

Top Level Newsletter: Connected Vehicle
(Published by IEEE Orange County Section)

August/September 2023
Volume 43.0

Editor-in-Chief: Kay Das, IEEE Life Member,
Systems Research Development and Innovation

Vol 43: This issue includes system features and competitive new developments

1. Autonomous testing calls for massive computing power
2. Lidar on a Chip Enters the Fast Lane: Sensors for Self-Driving Cars

General Notes

This series of newsletters is intended to provide the IEEE member with a top level briefing of the many different subjects relevant to the research, development and innovation of the connected vehicle.

The objective is to provide a platform for fast learning and quick overview so that the reader may be guided to the next levels of detail and gain insight into correlations between the entries to enable growth of the technology. Intended audiences are those that desire a quick introduction to the subject and who may wish to take it further and deepen their knowledge. This includes those in industry, academia or government and the public at large. Descriptions will include a range of flavors from technical detail to broad industry and administrative issues. A (soft) limit of 300 to 600 words is usually set for each entry, but not rigorously exercised.

As descriptions are not exhaustive, hyperlinks are occasionally provided to give the reader a first means of delving into the next level of detail. The reader is encouraged to develop a first level understanding of the topic in view. The emphasis is on brief, clear and contained text. There will be no diagrams in order to keep the publication concise and podcast-friendly. Related topics in the case of Connected Vehicle technology, such as 5G cellular and the Internet of Things will be included. The terms Connected Vehicle and Automated Driving will be used inter-changeably. Articles from other published sources than IEEE that add to the information value will occasionally be included.

This newsletter forms part of the regional Advanced Technology Initiative (ATI) of which connected vehicles form a constituent part. Technical articles solely from IEEE journals/magazines are referred to by their Digital Object Identifier (DOI) or corresponding https link. The link for each article is provided. Those readers who wish to delve further to the complete paper and have access to IEEE Explore (www.ieeeexplore.ieee.org) may download complete articles of interest. Those who subscribe to the relevant IEEE society and receive the journal may already have physical or electronic copies. In case of difficulty please contact the editor at kaydas@mac.com. The objective is to provide top level guidance on the subject of interest. As this is a collection of summaries of already published articles and serves to further widen audiences for the benefit of each publication, no copyright issues are foreseen.

Readers are encouraged to develop their own onward sources of information, discover and draw inferences, join the dots, and further develop the technology. Entries in the newsletter are normally either editorials or summaries or abstracts of articles. Where a deepening of knowledge is desired, reading the full article is recommended.

1 Autonomous testing calls for massive computing power

Automotive Engineering September 2023, JOHN KENDALL

Company goes to high-performance frontier by building liquid-cooled processors for ADAS testing.

As humans, we are quite good at driving, which, curiously, is one of the issues facing the development of autonomous vehicles (AVs). Another is the development of testing equipment that can handle the massive amounts of data needed to develop the failsafe systems that will guide AVs. This can be measured in teraFLOPS – floating point operations per second – a massive amount of data, which needs to be gathered and stored in the hostile environment of a development vehicle.

The in-vehicle devices need to have the right bandwidth. As Giancarlo Cutrignelli, head of product management at Edge Computing specialist Eurotech, told SAE Media, these devices need to sustain all the operations, storage and processing at the feed speed. “You need to have all of them packed into the car. They need to fit in the trunk, essentially. Then you have power constraints because we also need to supply them from the car, but there’s a caveat.

“All the OEMs we are talking to at the moment are designing their systems on electric cars, not ICE. Electric cars means the autonomy needs to be managed and the maximum power needs to be managed and there are some systems in electric cars which do not allow you to suck energy out if they are needed for safety systems.”

In addition to performance and what Cutrignelli calls the power budget, these devices will need to survive seven or eight or more hours of driving per day, every day.

“You want to have high performance, ruggedization and small form factor altogether to basically operate,” says Cutrignelli, “When you are storing and logging so much data, even the time to move this data away from the vehicle and back to the data center is relevant because there’s a latency in starting a new log-in cycle. So even transferring this enormous amount of data quickly is an issue.”

Eurotech’s solution to these problems is multifaceted. The company acquired Munich-based InoNet Computer last fall. InoNet Computer specializes in configurable systems in the medical, industrial and automotive sectors. Eurotech has historically supplied two other technologies: liquid cooling and ruggedization. “We’ve been using industrial-grade liquid cooling,” says Cutrignelli, “We have expertise in how to ruggedize systems and put them in a small form factor. This is summarized in a portfolio of products which have a varying level of ruggedization, a varying level of logging capabilities available in different form factors. So, it’s not a one-size-fits-all type of approach.”

