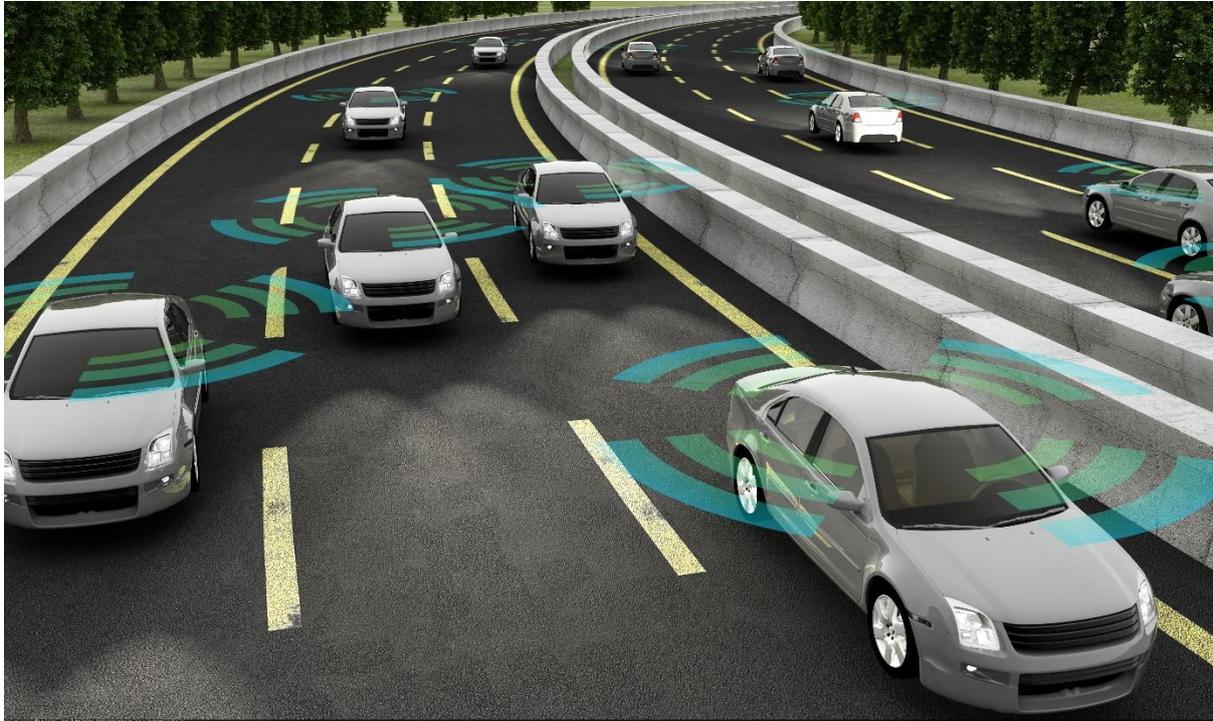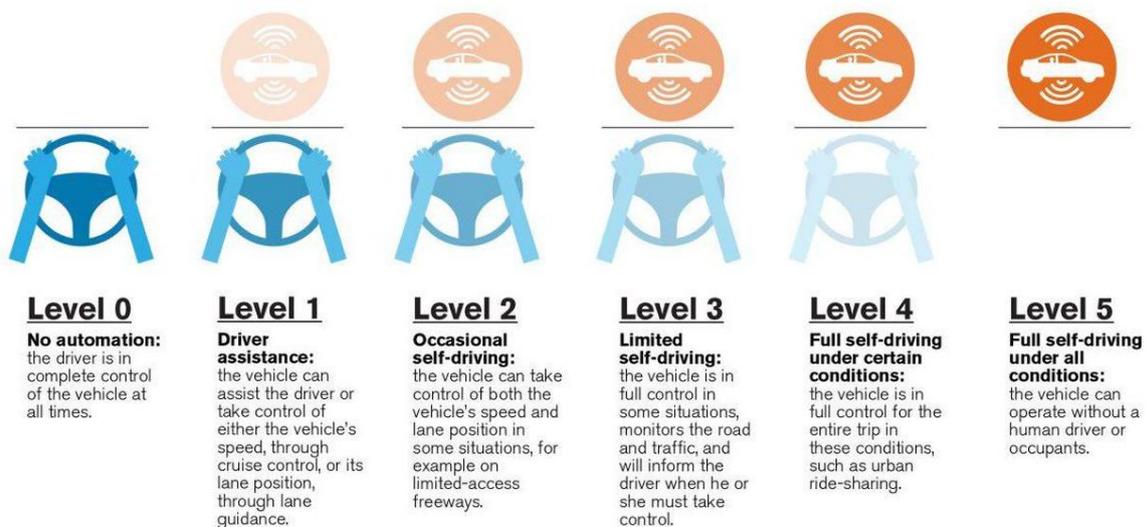# VSORA Simplifies ADAS

