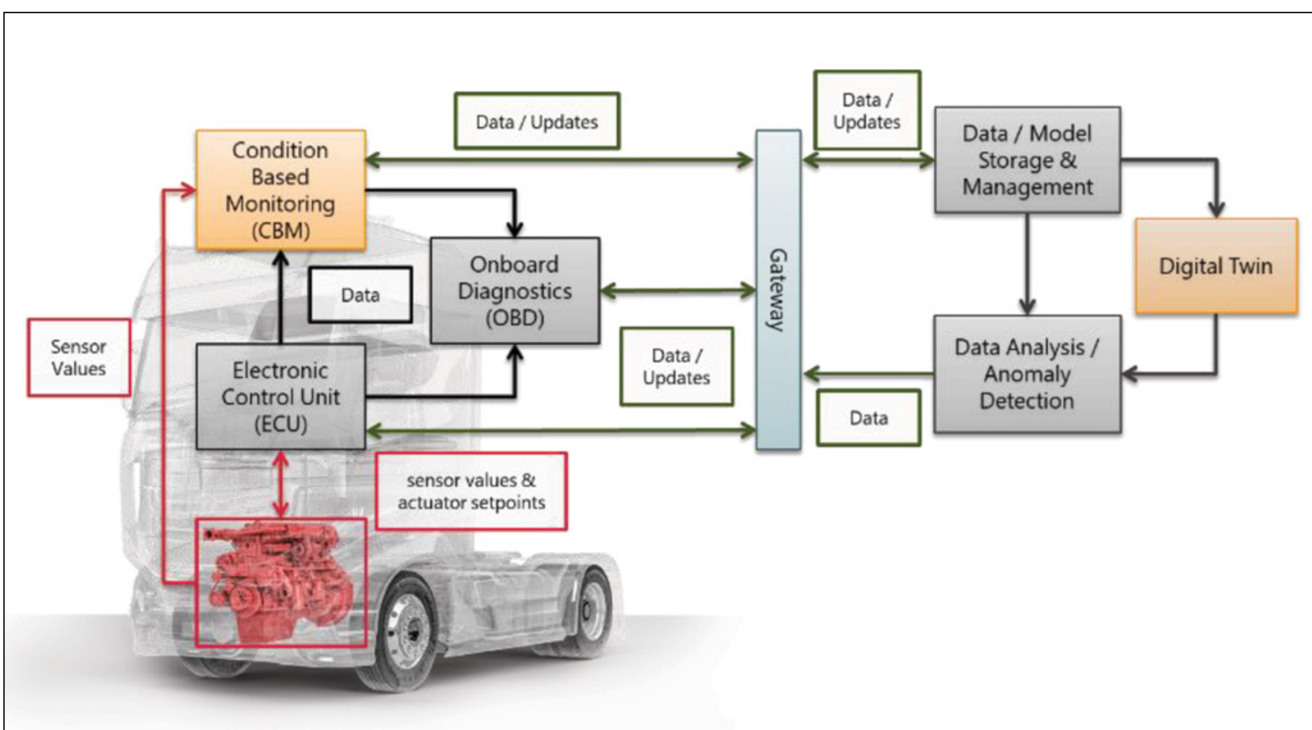

FEV

FEV employs a digital twin fed with data from the fleet to identify potential failures early in the development process.

Embark

Embark uses "active learning" techniques to identify the most relevant detections and provide the most useful insights into critical edge cases.

solution that handles up to 40 TB of data per vehicle per day in an L3 Pilot project taking place in Europe. The amount of time and effort involved in this part of the process can be underestimated, he said.

"Once you start developing and looking at the data, you see it might be poor quality and how much additional software effort is required," Tasky said. "To really understand the sensors, the quality output, any limitations in certain environmental conditions, things like that really need to be factored in to make sure you have time to account for it, prior to running some expensive tests."

Field data analytics are extremely valuable for the development and production of optimized components, as well as to identify design weaknesses, status analysis, and predictive maintenance. "Especially if you look at ADAS features, these are new components being developed and you don't necessarily have the same reliability data in different applications being introduced," Tasky explained. "Discovering that information is extremely valuable to [determine] trends and the life cycle of these new components."

The digital twin also is a useful tool in identifying potential failures and determining root causing issues. A complicated example, according to Tasky, is when there is onboard monitoring with offboard failure analysis and a digital twin.

