Springer Vieweg

ATZ extra

SEMICONDUCTORS | OPTOELECTRONICS

Full-Stack Solutions for Autonomous Mobility

Continental

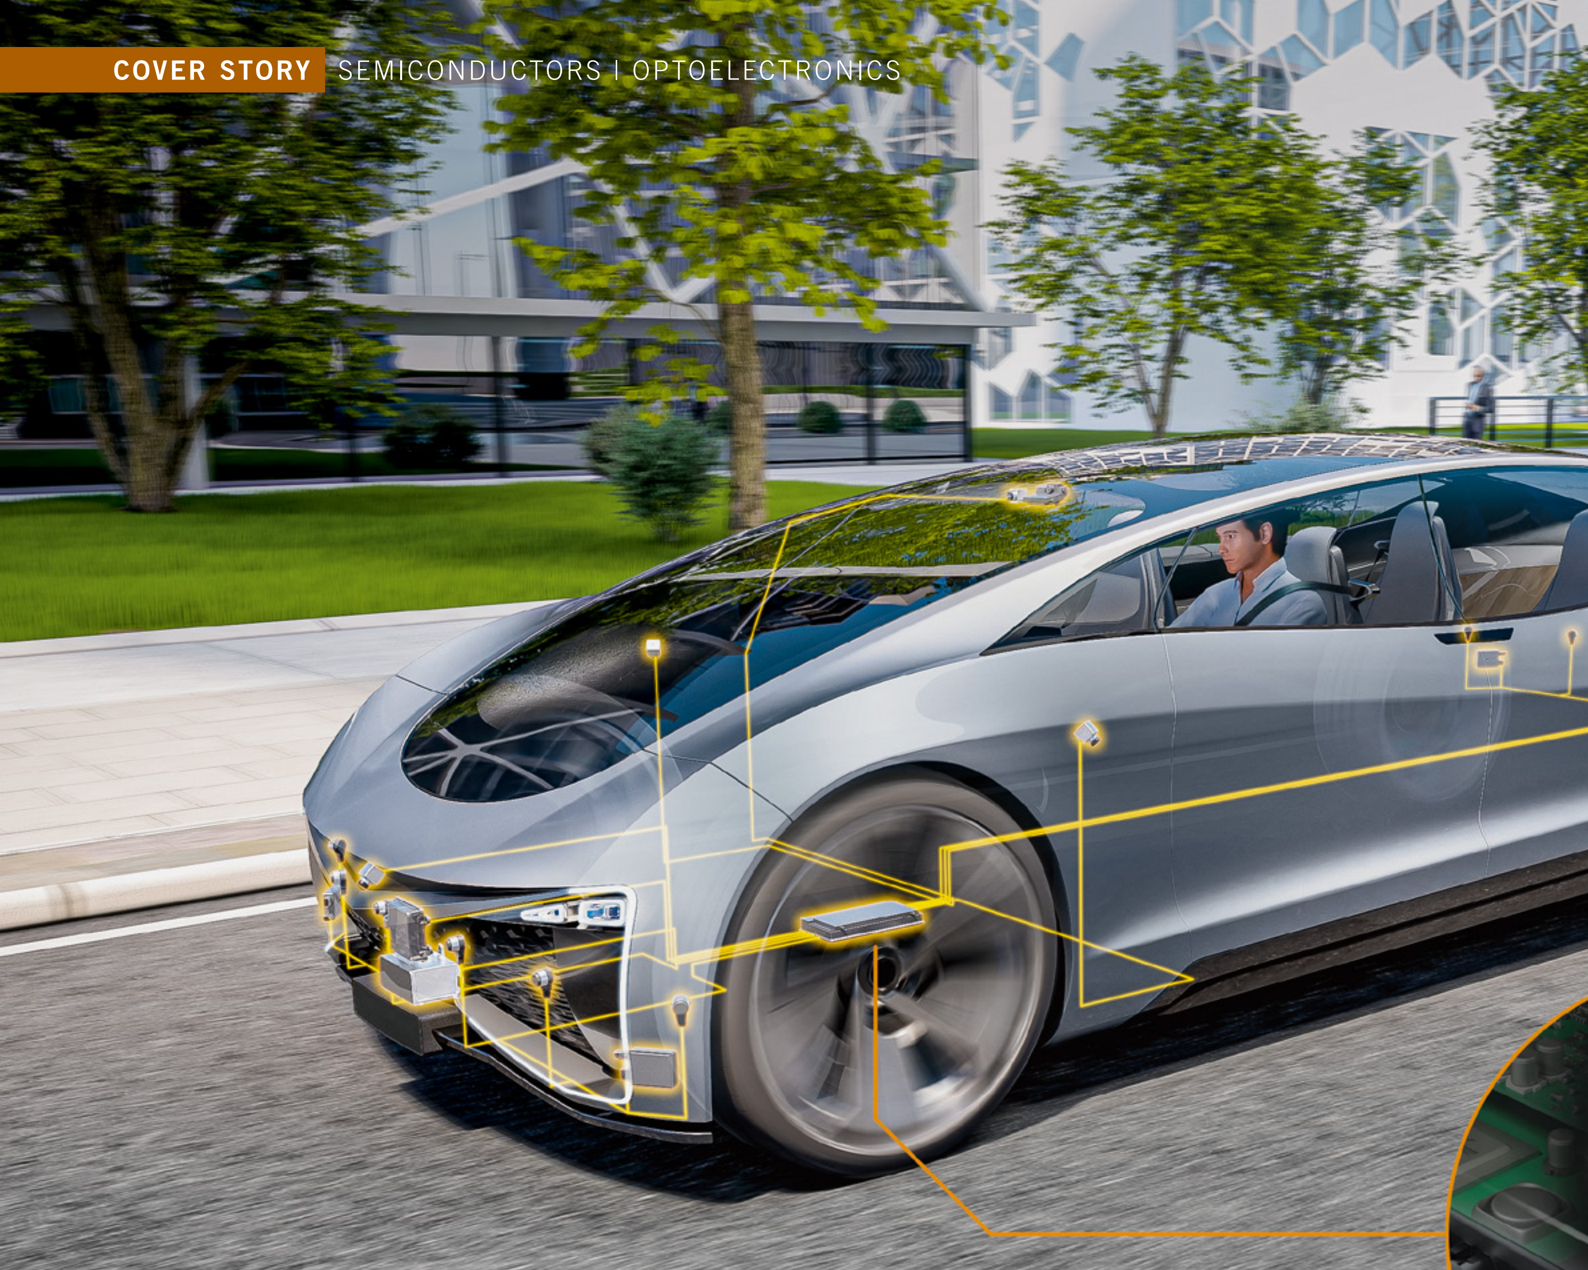

Scalable System Solutions Pave the Way to Autonomous Mobility

The combination of artificial intelligence (AI) and the ever-increasing number of environmental sensors in new vehicles is the basis for assisted and automated driving functions. However, scalable and cost-efficient end-to-end solutions are needed. Thus Continental and Ambarella are developing full-stack solutions based on Ambarella's CV3-AD system-on-chip family of centralized AI domain controllers.

WRITTEN BY

Peter Sättele
is Head of Product
Technology/ADCU
Concepts at Continental
in Ulm (Germany).

Prof. Alberto Broggi
is General Manager of VisLab
Automotive R&D at Ambarella
in Parma (Italy).

Powerful support functions such as adaptive cruise control (ACC), lane change assist and map-based assisted as well as automated driving are critical components for vehicles with autonomous driving capabilities. Global automotive experts have developed a classification system that defines the various levels of automation in a vehicle, ranging from 0 to 5 (L0-L5). The automated market segment will grow strongly in the coming years, as the technical requirements for autonomous solutions are increasingly being met. The number of environmental sensors installed in the vehicle, for instance, is steadily increasing, improving the data resources with which accurate and comprehensive AI modeling of the vehicle's surroundings is made possible.

