


Representation of lidar sensing



Planning the route for the uncrashable car

■ New sensor and processing technologies will lead to safer roads, says Chet Babla

Road traffic accidents kill 1.3 million people every year, with countless others injured. Many of these accidents involve human error and are potentially avoidable. While the ultimate goal of the ‘uncrashable’ car may still be some way off, the level of vehicle automation is increasing rapidly and paving the way to fully autonomous vehicles.

Choosing and fusing the optimum combinations of sensor technologies and then processing the huge amounts of data quickly and reliably will be pivotal to the successful development of safe vehicles.

Automation will be driven in part by increased safety regulations and changes to laws, but will only truly be possible by developing and implementing new sensor and processing technologies.

The ADAS spectrum

Vehicle automation spans a spectrum of functional capabilities, from advanced driver assistance systems (ADAS), which include lane-keep and lane-change assist features and adaptive cruise control, up to fully autonomous vehicles. Full autonomy may be several years away technically, politically and in terms of consumer acceptance, but new and emerging sensor technologies enable the progression through the levels of conditional ADAS automation. Fundamental to this will be the

integration of different sensor modalities, each offering different benefits.

Compared to vision and radar systems, lidar (light detecting and ranging) technology has superior resolution, range and depth perceptions, making it an important enabling technology for driver automation. In particular, new coherent detection solutions are delivering important benefits compared with established time of flight (ToF) techniques.

ADAS modalities

ADAS applications use various sensor modalities to gather data about a vehicle’s exterior and interior environment. Due to the dynamic nature of a car’s surroundings, this data must be processed rapidly to enable timely decision-making and safety responses such as braking, steering, or driver alerts. Modalities such as ultrasonic, radar, lidar and vision are fundamental to a vehicle’s sensing capability, with each bringing benefits to maximise the environmental awareness of the vehicle (*Figure 1*).

Ultrasonic sensors provide a practical, robust, inexpensive, short-range solution for low spatial applications such as park-assist. Initially used alone, they are now often combined with parking cameras.

Radar has been common in high-end vehicles for many years, providing

collision detection warning capabilities. Depending on the operating frequency selected from the 20-120GHz radio spectrum, it can perceive objects and relative speeds over a range spanning below 1m – ~200 metres. More recently, radar has been coupled with active safety features such as automatic emergency braking (AEB) and combined safety/convenience applications such as adaptive cruise control (ACC), lane-change assist and blindspot detection. Operating at 120GHz, this modality also delivers the performance and resolution needed for in-cabin sensing, supporting requirements for driver and occupant monitoring systems, including vital sign and respiratory detection.

Vision sensors operate in either visible light or infrared spectrums. They offer excellent spatial resolution which, when enhanced by perception-enabled algorithms and combined with colour, contrast and character recognition capabilities, provide support scene interpretation, object classification and driver visualisation. Vision sensing can support surround view, park-assist, digital mirrors, AEB, lane centring, and ACC. Emerging systems for driver and occupant monitoring are also supported.

Lidar operates in the near-infrared (NIR) spectrum and offers similar capabilities to radar. However,

at the upper end of the spectrum (approximately 1,550nm) lidar achieves spatial resolutions that are up to 2,500 times finer than radar. This allows the technology to render accurate 3D depictions of a vehicle’s surroundings.

Lidar uses a laser source to illuminate a scene with light, which is scanned in the horizontal and vertical axes. The resulting reflections are detected by receiving photodetectors creating the 3D image or ‘point cloud’.

Automotive lidar was initially earmarked for fully autonomous vehicles. Due to its resolution, range, and depth perception capabilities, however, it has become immediately relevant to ADAS in mass market vehicles.

The rise of coherent lidar

The two approaches to lidar are incoherent and coherent systems. Incoherent, also known as direct lidar, typically uses ToF, whereas the emerging coherent systems employ a frequency modulated continuous wave approach to detection.

Coherent detection measures the time taken for emitted light pulses to return to the photodetector. The distance to the object is then calculated using that measurement. With coherent detection, a source signal is modulated onto an emitted coherent laser wave, then

detected at the receiver by comparing the phase and frequency of the received signal with the original acting as a reference. Measuring the Doppler shift of received light achieves rapid and accurate positioning, classification and path prediction of objects over long ranges.