"You enable this through the use of a gateway that has the connectivity aspects to communicate to the digital twin in the cloud," he said. "There's a lot involved with setting up this infrastructure, which we help our customers with today. But the value of this is identifying failures well in advance during development or even field data in fleets."

Smart Diagnostics

As automated-vehicle (AV) developers began combining more advanced sensors such as lidars, infrared cameras, and L4-specific sensors for redundancy and higher integrity, diagnostic capabilities lagged initially, according to Ananda Pandy, technical specialist for ADAS and autonomy at ZF.

"The focus was on improving the computing ability and hardware development for the 'virtual driver' and how to scale that development," Pandy explained. "The diagnostic capabilities of the vehicle actuation systems were still at the same level as how it was developed for the core ADAS functions and were dependent on the safety driver in the vehicle during these development phases."

As AV development enters the "shakeout" phase, where the integrated vehicle platform accumulates significantly more mileage and infrastructure setup begins, the focus shifts more to creating the reliability metrics necessary to build confidence prior to rollout.

> *Once you start developing and looking at the data, you see it might be poor quality and how much additional software effort is required.*

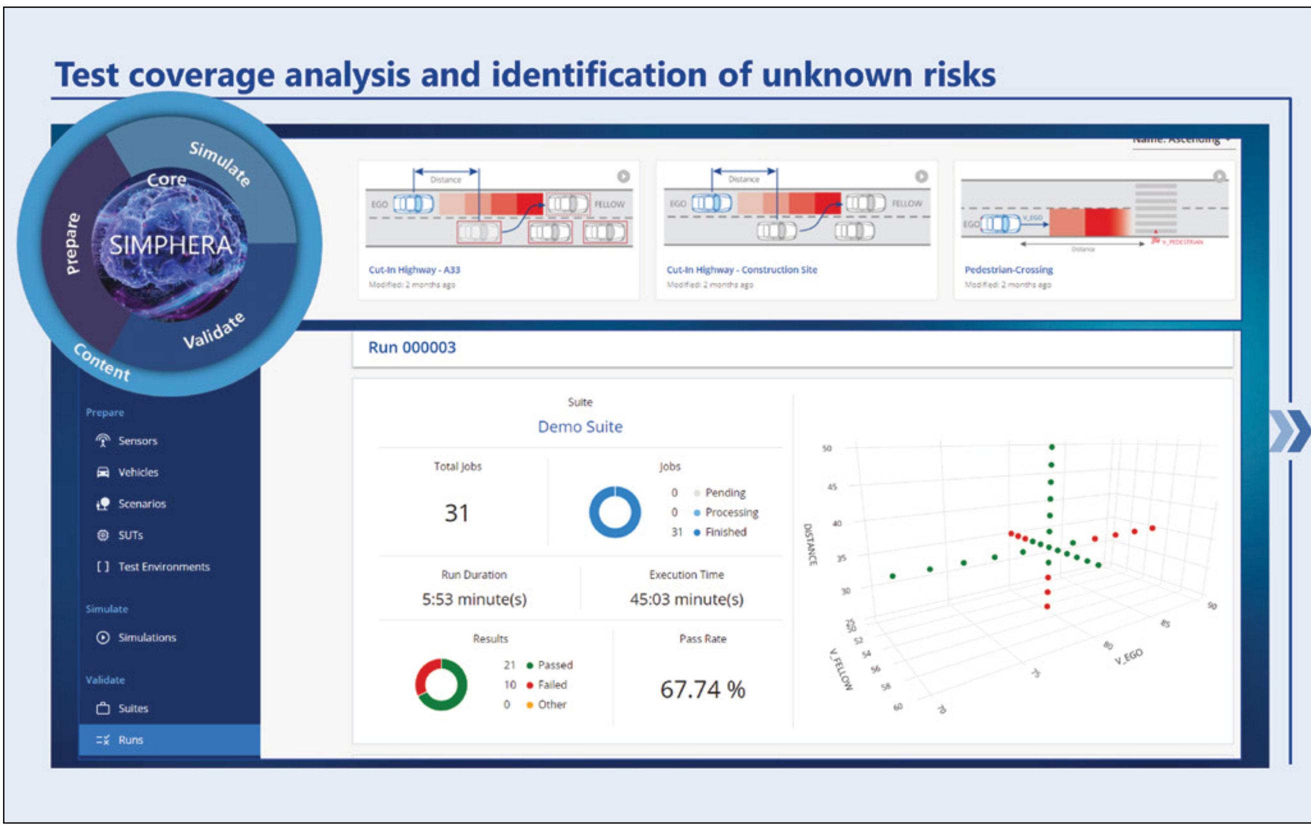

dSPACE

SIMPHERA can simulate and validate millions of test cases in the cloud, offering scenario-based evaluation of AV-system safety according to SOTIF.