Eurotech tends to use two kinds of liquid cooling. The first uses a closed circuit. “That’s what we use, for example in Dynacore 6110 and 6210 series of products, where you have a hybrid approach. So, the most thermally critical elements are cooled down with the liquid.”

Where more cooling is required, Eurotech would use an external heat exchanger supplied with the machine, giving it complete independence. “It’s like an external fridge that’s cooling down and feeding the system,” continues Cutrignelli, “If form factor is more of an impact and if you can rely on connecting the liquid cooling system to the air-conditioning system of the car, then you can use that as a heat exchanger. In that case, we can offer solutions which are way more compact.”

Eurotech will pack a high-performance edge computer. “The challenge is really making sure that the system is well-dimensioned,” says Cutrignelli, “This is harder than people may think. You can be a data scientist and just look at the computational and the modelling part of it, or you can be more of a mechanical engineer, looking at the capability of surviving hostile conditions, with requirements on temperature ranges and so on. It's a journey together with the customer because the requirements are changing all the time.” (622 words,

2. Lidar on a Chip Enters the Fast Lane: Sensors for Self-Driving Cars and Robots will be Tiny, Reliable, and Affordable , by (Michael R. Watts et al)

Published in: [IEEE Spectrum](#), (Volume: 60, [Issue: 9](#), September 2023)
Page(s): 38-43.

Abstract:

A switch to self-driving cars and trucks with various types of electronic sensors and sophisticated computers at the helm could save countless lives. But getting this promising technology into people's hands has been difficult, despite massive research investments and considerable technical progress. So when will self-driving cars really come to a driveway near you? The answer depends in part on whether such cars require a type of sensor called lidar, short for “light detection and ranging.” Most groups developing autonomous vehicles see lidar as a critical part of the sensor suite required for safe operation, because it allows a detailed 3D map of the vehicle's environment to be constructed with much more fidelity than can be done with cameras.

Elon Musk, though, has been pushing Tesla to adopt a controversial cameras-only approach to autonomous driving. “Humans drive with eyes & biological neural nets, so makes sense that cameras & silicon neural nets are only way to achieve generalized solution to self-driving,” Musk tweeted in 2021. The mechanical complexity and high cost of most lidar sensors—which not long ago would have added tens of thousands of dollars to the price of each vehicle—no doubt helped shape Musk's views. As early as 2016, he declared that “all Tesla vehicles exiting the factory have hardware necessary for Level 5 autonomy” —meaning that cars with cameras and computers alone have what's needed for fully autonomous driving.

Seven years and many crashes later, Tesla has not progressed past Level 2 autonomy, and traffic-safety specialists are questioning Musk's rejection of lidar. Requiring pricey sensors, though, would slow the widespread rollout of both advanced driver-assistance systems and fully autonomous driving. But reducing the cost of these sensors to a level that would satisfy automakers has remained an elusive goal for lidar manufacturers, which must also consider how to add their devices to cars without detracting from vehicle aesthetics.

We and others at our company, Analog Photonics, which spun out of MIT in 2016, hope to break this impasse. We are developing a tiny, chip-scale phased-array lidar that promises to slash costs and simplify integration. Here we'd like to explain some of the technical challenges we've encountered and how very close we are to commercialization.

Today, more than half of new cars are equipped with one or more radar sensors. These sensors are solid state, cost manufacturers less than US \$100 each, and are small enough to be inconspicuously placed around the vehicle. They are used for a variety of things, including automatic emergency

braking and adaptive cruise control, as well as lane keeping and other advanced driver-assistance functions.

But this wasn't always the case. Early automotive radars were large, mechanically steered, emitted short pulses of radio waves, and had limited performance. But the move to electronic scanning and continuous-wave emissions in automotive radars brought performance advancements and cost reductions, which in turn ushered in their widespread use.

Lidar is now undergoing this same evolution. The technology began making headlines around 2016 as a slew of companies, spurred on by the success of lidar sensors on vehicles entered in the DARPA Grand Challenge a decade earlier, began developing custom systems for autonomous vehicles. These systems tended to be pieced together from off-the-shelf components.