by Jan Pantzar, VP Sales & Marketing, VSORA

What is ADAS?

ADAS, or Advanced Driver Assistance Systems, are systems intended to help the driver in the driving process, to make it a safe, comfortable and simple experience. At the end it is not just the driver and the passenger safety that this relates to, but the safety for everyone on the road.

There are different levels of autonomy in the driving. We are talking about 6 levels, where the lowest level, L0, means no automation, and where the driver is in charge of all the driving and the vehicle only responds to inputs from the driver, but can provide warnings about the environment. The highest level, L5, means there is no human driver (so the steering wheel is optional) and everyone in the vehicle is a passenger. The vehicle is in charge of all the driving and can operate in all environments without any need for human intervention. This type of autonomy requires advanced sensor fusion and significant processing power as well as built-in redundancies if a component should fail.

Level 0
No automation: the driver is in complete control of the vehicle at all times.

Level 1
Driver assistance: the vehicle can assist the driver or take control of either the vehicle's speed, through cruise control, or its lane position, through lane guidance.

Level 2
Occasional self-driving: the vehicle can take control of both the vehicle's speed and lane position in some situations, for example on limited-access freeways.

Level 3
Limited self-driving: the vehicle is in full control in some situations, monitors the road and traffic, and will inform the driver when he or she must take control.

Level 4
Full self-driving under certain conditions: the vehicle is in full control for the entire trip in these conditions, such as urban ride-sharing.

Level 5
Full self-driving under all conditions: the vehicle can operate without a human driver or occupants.

Source: Forbes, SAE & NHTSA

The world today is seeing the first L3 commercial vehicles, for example in the form of the 2019 Audi A8L. L3 is classified as conditional automation, where the vehicle can take full control over steering, acceleration and braking, such as accelerating past a slow-moving vehicle. However, the driver must be ready to take control should the self-driving systems be unable to complete its task.

In the next level, L4, there is a high level of automation, where in most cases there is no need for human interaction, even though the human still can manually override. Legislation and infrastructure are not yet in place for this to be generally available, although in some places within a limited area (geofencing) it is possible. For example, Waymo has a L4 taxi service in Arizona and Navya is selling electric L4 shuttles and cabs in the US.

ADAS Adoption

Euro NCAP (the European New Car Assessment Programme) has embraced ADAS and continuously adjust and adapt its assessment procedures to handle the increasing number of systems and technologies. Consumer awareness for occupant and pedestrian protection has become widespread through the championing of Euro NCAP over the past decade or more. It is expected that Euro NCAP will play a similar role for the ADAS promotion and adoption.

Already today the Euro NCAP has a "safety assist" test, where they assess forward collision warning/autonomous emergency braking systems, as well as lane departure warning systems, lane-keep assist systems and speed assist systems. These are all included in the ratings for new car assessments.