ToF-based lidar systems operating in NIR use laser sources and silicon-based detectors that are readily available from existing datacoms and consumer applications. Coherent lidar systems operate at short wave infrared (SWIR) wavelengths of 1,300nm-1,600nm, enabling longer range detection. An additional benefit is that the components used are available from current telecoms applications.

Coherent detection has several performance benefits over incoherent ToF that have led to it gaining traction in automotive applications. These include support for per-point velocity and better immunity to interference from other light sources including sunlight and other lidars around the vehicle. An inherent sensitivity that is several orders of magnitude higher gives a longer effective range, or, depending on the application needs, the option to lower the optical transmit power for a required performance level. Another benefit is the better eye safety of SWIR wavelengths compared to NIR, meaning higher optical power can be used safely.

From a design and manufacturability perspective, the ease of integration of coherent lidar components into silicon photonic ICs allows smaller form factors. At the same time leveraging the economies of scale that silicon brings enables mass market deployment.

Scalable processing

To date, addressing the complex and compute-intensive demands of coherent lidar systems has been achieved using high-end FPGAs. These are versatile but have a high price point and are power hungry, thermally demanding devices. These characteristics make them unsuited to mass market automotive applications.

Though coherent detection does require more complex processing, innovative digital signal processing techniques can address this and meet the commercial and technical needs of increasingly complex and high performance ADAS. One way to achieve this is with an asic, which if designed from the outset for the end




Figure 1: Fusing the benefits of different sensor modalities provides a comprehensive view of a vehicle and its surroundings

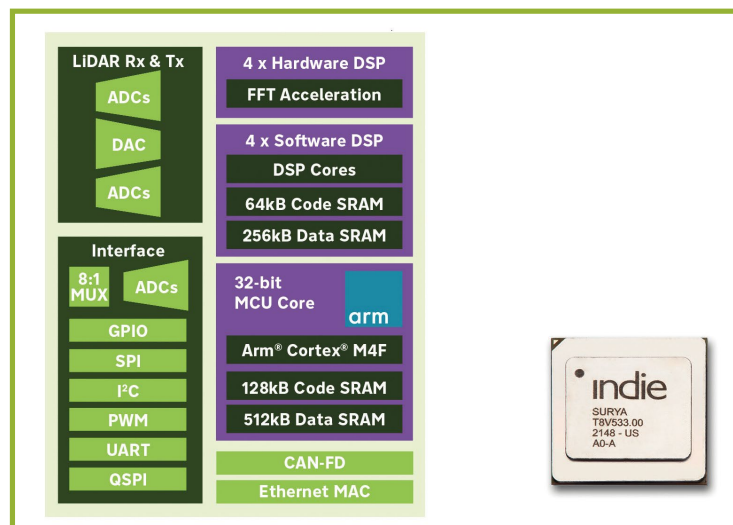


Figure 2: Indie Semiconductor's iND83301, Surya lidar SoC is claimed to be the world's first merchant, coherent detection lidar asic

application, has the potential to deliver major power efficiencies, a simplified external bill of materials and reduced system form factor at a far lower cost than an FPGA-based solution. Although complex asics typically require several years of development and major upfront investment, adopting them is compatible with automotive volumes that will grow further as automakers push toward the goal of the uncrashable car.

Multiple sensor modalities

The exponential increase of deployed sensors for ADAS is leading to a corresponding increase in the volume

of data to be transported and reliably processed. This places growing demands on the computational power of the vehicle.

Automotive engineers can address this challenge using either centralised or distributed approaches. Centralised processing architectures have a more straightforward software architecture, a lower component count and can help reduce processing latency. They are, however, power hungry, which is a particular challenge as powertrains are increasingly electrified, necessitating large form factors and the need for complex thermal management. These

disadvantages are compounded by the high bandwidth necessary to accommodate the flow of raw sensor data to a central point. Furthermore, a centralised compute architecture may not be suited to mass market vehicle classes where the technical and commercial burdens of 'overprovisioned' compute may be practically prohibitive.

Distributed computing architectures have some degree of local pre-processing integrated in or near the sensors in the case of a local sensor fusion processor and a zonal controller for data aggregation, processing or fusion with other sensor data. This eases the central processing compute requirements and reduces the need to support multi-Gbit communications across the vehicle. It also allows scalability of deployment across an OEM's vehicle classes. The trade off is that software management of distributed nodes and 'sense-think-act latency' needs to be carefully considered and resolved, but this challenge is being addressed today. □

About the author

Chet Babla is senior vice-president of strategic marketing, Indie Semiconductor