Multi-Die Design for Automotive Applications

Author

Hezi Saar
Executive Director of Product Management

Introduction

Electronics for automobiles sit in the middle of two key trends in the semiconductor industry. The first is the increasing need for functional safety in many applications, bolstered both by user demands and formal standards. Safety considerations affect many aspects of chip development, from technology choices to system implementation. The other key trend is the move toward multi-die designs where homogeneous dies (also called chiplets) or heterogeneous dies implemented in different process technologies are integrated in a single package using standard or advanced package technologies. Yield considerations, ability to mix and match dies, optimizing power, performance, and area (PPA), and other factors are making multi-die design more attractive. This white paper examines the intersection of the two trends, as multi-die design is becoming an inevitable choice for automotive applications. This approach offers several advantages but is not without its challenges, and some of these are also discussed.

Overview of Automotive Chips

Vehicular electronics have always had some of the most stringent requirements across all applications. The operating environment is dirty and noisy, subject to vibration and beset by temperature and moisture extremes. Even the impact of potential bit flips can be more harmful due to safety-critical applications. Since the earliest appearance of electronic parts in automobiles, manufacturers realized that they needed to meet a higher bar than chips used in controlled locations. Only aerospace and government applications have stricter requirements.

At the same time, automobile users have high expectations for their vehicles in terms of reliability and safety. Drivers expect their cars to be ready for use anytime, anywhere, and to operate immediately. They want to be able to drive long distances, with only very rare failures that do not compromise their safety. They expect repairs to be quick and easy, minimizing their downtime. All these expectations go well beyond those for most other consumer devices. It is certainly frustrating if one's smartphone stops counting steps, but the result can be fatal if a lane correction system fails and steers the car into another vehicle.

Reliability and safety concerns are amplified by the increased role of electronics in modern automobiles. Consumer expectations are also growing; the more that a car can do the more the consumer demands perfection. In particular, the compute requirements of vehicles have grown immensely, and continue to grow every year. Users expect smartphone-class in-vehicle infotainment (IVI) experiences and semi-automated driving capabilities such as advanced driver assistance systems (ADAS) today, with fully autonomous vehicles soon. Figure 1 shows the standard industry definition for the various levels of autonomy.

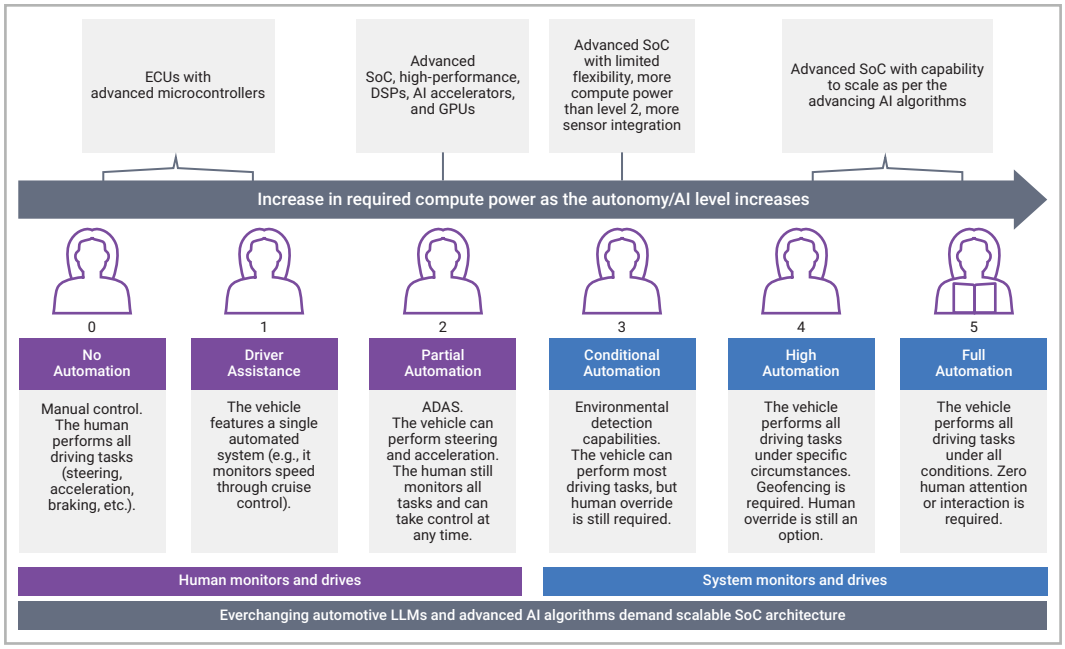

Figure 1: Levels of driving automation

The six levels of autonomy each place different requirements on the complexity of chips required to deliver them. Levels 0 and 1 have been available for years, using electronic control units (ECUs) with advanced microcontrollers. Achieving level 2 and higher requires complex system-on-chips (SoCs) with a combination of high-performance central host processors, digital signal processors (DSPs), AI accelerators, and graphics processing units (GPUs). The step from level 2 to level 3 is the most challenging. Level 3 requires orders of magnitude more compute power than Level 2, as well as a much broader array of sensors with higher resolution. Level 3 can be accomplished in a single SoC, although flexibility is limited because the hardware accelerators are tuned precisely for the software algorithms running on them. Levels 4 and 5 demand more compute power than Level 3, and really pushes the limits of what can be done in a monolithic SoC.