A prerequisite is hardware and software optimized for this special application with high computing power and at the same time the lowest possible energy consumption. In addition, the market penetration of autonomous driving functions requires a scalable solution that can handle large and diverse sensor suites, and that can be designed for use cases of varying complexity even in case of largely the same sensor mix. The core of such a solution is a central AI domain controller that delivers data processing, deep sensor data fusion and a model of the vehicle's surroundings up to trajectory planning for all autonomous driving functions. The domain controller must have a large number of signal inputs and sufficient computing power to allow real-time or low-latency surroundings modeling. In view of the ever-increasing electrification of the powertrain, it is also essential that this high computing performance is delivered as efficiently as possible.

Continental and Ambarella are developing such a solution, combining sensor hardware with the corresponding software. The hardware core is Ambarella's CV3-AD AI system-on-chip (SoC) central domain controller family. This family includes Ambarella's third generation CVflow vector processor for artificial intelligence (AI), which processes and fuses data from up to 24 data streams from a host of cameras, radars, lidars and ultrasonic sensors. The first application of this joint solution is as a complete L4 fallback system in a commercial vehicle, which is planned for series production in 2027, **FIGURE 1**. Additionally, the companies continue to target 2026 global series production for a broad range of applications for their scalable full-stack systems, from ADAS and Level 2+ up to L4.