These first-generation lidars went only so far. Spinning or scanning mirrors contributed to their high costs and made their integration into vehicles difficult. They also suffered from reliability issues, and their pulsed operation led to problems in the presence of direct sunlight and resulted in an inherent susceptibility to interference from neighboring lidars. As a result, the available lidar sensors have not met the stringent performance, reliability, and cost goals of the automotive industry.

Carmakers are looking for high-performance, long-range lidar sensors that will cost them less than \$500 each. While lidar manufacturers have made progress, the industry isn't there just yet.

Our company chose to attack these problems head-on by designing lidar sensors that are built entirely on a chip—a photonic integrated circuit made of ordinary silicon. It has no moving parts and generates, emits, and receives light with no external hardware. And its tiny size makes it easy to incorporate into the bodies of even the sleekest cars on the road.

Lidar is a lot like radar, but it operates with wavelengths typically between 905 and 1,550 nanometers (compared with a few millimeters for automotive radar). This difference gives lidar much better spatial resolution.

Most **early** automotive lidars, like most early radars, used what is called time-of-flight (ToF) detection. A short pulse of electromagnetic energy is sent out, hits an object, and then reflects back to the sensor, which measures the time it takes for the pulse to complete this round trip. The unit then calculates the range to the object using the known speed of light in air. These systems all suffer from some inherent limitations. In particular, lidars built on this principle are prone to interference from sunlight and from light pulses coming from other lidars.

Most **modern** radars systems work differently. Instead of sending out pulses, they emit radio waves continuously. The frequency of these emissions is not fixed. Instead, they are swept back and forth across a range of frequencies.

To understand the reason for doing that, it's important to know what happens when signals of two different frequencies are combined in a way that isn't purely additive. Doing so will generate two new frequencies: the sum and difference of the two frequencies you initially mixed. This process invented 1901 is called heterodyning,

Frequency-modulated continuous-wave (FMCW) radars take advantage of the fact that signals of two different frequencies, when mixed in this fashion, give rise to a signal whose frequency is the difference of the first two. In these radars, the mixing is done between the outgoing signal (or, in truth, an attenuated version of it, often called the local oscillator) and the reflected signal, which differ in frequency because the outgoing signal is, as we mentioned, being swept across a range of frequencies. So by the time the reflected signal makes it back to the sensor, the outgoing signal will have a different frequency from what it had when the now-reflected waves first left the radar antenna.

If the reflected signal took a long time to make the round trip, the difference in frequencies will be large. If the reflected signal took only a short time to bounce back, the difference in frequencies will be small. So the difference in frequencies between outgoing and reflected signals provides a measure of how far away the target is.

While they are more complex than ToF-based systems, FMCW systems are more sensitive, essentially immune to interference, and can be used to measure the velocity of a target in addition to its distance.

Automotive lidar is now adopting a similar approach. FMCW lidar involves slightly altering the frequency, and thus the wavelength, of the transmitted light and then combining the backscattered light with a local oscillator at the frequency of the transmitted light. By measuring the frequency difference between the received light and the local oscillator, the system can determine the range to target. What's more, any Doppler shifts from a moving target can also be extracted, revealing the target's velocity toward or away from the sensor.

This capability is useful for quickly identifying moving targets and discriminating among closely spaced objects that are moving at different speeds. The velocity measurement can also be used to predict other vehicle movements and can even sense a pedestrian's gestures. This additional dimension to the data, not available from ToF systems, is why FMCW systems are sometimes called 4D lidar.

As you might imagine, FMCW lidar systems use a very different laser source than ToF systems do. FMCW lidars emit light continuously, and that light has comparatively low peak power. The laser power levels are similar to those used in many communications applications, meaning that the light can be generated and processed by *photonic integrated circuits*. This tiny laser system is one of the key factors that has enabled chip-based lidars.

The photonic integrated circuits we designed can be fabricated on standard 300-millimeter-diameter silicon wafers using photolithography, just as is done for most integrated circuits. So we can take advantage of the maturity of the CMOS semiconductor-manufacturing industry to combine all of the various on-chip optical components needed for a full lidar system: lasers, optical amplifiers, waveguides, splitters, modulators, photodetectors, and, in our case, optical phased arrays.

The economies of semiconductor manufacturing slash the cost of each of these components. Having all of them integrated on a single chip helps, too. You see, all lidar systems both transmit light and receive light, and the transmitting and receiving optics must be well aligned. In systems built with discrete optical components, the need for precise alignment adds complexity, manufacturing time, and cost. When things slip out of alignment, the lidar can fail. With integrated photonics, the precise alignment is inherent, because the waveguides carrying the light are lithographically defined.