"Diagnostics play a major role in supporting these reliabilities, and it's not just the number of miles driven without interventions or the number of trips being completed," Pandy said. "It's imperative to have the predictive diagnostics and not have any latent failures in the system in order to make the call for the driverless launch."

Smart diagnostic features should include a "self test" at the beginning of the journey to check for any pre-existing faults, Pandy said. A passing test can be a condition required to enter autonomous configuration. Self-test can be done either by simulating inputs internally within the fail-operational steering system, for example, or can be executed by the virtual driver before requesting autonomous configuration.

After a successful self-test, "dynamic diagnostics" can help continuously check for validity of the inputs required to detect for normal steering effort. Feedback can be provided to the virtual driver and vehicle trajectory planning can be enabled to handle any latent faults. "For example, by going to a safer speed before entering a tighter curve or ramp. This can help ensure safe maneuvering of the vehicle," he said.

Another important feature is a "fail operational window" that can be different for different actuation systems. "A general best practice could be to report the validated time that is available for the virtual driver to bring the vehicle to a safe stop when fail operation is initiated by a particular actuation system,"

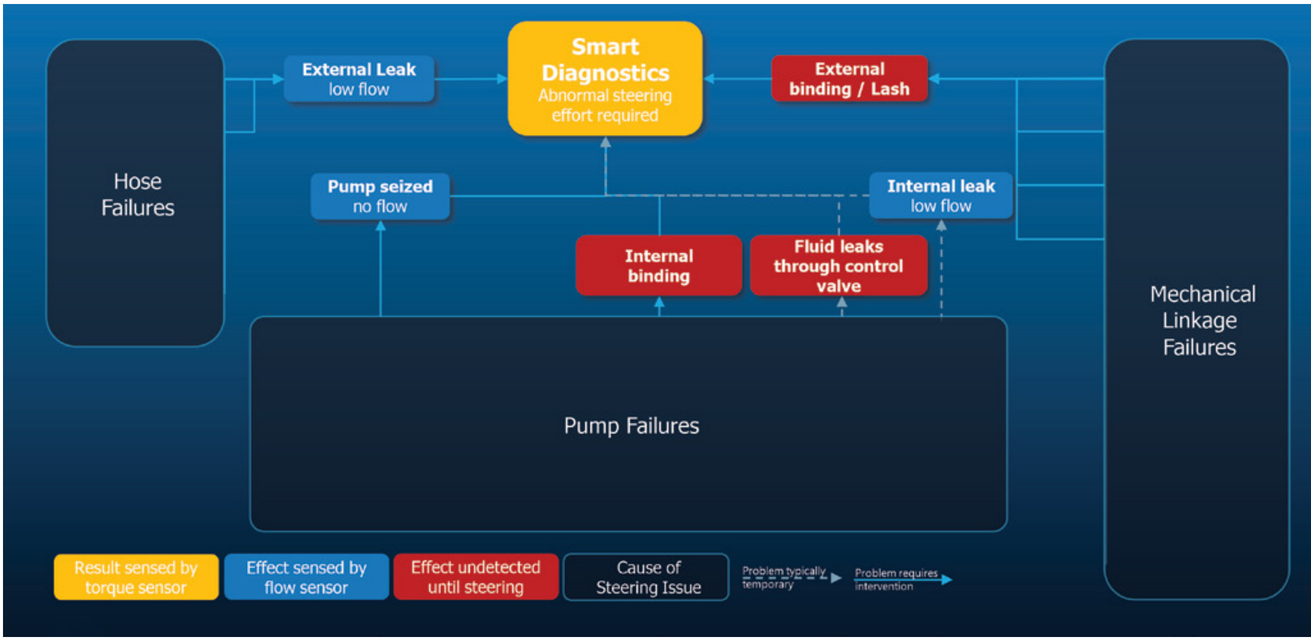

ZF

A conventional steering system is transformed for higher levels of autonomy. It's critical to understand how much of the existing diagnostics from ADAS can be reused, said ZF's Ananda Pandy.