There are already ADAS functions that are mandatory on new vehicles. For example, ABS has been mandatory on all new vehicles since 2004 in the EU; electronic stability control since 2014. Autonomous emergency braking has been mandatory on all commercial vehicles in the EU since 2015, and rear-facing or reversing cameras with a dashboard display screen are mandatory in the US since May 2018.

In May 2018 the European Commission published a list of 11 new safety features they would like to have mandated on new cars by 2021. Some of these are existing ADAS functions – autonomous

emergency braking, lane keeping assistance, intelligent speed assistance and reversing camera/rear detection system.
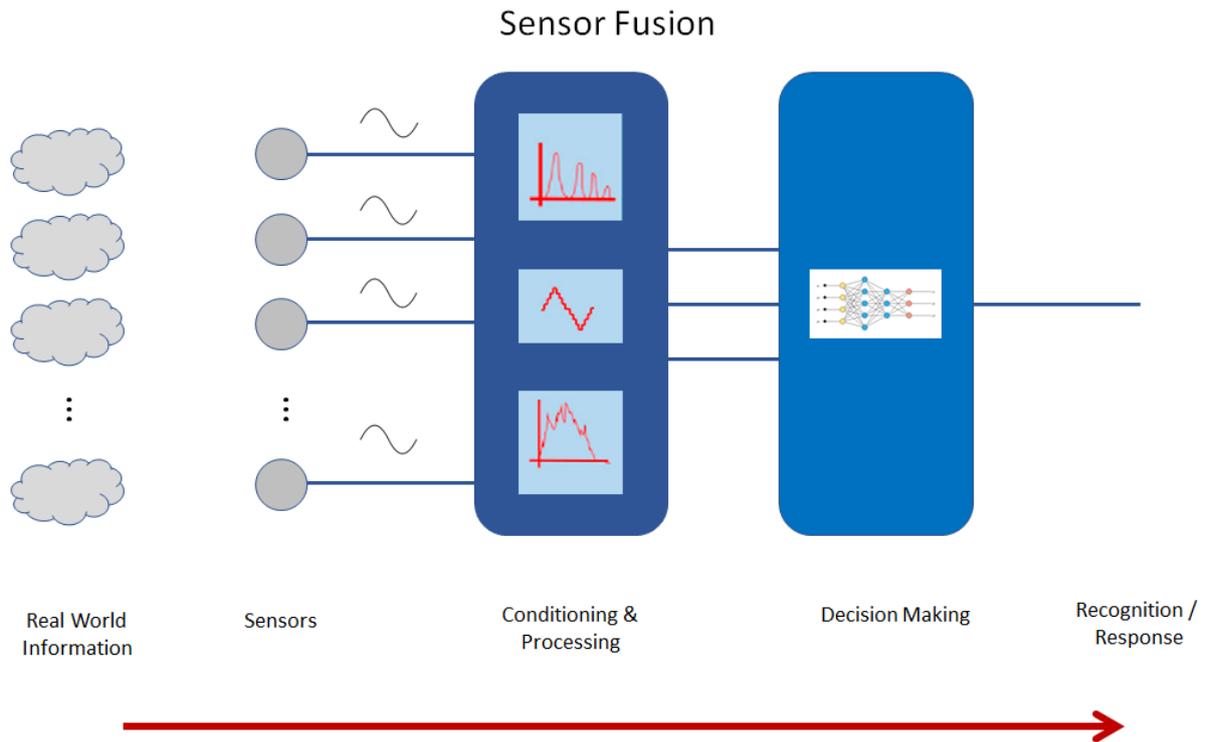
ADAS Today

