System on chip methods for Autonomous Cars

Mr. V. P. Sampath
Consultant & Sr. Member, IEEE
ramsampath78@rediffmail.com

Design Methodologies for the Automotive sector is changing.The average number of IP cores integrated into automotive SoCs is growing from nearly 20 about five years ago to more than 100 within the next few years. Modern trends are making an impact on the semiconductor IP providers who offer the functional ingredients that make up a chip. Some of the new application areas for automotive SoCs include ADAS, sensor fusion for autonomous driving, vision processing (front camera, object detection, and recognition, surround view, etc.), advanced sensor control and processing (LIDAR, RADAR, etc.) and machine learning for decision making functions in all these domains. Lidar units that Google uses in its self driving cars cost up to US$ 70,000 per unit, though there are now units that cost as little as US$ 250. GPS sub-systems are now available as sophisticated system-on-chip (SoC) ICs or multi-chip chipsets that require only power and antenna, and include an embedded, application specific compute engine to perform intensive calculations. Automotive Enhanced, meaning the chip has attributes targeting specific automotive use cases. Arm plans to extend its Automotive Enhanced IP solutions to a wider range of its products in the future.

Deep sub-micron processing technologies have enabled the implementation of new application-specific embedded architectures that integrate multiple software programmable processors and dedicated hardware components together onto a single cost-efficient IC. Application-specific architectures are emerging as a key design solution to today's microelectronics design problems, which are being driven by emerging applications in the areas of wireless communication, broadband networking, and multimedia computing.

SoC interconnect plays a vital role in facilitating functional safety because it interacts with all the data on chip. Consequently, on-chip communications are a critical building block in meeting the overall functional safety requirements. Selecting interconnect IP that is developed in accordance with the ISO 26262 functional safety specification can save OEMs and Tier-1s several man-months spent qualifying an automotive chip that must meet functional safety specifications.

Advances in hardware and networking will enable an entirely new kind of operating system, which will raise the level of abstraction significantly for users and developers. Such systems will enforce extreme location transparency. Any code fragment runs anywhere, any data object might live anywhere and the system manages locality, replication, and migration of computation and data and Self-configuring, self-monitoring, self-tuning, scaleable and secure.

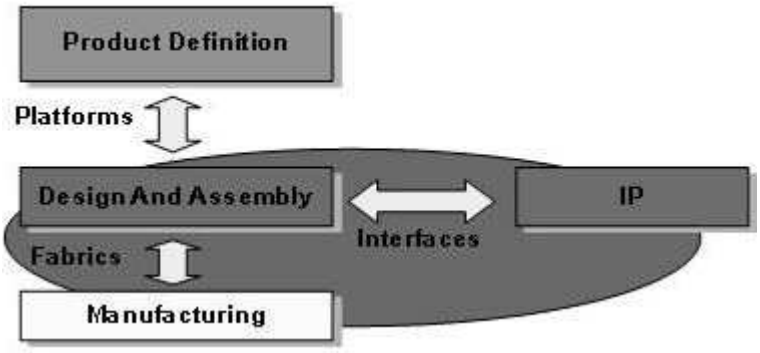

Figure1. Electronic System Design

New automotive SoCs utilize multiple specialized processing units on a single chip to perform multiple simultaneous tasks like camera vision, body control, and information display. The on-chip communications infrastructure is key to ensuring efficient data flow on the chip. And as the types and numbers of processing elements increase, the role of interconnects and memory architecture connecting these processing elements becomes crucial.

Issues in SoC limits:

- Economics
- Productivity
- Process
- IP Delivery & Reuse
- Tools & Methodology
- Manufacturing

Memories that are closely coupled to a single processing element are often implemented as internal SRAMs and are usually transparent to the running software. This approach works well for smaller systems, but an increase in the number of processing elements necessitates a corresponding increase in closely coupled memories. Another approach is to have RAM buffers that can be shared with multiple processing elements. However, in this case, access must be managed at the software level, which, in turn, can lead to software complexity as the system scales up. This software complexity can lead to systematic errors that can lead to errors and faults that affect ISO 26262 safety goals. Finally, as the systems become larger, it's often useful to implement hardware cache-coherence technology. It allows processing elements to share data without the overhead of direct software management. And there is a new technology for cache coherence, now widely implemented in automotive SoCs, that allows processing elements to efficiently share data with each other and as peers in the coherent system using a specialized configurable cache called a proxy cache.

Beyond memory architecture, whether it achieves data locality with buffers or is cache-coherent, what also matters is the on-chip interconnect. It optimizes the overall data flow to guarantee the quality of service (QoS) and thus ensures that automotive SoCs meet the bandwidth and latency requirements. Bandwidth allocation and latency requirements are a critical factor in mission-critical automotive designs, especially when some of the processing may be non-deterministic, such as for neural-network and deep-learning processing. Automotive designs are also providing the impetus for implementing new technologies like artificial intelligence (AI) because it is impossible to manually create "if-then-else" rules to deal with complex, real-world scenarios.

AI algorithms that can handle highly complex tasks are being incorporated into automated driving systems and other life-critical systems that must make decisions in near-real-time domains. That's why machine learning, a subset of AI, is the most publicly visible new application in self-driving cars.

Machine learning enables complex tasks in ADAS and automated driving through experiential learning that are otherwise nearly impossible using rule-based programming. But machine learning requires hardware customization for algorithm acceleration as well as for data-flow optimization. Therefore, in machine-learning-based SoC designs, the ADAS and autonomous car architects are slicing the algorithms more finely by adding more types of hardware accelerators. These custom hardware accelerators act as heterogeneous processing elements and cater to specialized algorithms that enable functions such as real-time 3D mapping, LiDAR point cloud mapping, etc.These highly specialized IP accelerators can send and receive data within the near-real-time latency bounds and deliver the huge bandwidth required to identify and classify objects, meeting stringent and oftentimes conflicting QoS demands.

Designers can compete and differentiate by choosing what to accelerate, how to accelerate it, and how to interconnect that functionality with the rest of the SoC design. Regarding new technologies, it's also worth mentioning that neural net networks have become the most common way to implement machine learning. What neural networks do here is implement deep learning in autonomous driving systems using specialized hardware accelerators to classify objects like pedestrians and road signs.

About the author

Mr. V. P. Sampath works as a consultant that develops hardware/software co-design tools. Among his publications are technical articles and papers on FPGA and Embedded systems and methods as well as textbooks. He is an active Senior Member of IEEE and Member of Institution of Engineers. He is a mentor for the semiconductor industries.

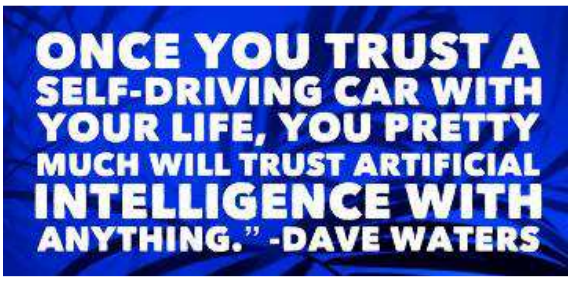

"self-driving cars are the natural extension of active safety and obviously something we should do."
-elon musk