FIGURE 1 The first full-stack solution from Continental and Ambarella will go into series production as an L4 fallback system in a commercial vehicle in 2027 (© Continental)

FIGURE 2 Powerful image data processing increases detection performance, even in unfavorable conditions (© Ambarella)

SCOPE OF AN END-TO-END SOLUTION

The typical mix of sensors in a highly equipped vehicle with comprehensive support makes it clear how crucial the computing power of a domain controller actually is. Sensor suites with more than 30 sensors are expected in the future. In addition to a dozen ultrasonic sensors, these will include 10 to 15 cameras and around 10 radars, plus several lidars. The data flow from these sensors is bundled into separate feeds, typically in dedicated interfaces for camera signals and, depending on the architecture, in one (for a zonal architecture) or more ethernet interfaces for the radars and lidar sensors. In the case of camera signals, the usefulness of the data depends heavily on image processing. Only with powerful algorithms can reliable information be obtained from the raw data, including in unfavorable conditions, **FIGURE 2**.

To interpret the surroundings, both the infrastructure for general vector data processing and processors for neural networks (NN) need to be present in the domain controller. The CV3-AD SoC family combines high computing performance with low energy consumption (power and cooling requirements), which is partly due to the small structure width of 5 nm, but even more to a design optimized for the above tasks. For example, the CV3-AD685 SoC's typical energy demand is as low as 30 W. Compared with other SoCs, this can mean an estimated range gain of a minimum of 30 km for an electric vehicle with the same battery capacity. Alternatively, the battery cost can be reduced significantly, and battery mass can be reduced by several kilograms while maintaining the same range.

A full-stack solution for ADAS to autonomous driving requires a wide range of different core competencies. These include in-depth sensor know-how covering the entire sensor mix and, consequently, the essential data resources for environmental perception around the vehicle. Further requirements include specific software expertise in the ADAS domain, such as computer vision and data fusion as well as trajectory planning and control for ADAS applications from NCAP, and from L2+ and L3 through to L4 systems, **FIGURE 3**.

FIGURE 3 The overall solution is integrated into applications ranging from ADAS through to L2+, L3 and L4 systems (© Continental)