Pandy noted. "This information can be used by other actuation systems, or it could be used for compliance purposes to ensure that the vehicle was indeed brought to a safe state within the fail-operational window that was provided."

Over-the-air troubleshooting during the early stages of development and integration work also is key. "A master diagnostic message that consolidates the status of a fail-operational system, including the fault codes, fail-operational time window, and performance measures, can help other subsystems to plan for safe actions as well as roadside inspections," Pandy said. "This can be similar to the existing diagnostic message, but it is tailored for AVs such that it can be easily integrated with the V2V or V2I [vehicle-to-vehicle or vehicle-to-infrastructure] communications."

Release Testing and Homologation

A sequence of testing — from model-in-the-loop (MIL), software-in-the-loop (SIL), component, domain, vehicle hardware-in-the-loop (HIL), and finally, real-vehicle testing — is necessary to achieve the software quality and robustness required for ADAS and AV systems.

"It's about running a lot of tests and finding the critical edge cases that you need to validate the software that needs to be deployed," said Jace Allen, director of ADAS/AD Engineering and Business Development for dSPACE, Inc. "Integral to all of this is really trying to get a truck to certification or homologation."

dSPACE participates with different organizations including ISO and collaborates with companies such as TÜV to offer broad expertise from system requirements to release testing and homologation. The company also expanded its partnership with BTC Embedded Systems to offer a new web-based solution, SIMPHERA, that uses simulation to validate and homologate autonomous-driving systems.

Available for use in general customer projects starting second half of 2021, the SIMPHERA simulation and validation environment integrates the BTC ScenarioPlatform that creates scenarios and generates and evaluates tests based on coverage. The high level of abstraction makes it possible to express thousands of test cases with just one abstract scenario, the companies claim.

BTC's automated test generation functionality uses advanced technology such as model checking, AI, and intelligent weakness detection, which allows test cases to be generated based on statistical methods and meaningful coverage metrics. Compared to random or "brute force" test-generation approaches, this strategy "considerably" reduces the amount of test data and delivers clear metrics, says BTC, even with regard to future homologation criteria.

"It's the same methodology we've talked about for SIL and HIL — test asset reuse," Allen explained. "I can define my sensors, my vehicles, my scenario, the interface to my SUTs [systems under test] and then run whatever simulations I want so that I can find those edge cases, so that I can evaluate my AV-system safety according to SOTIF [Safety of the Intended Functionality] and so forth." ■

RESOURCES

Autonomous Drive Emulation

Applying C-V2X test solutions across the automotive workflow with the ADE platform.

Autonomous driving holds enormous potential for transforming the way we get around. Further, it promises to forever change not only vehicle design and manufacturing but also automobile ownership and, indeed, the overall business of transportation.

Achieving the goal of fully autonomous driving involves the development of highly complex software infused with artificial intelligence (AI) that can correctly interpret and act upon streams of real-time data arriving from the surrounding infrastructure and emanating from arrays of in-vehicle-based sensors. One consequence is that thorough verification of the functionality, performance and safety of such systems will increasingly depend on detailed simulation and testing in the lab. This can be utilized throughout the automotive workflow, beginning long before new advances are deployed in vehicles operating on public roadways.

The types of sensors most associated with advanced driver-assistance systems (ADAS) include radar, LIDAR, ultrasonic, and cameras. External inputs arriving via vehicle-to-everything (V2X) radio add important data into the mix. Per the relevant standards, the functionality and safety of systems that incorporate V2X must be verified across a variety of situations and conditions. As the breadth and depth of such testing increases, it quickly becomes too expensive, impractical, and risky to contemplate the use of actual vehicles operating on closed tracks or public roadways.

The amount of testing applied to any new vehicle design has always been massive but that burden becomes enormous when the amount of technology in a car is quadrupled, with terabytes of data moving within a single car every day. OEMs and their suppliers need a way to test in a closed-loop environment using real-world signals in a lab.