To achieve a reasonable functionality ADAS require the combination of multiple technologies. The general principle is that one or more sensors will monitor something, and then the provided result needs to be processed and analysed, after which a decision on action can be taken, for example, braking or steering. Sometimes the information requires a combination of inputs from various sensors, so called sensor fusion, for it to function in a reliable way.
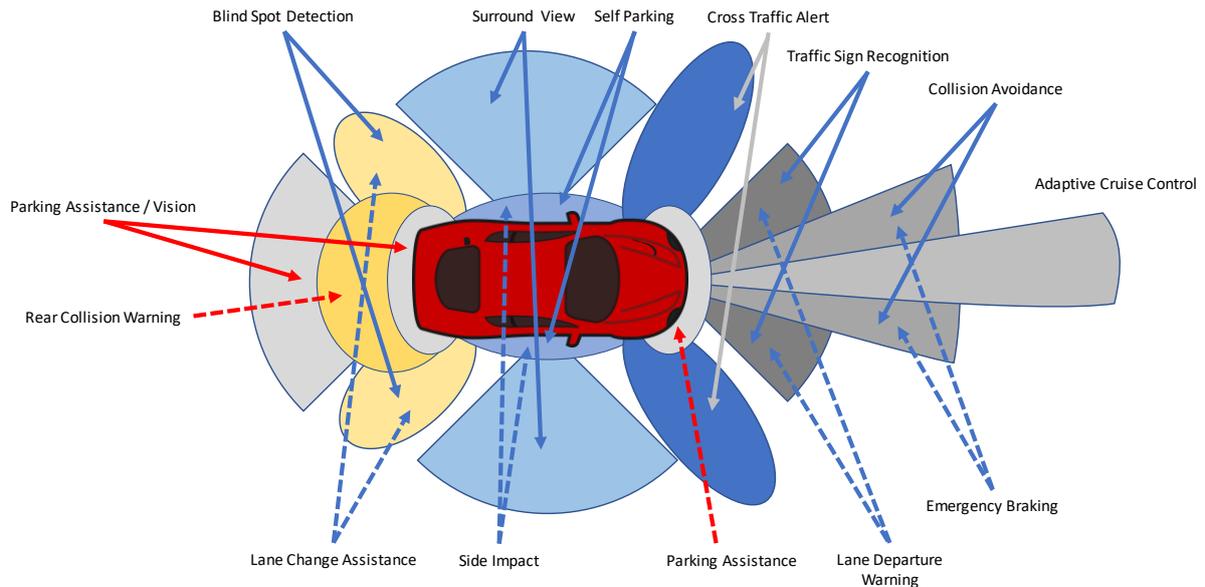
Typically, an action today would be a driver alert. This alert can be the beeping from the parking sensor or a blinking light on the dashboard. The forward collision warning usually provides audio (escalating warning beeps), dashboard warning and a steering wheel alert (vibration).

ADAS Sensors

Sensors are the foundation for ADAS, and they provide a continuous stream of information about the environment surrounding the vehicle. The sensor must not only detect what the driver can see, but also what cannot be seen or hasn't been noticed. Currently four types of ADAS sensors are being used, each with its benefits and each with its drawbacks. As a result, many of the ADAS functions rely on a combination of information from multiple, different sensors where sensor fusion is used to generate an actionable output.

Sensor Fusion

| Real World Information | Sensors | Conditioning & Processing | Decision Making | Recognition / Response |

The four sensor types used in ADAS currently are camera, radar, lidar and ultrasound.

Ultrasound

The ultrasonic sensor is the oldest sensor, and it uses reflected sound waves to calculate distance to objects. These have a comparatively short effective operating range, which means they are only really useful in low-speed systems. Ultrasonic sensors are used in parking sensors, and also in park assist, self-parking as well as blind-spot monitoring. They are not affected by poor light conditions (at night; bright, low sunlight) or by bad weather.

Some manufacturers have decided to forego ultrasonic sensors and have replaced them with short-range radar, which makes it possible to use the sensor for rear cross-traffic/pedestrian alert systems as well.

Radar

Radar is well established and is based on detecting objects by measuring the time it takes for a transmitted radio wave to reflect back from any object in their path. For ADAS applications radar is divided into three categories: long-range radar (LRR), mid-range radar (MRR) and short-range radar (SRR). Short- and mid-range radar systems are typically using microwaves in the 24GHz spectrum, although there has been a shift to 77GHz for bandwidth and regulatory reasons.