On the safety side, satisfying both consumer expectations and industry standards gets harder with each level. The main requirement is that all systems, modules, and SoCs must meet the ISO 26262 functional safety standard, which aims to reduce hazards caused by malfunctioning electrical and electronic systems. A wide variety of faults must be detected during real-world driving operation, including manufacturing defects not detected in production test, the effects of silicon aging, and damage from the harsh operating environment. Corrective action must be taken, by compensating for the fault if possible or by guiding the vehicle to a safe stop and refusing to start it again until the problem has been corrected by a reboot (transient fault) or a repair. Assessing the health of automotive chips over their entire lifetimes is mandatory.

In terms of reliability, requirements have grown dramatically due to the popularity of electric cars and plug-in hybrids. These have an "always on" profile since they are plugged in for much of their non-driving time. Since cars can last at least 15 years, their electronic systems must also perform reliably, safely, and securely over this long period of time. Building in fault tolerance and security (aligning with the ISO/SAE 21434 cybersecurity standard for road vehicles) is also mandatory. For bug fixes and vehicle upgrades, the use of over-the-air (OTA) software updates enables vehicle maintenance and extends their lifetimes.

New safety requirements, advanced node adoption, and zonal architecture are the three main reasons for the growing demands on chips in automotive applications. In addition to the standards-driven safety requirements, the higher likelihood of defects using advanced technology nodes makes test, yield, and field detection of faults even more detailed tasks. The third factor is the industry move toward zonal architectures, in which electronic components and functions are grouped within specific areas (zones) of the vehicle and connected to a central compute cluster. This makes it easier for the system to implement software-defined vehicle (SDV) capabilities.

Overview of Multi-Die Design

The definition of a multi-die design is simple: homogeneous or heterogeneous dies in different process technologies integrated in a single package. Although multi-die design has been used in some specialized applications for years, it is being adopted more widely due to two contrasting trends in the chip industry.

The first trend is scalability. New products can be assembled quickly from existing dies, especially variants targeted for specific applications, enabling a flexible solutions portfolio. Developing a multi-die design reduces time to market (TTM) over integrating the function into a new SoC and fabricating it. This also reduces project risk by reusing proven dies rather than relying on a new SoC. Even when a new die must be developed as part of the solution, it is smaller and less complex than an SoC and thus still reduces risk.

There are additional benefits to multi-die design. Die-to-die connectivity within a single package provides better throughput than chip-to-chip connectivity on a PCB. More functionality per unit area is achieved, especially when dies are stacked in a 3D configuration. Overall power is reduced significantly by reducing the number of PCB signals, although power per unit area increases and this must be managed properly. One common example of effective multi-die design is stacking memory modules on top of a processor within a single package.

The second trend driving adoption of multi-die design is aggregation, in which a set of discrete chips on a PCB is integrated into a single package instead of one large monolithic SoC. Delivering smaller dies provides several benefits. Yield may be better for multiple smaller dies than for a large monolithic die, especially in advanced technology nodes that stress the reticle limits of semiconductor manufacturing equipment. This may make aggregation of discrete chips or disaggregation of a larger monolithic die a more economical choice. Multiple dies make it possible to mix heterogeneous designs from different foundries, nodes, vendors, or technologies within a single package. Designers can scale the digital parts of the chip to the latest node while keeping analog functions in the same technology used for the SoC to avoid redesign effort and reduce risk.

Some configurations enabled by multi-die design include ADAS AI scaling using homogeneous dies for increased computing power for ADAS functionality, heterogeneous computing for an accelerator companion mode, and, as can be seen in Figure 2, I/O, memory, or any function split. This provides the ability to have highly scalable platforms able to compose the required set of multi-die features with different functional dies.

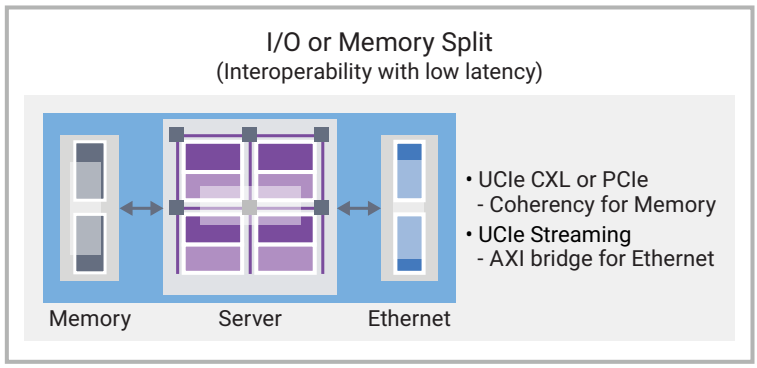

Figure 2: Example of I/O or memory split

Types of Multi-Die Package Design

A multi-die design offers much higher I/O density than a traditional PCB, helping to increase throughput and product performance. Figure 3 shows examples of the different forms of multi-die package designs available today, along with the relative I/O density achievable. Foundries offer two types of packaging—traditional packaging with standard bumps and advanced packaging with microbumps—with 2D and 2.5D options which is the focus for automotive applications. With advanced package technologies such as 3D stacking, through-silicon vias (TSVs), interposers, direct bonding, and finely spaced microbumps, a million times more density than a 2D PCB is possible. Many forms of multi-die design use the standard Universal Chiplet Interconnect Express (UCIe) for die-to-die connectivity.