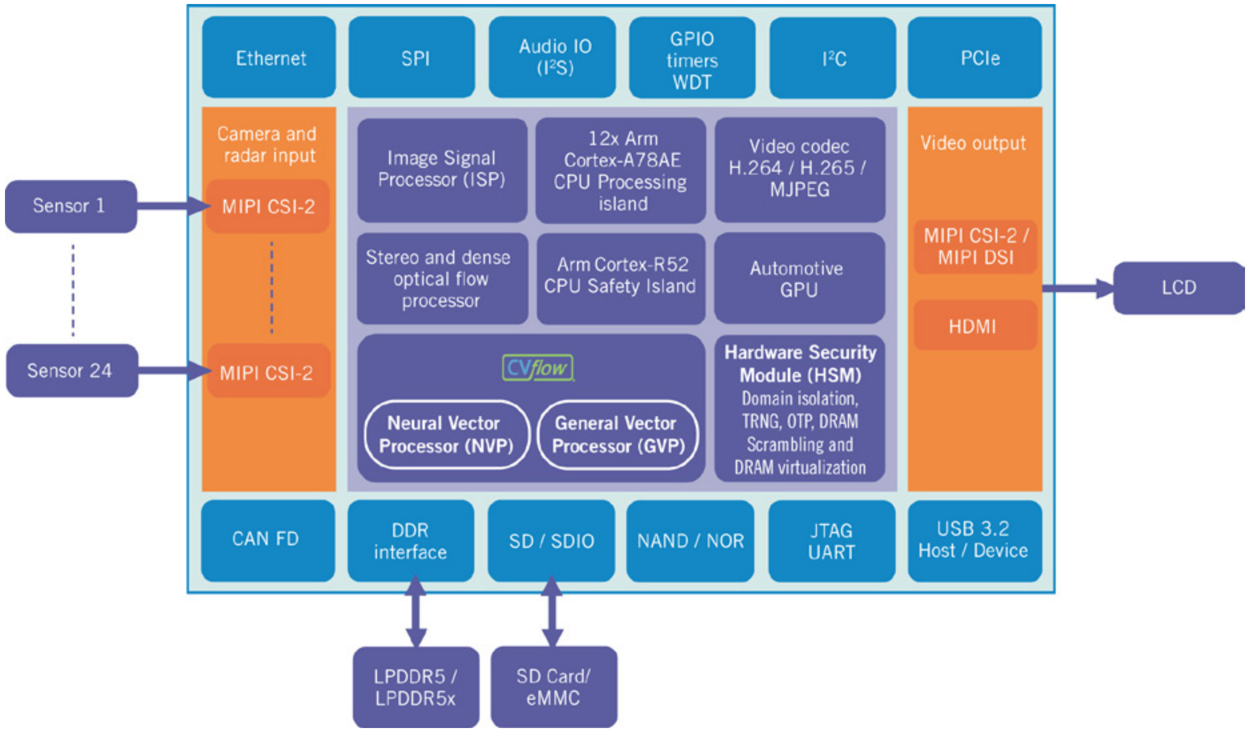

FIGURE 4 Block diagram of the CV3-AD685 SoC (© Ambarella)

SOC CHARACTERISTICS

The first version of the CV3-AD SoC family to go into series production will be the CV3-AD685 AI domain controller. 12 Arm cores of the Cortex-A78AE type are integrated in this SoC as a CPU. Further key hardware and software includes the image signal processor (ISP), a processor for stereo camera signals, a video codec, an automotive GPU for visualizations, a hardware security module, a security processor also based on Arm Cortex technology, and extensive peripherals and interfaces. As a central element of the design, two powerful vector processors are responsible for AI as well as parallel data processing, **FIGURE 4**.

The efficient hardware architecture, which is the result of two approaches, speaks in favor of the CV3-AD SoC family from a Tier 1 perspective:

- On the one hand, the hardware has been consistently designed according to the requirements of the software ("algorithm first"). Functions such as image processing and fusion through to trajectory planning therefore predetermine the hardware architecture. These specific use cases were taken into account during chip design.
- On the other hand, the SoC is highly energy-efficient and fast because bottlenecks as they can be found in some other SoCs have been prevented in the CV3-AD family. For example, the use of the RAM interfaces in this SoC is particularly efficient. In general, the internal communication and data flows within the SoC are optimized down to the last detail.

FIGURE 5 shows a comparison between three performance parameters of the CV3-High-Dev SoC with similar SoCs available on the market. In terms of

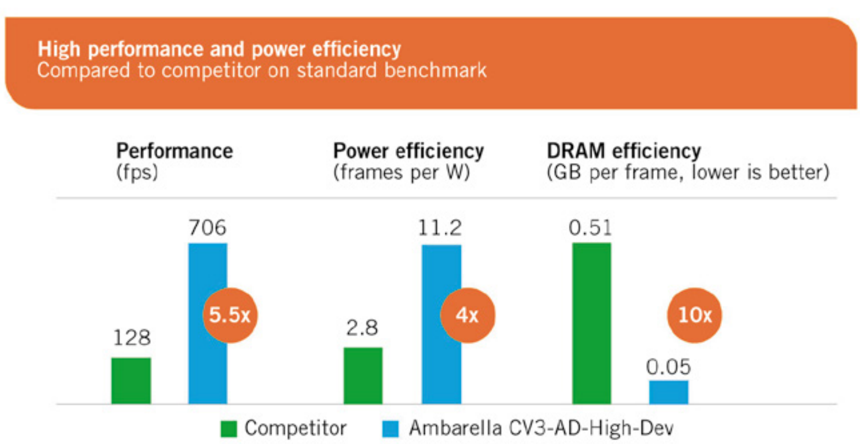

FIGURE 5 Comparison of key performance features of the CV3-AD-High-Dev SoC with similar SoC available on the market (© Continental)

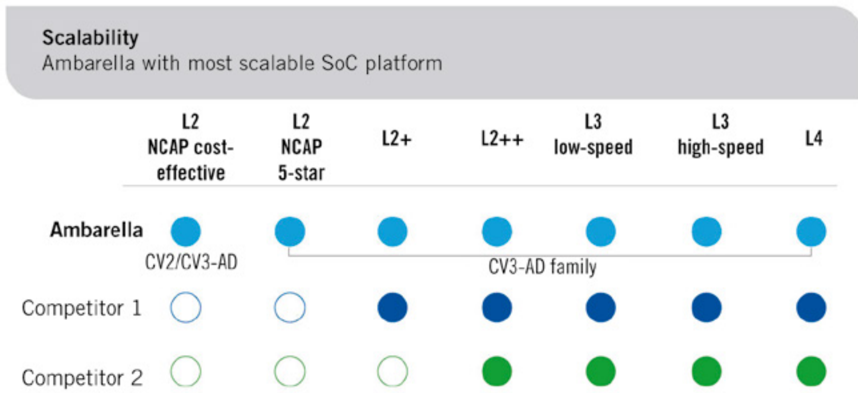

FIGURE 6 Scalability of various SoCs as comparison (© Continental)

computing power as well as energy and RAM efficiency, the extent to which the developers have optimized the silicon design in the CV3-AD SoC family becomes clear. In addition, hundreds of network types were analyzed during development to ensure that the SoC family is compatible with a wide variety of vehicle system architectures.