SRR has a useful range of 10-30 meters, which makes them useful in blind-spot detection, lane-change assist, park assist and cross-traffic monitoring.

MRR operates between 30 and 80 meters, and LRRs have a range of up to 200 meters. LRR is used in adaptive cruise controls, forward collision warning and automatic emergency braking. LRR suffers from having its measurement angle decrease with range, so in some cases, like adaptive cruise control, inputs from both SRR/MRR and LRR sensors are used.

Radar can be used to detect speed.

Like ultrasonic sensor, radar sensors are not weather and light dependent. They work at night and in rain, snow and fog. The drawback with radar is that is will detect an object, but cannot identify the object, requiring additional sensors to help with this. The 24GHz operation has further limitations since it is unable to differentiate between multiple targets. Also, because of the limited field of view in automotive applications, multiple sensors are required in order to get proper coverage.

Lidar

Lidar is similar to radar but uses light waves from lasers instead of radio waves. It calculates the time it takes for the light to hit an object or surface and reflect back to the scanner. Using lidar makes it possible to calculate a reflected point's position in 3D with high accuracy, and the collected data points are represented as 3D point clouds. The 3D position ability also enables the ability to establish movement direction, i.e. an object moving towards the sensor; away from the sensor; across etc.
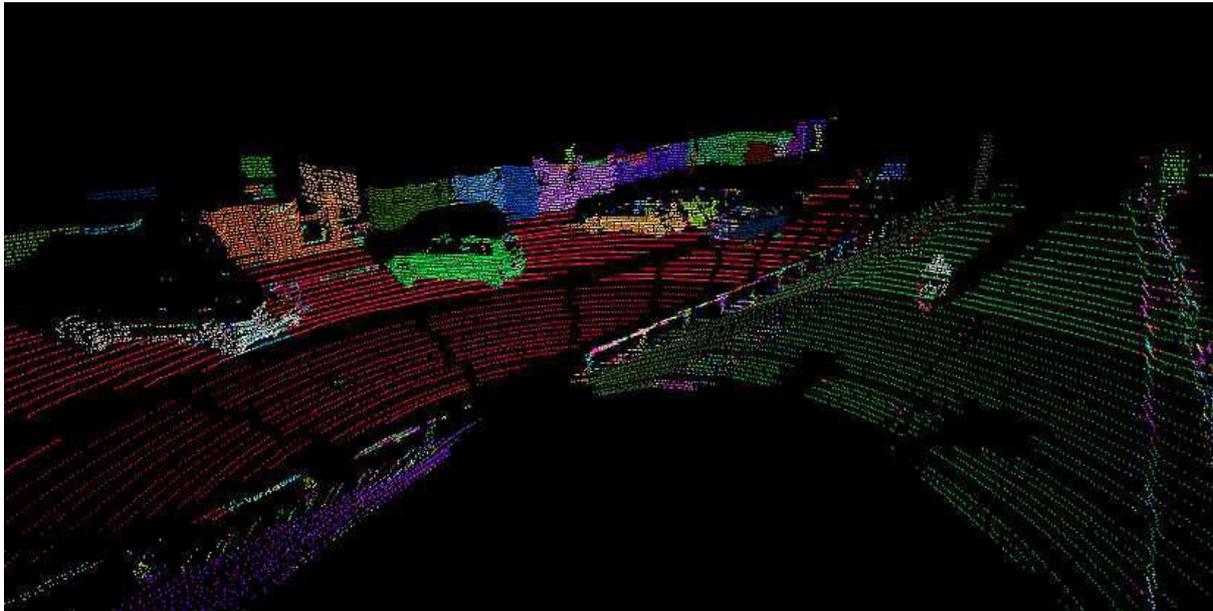
Source: LiDAR for Autonomous Vehicles, Yu Huang

The challenge lies in converting the point clouds into real-time 3D graphics and to detect objects. Currently there are two types of lidar being used – electro-mechanical lidar and solid state lidar. The electro-mechanical lidar is based on a spinning sensor assembly typically mounted on top of a vehicle. By continuously rotating it scans the vehicle surroundings and has long range coverage. The solid state lidar has all components – emitter, receiver and processor – typically integrated on a single chip, making them compact. Since they have no moving parts several units need to be used in order to achieve 360° coverage. Even though lidar can work in poor light conditions, in some cases it will be affected by bad weather, like rain for example.