Keysight's ADE platform brings the road to the lab, thereby enabling the testing of real sensors with real data in a closed-loop system. Ultimately, this translates into confidence, cost savings, and a competitive edge in the race to achieve fully autonomous transportation over the roadways. ■

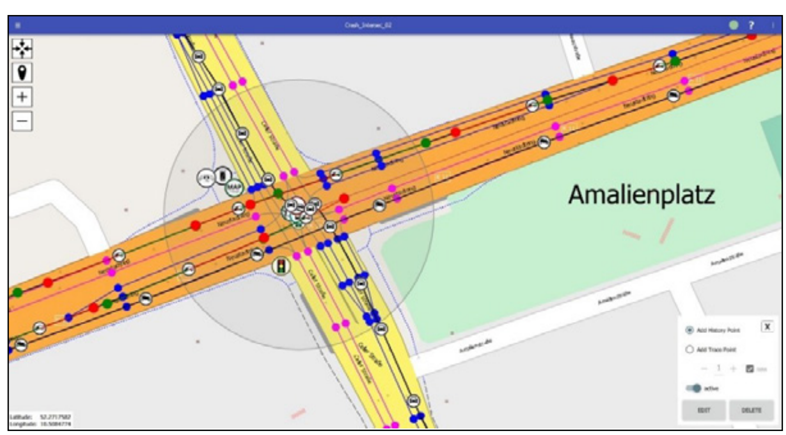

Download The Complete White Paper

Accelerate the Development of Advanced Driver-Assistance Systems

The ADE platform closes the loop for highly realistic testing.

Forecasts for the adoption of connected autonomous vehicles (CAVs) remain optimistic. However, those predictions are dependent on unshakable confidence in the minds of consumers, regulators, and the insurance industry. Building confidence in advanced driver-assistance systems (ADAS) means hundreds of millions of miles of road testing — actual or simulated — to fully explore corner cases and thoroughly validate new designs.

Keysight's Autonomous Drive Emulation (ADE) platform is the environment emulator for in-lab testing versus realistic roadway scenarios, from mundane to one-in-a-million. Using total scene generation, the platform exercises ADAS software using time-synchronized inputs to the actual sensors. Its open architecture also closes the loop with your existing HIL systems and 3D modelers, enabling you to keep pushing ADAS towards Level 5.

ADE generates total roadway scenes with multi-faceted environment emulation. It fills the gap between software simulation and on-the-road testing of ADAS capabilities by performing holistic validation of the actual sensors, ECU software, AI logic, and more.

Accelerate the development of new ADAS software features and gain deeper insights into the ADAS software behavior earlier in the development cycle. Develop greater confidence in the validation process by thoroughly validating line-of-sight-based sensors such as radar and cameras and with synchronous testing of communication-based systems (e.g., C-V2X).

Cover dangerous situations, risky corner cases, and more in the lab by emulating complex situations with time-synchronized simulation of the real sensors to be used in vehicles. Find potential issues earlier in the development process, reducing the likelihood of post-release failures. Ultimately test more scenarios sooner and achieve greater confidence in required ADAS functionality.

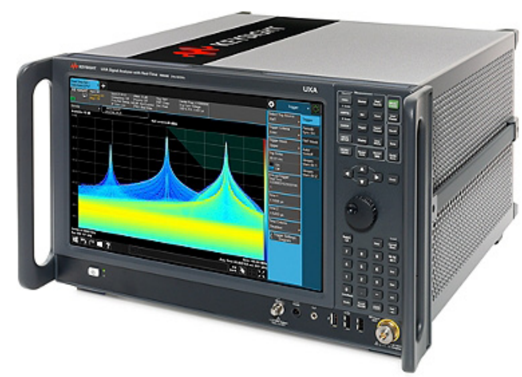

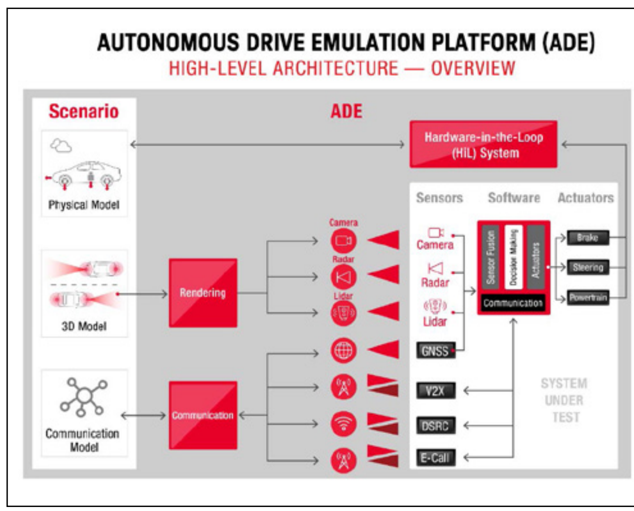

The ADE platform allows testing of an ADAS starting from the component level and progressing to the sub-system level, long before a vehicle is tested on the road. The goal is to decrease the overall cost of testing while improving test coverage at an early stage in verification. Ultimately, this helps improve the performance of ADAS when deployed on public roadways. ■