Compared with Ambarella's previous generation, the CVflow AI engine's neural vector processors (NVPs) within the CV3-AD family, for example, have up to 40-times higher performance. The CVflow AI engine is a specialized hardware module inside of all CV3-AD family SoCs. In addition, there is the computing capacity for the general vector processor (GVP), which is used to accelerate classical computer vision and radar processing. Optimization functions for the processing of high-resolution radar signals are also integrated here. An integrated hardware security module provides advanced cybersecurity elements, such as asymmetric/symmetric cryptography, secure memory, key management, encrypted computer vision tasks, a true random number generator as well as other features.

Deep data fusion is an example of the kind of efficient processing that can be implemented on Ambarella's CVflow AI engine. Pixels from a large number of high-resolution cameras, together with data from radars and echoes from laser scanners or ultrasonic sensors, are all fused together at a low level, thanks to a neural network running efficiently on the CV3-AD. This SoC's very high computational power allows it to perform low level fusion of the raw data from the many high-bandwidth sensors mounted all around the vehicle, making it possible to produce a single reconstruction of the vehicle's entire surroundings. Since all sensors contribute to this single representation, the data loss from decentralized, high-level fusion (i.e., fusing preprocessed object level data from each sensor, with a limited and partial view of the surroundings) will become a thing of the past.

SCALABILITY

Decisive for the scalability of the respective SoC derivatives is the balance of resources integrated into the SoC (in addition to the CPU cores, these include, for example, DSPs and NN processors as well as the hardware accelerators for certain image processing interfaces like image de-warping and pyramiding,) which decides on reasonable scaling levels, **FIGURE 6**.

The hardware scaling is not only based on the type and number of sensors, but also on the requirements associated with the application. It makes a big difference, for example, whether a given camera set is only used to generate an all-round image for the driver's information, or whether a lane change is also to be supported via camera image data processing. In both cases, for example, four cameras provide the signals, but the processing of the image data differs significantly with regard to resource requirements.

In terms of use cases, a single camera, for example, ranges from simple ACC to NCAP24/26 solutions with autonomous braking for pedestrians in more complex situations. The varying performance of person detection, alone, has a direct impact on the resource balance, depending on the number of variations in size, clothing and posture that need to be perceived. With this in mind, it is highly advantageous that in the CV3-AD SoC family, the internal units are easily scalable based on whether a particular resource is needed to a greater or lesser extent.

OUTLOOK

The aim of the partnership is to develop scalable full-stack vehicle system solutions with the highest performance for ADAS through to autonomous driving. The partners combine software and hardware expertise as well as Continental's broad portfolio of automotive system solutions with Ambarella's expertise in the AI field of computer vision, along with powerful SoCs and additional software modules.

In addition to the development of camera-based perception solutions for ADAS, the partners are focusing on end-to-end, AI-based solutions for further automation levels from Level 2+ up to L4. These full-stack systems take a multisensor approach that includes high-resolution cameras, radars, lidars and ultrasonic sensors, as well as the associated control units from Continental, plus the necessary software. The aim is to achieve more comprehensive and more precise environmental perception as the basis for high-performance driving functions, in the interest of further enhanced safety. A first production solution in commercial vehicles is planned for 2027, and the companies continue to target 2026 global series production for a broad range of automotive applications.

IMPRINT

Special Edition 2023 in cooperation with Continental Autonomous Mobility Germany GmbH, Ludwig-Erhard-Straße 3, 65760 Eschborn;

Springer Fachmedien Wiesbaden GmbH, Postfach 1546, 65173 Wiesbaden,; Springer Fachmedien Wiesbaden GmbH, Postfach 1546, 65173 Wiesbaden, Amtsgericht Wiesbaden, HRB 9754, USt-IdNr. DE81148419

MANAGING DIRECTORS: Stefanie Burgmaier | Andreas Funk | Joachim Krieger

PROJECT MANAGEMENT: Anja Trabusch

COVER PHOTO: © shutterstock | Petair

Continental
The Future in Motion

Autonomous Mobility.
We make it happen!

Autonomous driving is one of the major future fields of the automotive industry. Cruising, parking, driverless, human vision and safety technologies offer unsuspected potential for making mobility more sustainable, more comfortable and, most importantly, safer.