Cameras

Cameras are widely used in ADAS applications, because of their ability to identify colour and contrast information. These are key for detecting road signs and road markings. They normally also have the resolution to classify objects, for example, pedestrians and bikers. Forward-facing cameras are used in lane-keeping assistance, cross-traffic alert in addition to traffic sign recognition. Rear-facing cameras are widely used for parking assistance. More recently stereo cameras have been used for active cruise control and forward collision warning applications. The stereo camera enables 3D image information that can be used to calculate distances to moving objects, for example.

Cameras are cost-effective but suffer from performance limitations in bad weather and under poor light conditions. Another issue with cameras can be determining contrasts / edges for example, when the object and the background have the same colour.

Thermal cameras can be used to detect humans and animals at ranges up to 300 meters, without being affected by fog, dust, sun glare or darkness. These first started appearing as passive night vision assist systems a decade ago.

Sensor Fusion

To benefit from the advantages of different sensor types and to overcome the disadvantages sensor fusion is typically used. For example, cameras can be combined with radar and/or lidar to ensure the data stream is unaffected by a wider variety of conditions.

The real challenge will then be how much compute speed will be required to handle the data throughput. An interesting example was brought forward by OSRAM at the 2018 Electronica in Munich.
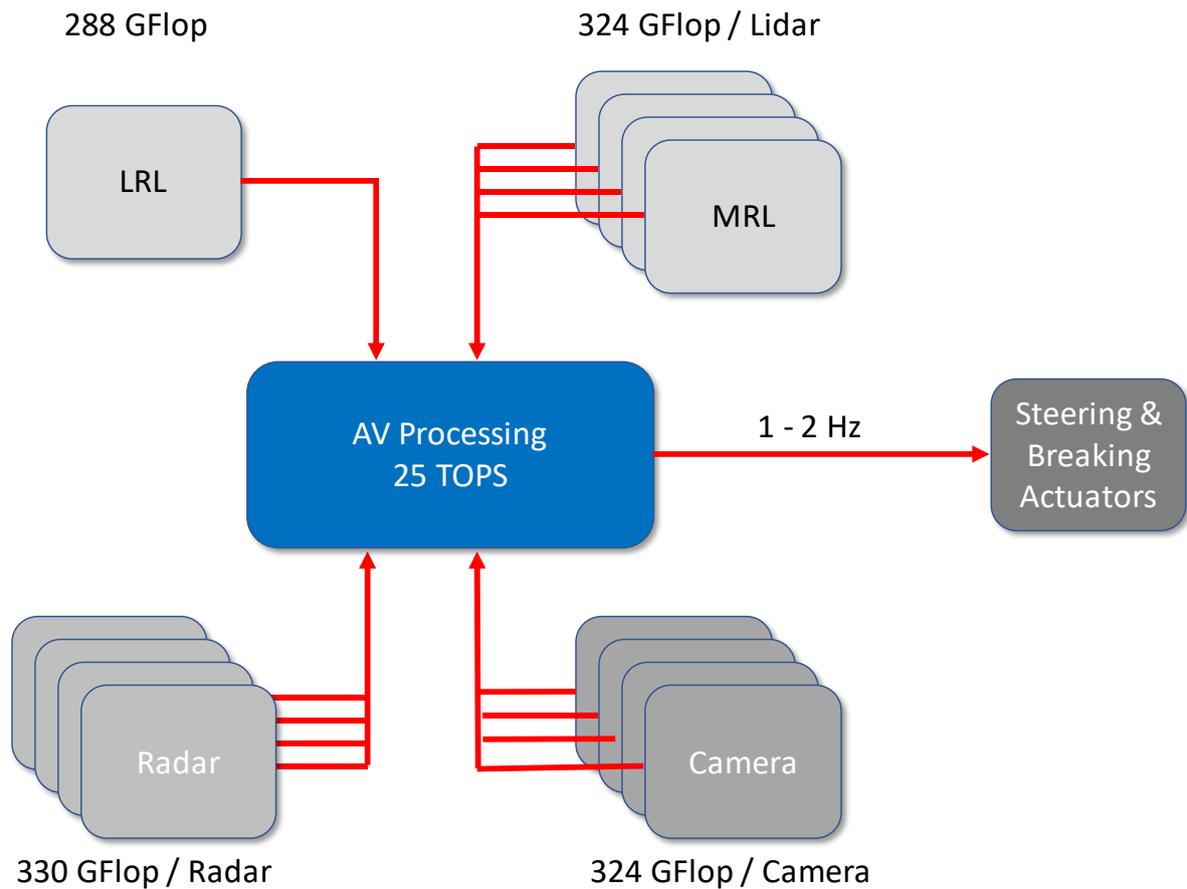
OSRAM made an assumption that you would have 1 long range lidar (LRL), 4 medium range lidars (MRL), 4 radars and 4 cameras. The inputs from these would then be processed and the output would be fed to steering and braking actuators. The steering and braking actuators has a natural frequency of 1-2 Hz for most cars, since anything faster will affect the comfort.