While a handful of companies are working to develop photonic IC-based lidars, only Analog Photonics has figured out how to eliminate the need to mechanically scan the scene with its single-chip lidar. Instead of mechanical scanning, we use what are called *optical phased arrays*, which allow the beam to be steered electronically.

Scanning is an essential aspect of lidar and one of the key challenges of the technology. Initially, lidar sensors scanned by either spinning the sensor or introducing rotating mirrors into the beam path. The resulting hardware was cumbersome, expensive, and often unreliable.

Although some radars also point their antennas mechanically—as you have no doubt noticed at airports and marinas—some steer the radar beam electronically using phased antenna arrays. This technique adjusts the phase of the signals leaving each of several antennas in such a way that radio waves interfere with one another constructively in one direction and destructively in other directions. By adjusting signal phases at each antenna, the radar can vary the direction in which these signals combine constructively to form a beam.

Electronically phased arrays are the beam-steering technology of choice for automotive radars. Recognizing that the physics of a phased array applies to all frequencies of the electromagnetic spectrum, including optical frequencies, we decided to use this approach in our solid-state lidar. Aided by DARPA through its Modular Optical Aperture Building Blocks program, and with help from several automotive partners (whose names we can't yet reveal), Analog Photonics has developed on-chip optical phased arrays.

The top surface of the chip is used as both a transmitting and receiving aperture—that's where the energy leaves and returns to the chip. The on-chip optical phase shifters and emitters are individually controlled with custom electronics to steer exceedingly tight optical beams, ones that are just several millimeters wide.

Achieving a range of steering that's large enough to be useful requires thousands of closely spaced phase shifters. For example, for a lidar that operates at a wavelength of 1,550 nm, the phase shifters must be placed 1.5 micrometers apart for a 60-degree steering range.

Optical phase shifting requires altering the optical properties of the transparent material inside the chip's many micrometer-scale optical wave-guides, which channel the light from the laser where it is generated to the aperture where it is emitted. If you can change the speed of light in that material, you will alter the phase of the light wave exiting the waveguide.

The material here is just silicon, which is transparent to light at infrared wavelengths. One way to alter the speed of light in silicon is to pass sound waves through it, a technique being pursued for use in lidar by researchers at the University of Washington. Another way is to change the temperature: The hotter the silicon, the more the light passing through it is slowed. This is the principle behind what are called thermo-optic phase shifters.

With thousands of phase shifters on a chip, it's critical that each one consume very little power, mere microwatts. And that's hard to do when you must heat things up. We sidestepped the need for heating by using electro-optic rather than thermo-optic phase shifters. This approach also enabled us to steer

the beam faster, allowing it to step across the field of view at rates exceeding one million scan lines per second.

There remained, though, the challenge of how to connect the many closely spaced optical waveguides with the electronics required to adjust the speed of light within them. We solved this using flip-chip technology: One CMOS chip has thousands of solder-coated copper bumps placed about 75 micrometers apart, or about half the width of a human hair. This scheme allows our silicon photonics chip to be permanently mated with a semiconductor electronic chip containing the needed digital logic and a matching set of copper bumps. Simple commands to the electronic chip then drive thousands of photonic components in the appropriate fashion to sweep the beam.

We have now built and delivered proto-types of the world's first all-solid-state beam-sweeping lidar to its industry partners, which are companies that supply automotive equipment directly to carmakers. We've solved most of the fundamental and engineering challenges and are now focused on increasing the lidar's performance to meet production specifications. We expect to be turning our creations into actual products and producing large numbers of samples for the automotive industry in 2025.

We are currently working on two different versions of our lidar: a long-range version intended to be mounted at the front of the car for use at highway speeds and a short-range version with a wider field of view to provide complete coverage all around the vehicle. The two sensors have different optical phased arrays in their photonic ICs, while sharing the same back-end signal processing.

We expect that relatively low-cost lidar sensors from some of our competitors, will begin showing up in some top-of-the line cars by next year. And driven by the availability of low-cost solid-state sensors like the ones we're working on, lidar will be common in new cars by the end of the decade.

But the future of lidar won't end there. Market forecasters expect lidar to be used for many other applications, including industrial automation and robots, mobile-device applications, precision agriculture, surveying, and gaming. (2009 words)

<https://doi.org/10.1109/MSPEC.2023.10234174>