Download The Complete White Paper

All Things 5G: C-V2X

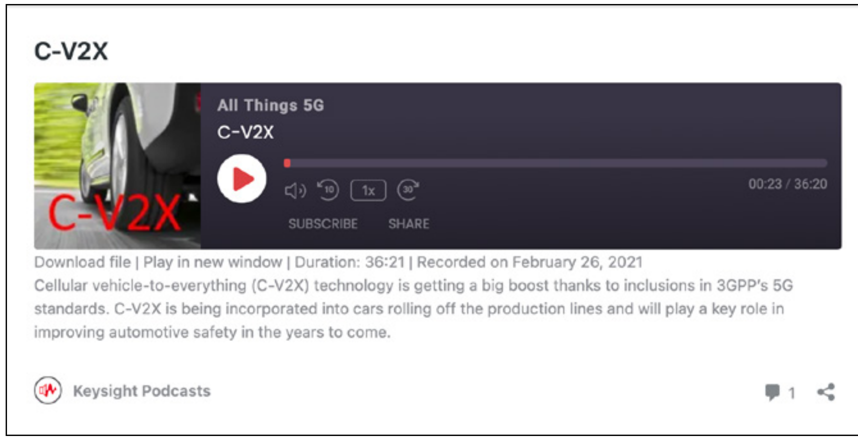

Cellular vehicle-to-everything (C-V2X) technology is getting a big boost thanks to inclusions in 3GPP's 5G standards. C-V2X is being incorporated into cars rolling off the production lines and will play a key role in improving automotive safety in the years to come.

Keysight's test and measurement solutions for C-V2X help accelerate the design and manufacturing of wireless technologies critical to autonomous driving. The latest episode of the All Things 5G podcast addresses how Keysight meets 5G test challenges for C-V2X module and communications systems developers.

Listen to the Podcast

Next-level total scene generation technology for stimulating ADAS sensors starts here

Moving to Level 3 autonomous driving and beyond demands a better, faster way to verify and train your advanced driver assistance systems (ADAS) algorithms. Training these algorithms requires testing against millions of complex radar scenes covering a wide variety of driving scenarios.

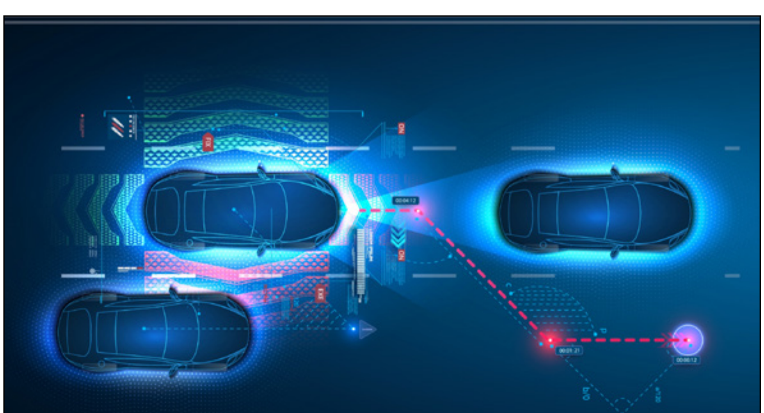

Join Keysight as we demonstrate a new technology for in-lab testing of realistic roadway scenarios, from mundane to one-in-a-million. Emulating 40+ objects simultaneously, the technology promises to help you shift testing of complex driving scenarios from the road to the lab to accelerate algorithm learning.

Register today to learn more about how you can sign off on your new ADAS functionality with confidence.

Register Today

Tech Talk: How to Generate Scenarios with Automated Processes

How can you generate scenarios with automated processes?

This is the question host Anna Jeske explores in a Tech Talk with Alexander Frings, team leader of product management engineering services at IPG Automotive and an expert on scenario generation.

He provides insights into the ScenarioRRR toolbox that supports extremely efficient scenario generation and testing, and advances the validation process for driver assistance systems.

Watch the video