The LRL was assumed to have an azimuth of 80^o and a range of 200 meters. The MRL had an azimuth of 180^o and a range of 100 meters. In both cases the elevation was 60^o, the angular resolution 0.1^o, the linear resolution 1 cm and the frame rate 30 Hz.

A few quick calculations provide the data throughput required for the LRL cloud points of data for the LRL: 288 GFlops. Calculated this way:

- Number of points per frame (ppf) = (azimuth / angular resolution) x (elevation / angular resolution); -- 80/0.1x60/0.1 = 480,000
- Number of frames in 1 data set (fpds) = range/linear resolution; -- 200/0.01 = 20,000
- Number of cloud points in 1 data set = ppf x fpds; -- 480,000 x 20,000 = 9.6×10^9
- Number of cloud points per second = (ppf x fpds) / frame rate; -- 9.6×10^9 x 30 = 288×10^9 (288 GFlops)

Similarly, the MRL data throughput is calculated to 324 GFlops. A combination of 1 LRL and 4 MRL will have a data throughput requirement of 288 + 4 x 324 = 1,584 GFlops. Assuming that the radar and camera needs are similar the data throughput will be 3 x 1,584 = 4,752 GFlops or approximately 5 TOPS for a single pass of data. At least 5 operations are required before objects are identified and referenced into the environment: 5 x 5 = 25 TOPS will be required under these conditions.

288 GFlop

324 GFlop / Lidar

LRL

MRL

AV Processing
25 TOPS

1 - 2 Hz

Steering &
Breaking
Actuators

Radar

Camera

330 GFlop / Radar

324 GFlop / Camera

VSORA Solutions

VSORA has developed a unique, algorithm-agnostic architecture that allows you to optimize power, performance and silicon size. With a single core scalable from 256 to 65,536 MACs and the ability to design a system with multiple cores, the solution gives the user a wide range of choices. An innovative combination of software and hardware can reconfigure a system in a single clock cycle, to accelerate all current and future algorithms without hardware changes.

For the specific example provided by OSRAM above, two cores using 4,096 MACs each would provide 25 TOPS computing power. There is plenty of room to scale this up, as a single core of 65,536 MACs can provide up to 290 TOPS, and the ability to run multiple cores in parallel makes the selection process very easy and extremely flexible.