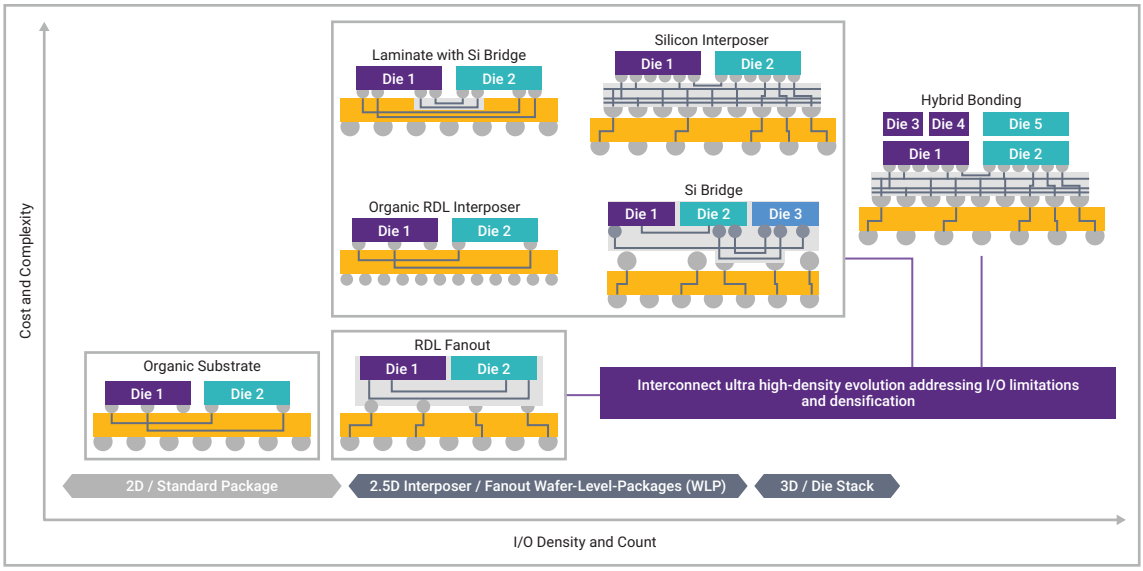

Figure 3: Types of multi-die package design

The Impact of Merging Trends

As noted earlier, automotive chips are affected both by increasing expectations and requirements, especially for reliability and safety, and general industry momentum toward multi-die design. These two trends are starting to merge as multi-die design is becoming an emerging solution for automative devices. This migration is driven by three key factors:

- AI-based ADAS features requiring more autonomy into Level 3 and beyond
- The digital cockpit, including driver monitoring, facial ID, and digital assistants based on generative AI (GenAI)
- Zonal architecture with virtualized apps, enabling scalable hardware and software architecture

The move to multi-die design has implications for many aspects of the chip development flow, including the electronic design automation (EDA) tools used and the IP selected for integration. The architecture phase shares many tasks with traditional monolithic chips. A digital twin is used for the multi-die architecture assessment "what if" as well as early software development. There are many ways to divide SoC level functionality across multiple dies, and selecting the optimal partition is a key step in the development process.

The design and verification tasks for multi-die design are also similar to monolithic chips, but the implementation phase is quite different. TSVs, interposers, microbumps, and die-to-die connectivity must all be incorporated into the physical implementation. This phase also involves planning for test and repair of devices in the field using silicon lifecycle management (SLM) techniques. IP such as process, voltage, and temperature (PVT) sensors, advanced ECC for on-chip memories, die-to-die monitoring, test and repair, and path margin monitors (PMMs) are part of this solution. All these elements must be incorporated into the physical implementation to monitor silicon health and support SLM. In multi-die design, SLM applies both to the individual dies and to the complete assembled package. A failure in a die-to-die interconnect can be as serious as a failure within a die.

When robust EDA and IP solutions are used for multi-die design, the resulting devices provide the power, performance, and flexibility needed for high reliability automotive applications. Auto makers and tier 1 subsystem suppliers may design their own chips, and this was more common when only a few manufacturers were developing advanced applications. Now that electric vehicles, ADAS, and IVI are appearing across the industry, semiconductor suppliers are developing new generations of automotive chips to be sold widely. As noted previously, using multi-die design makes it much easier for these suppliers to develop a family of chips that address different markets. Figure 4 provides a view of how this can work.

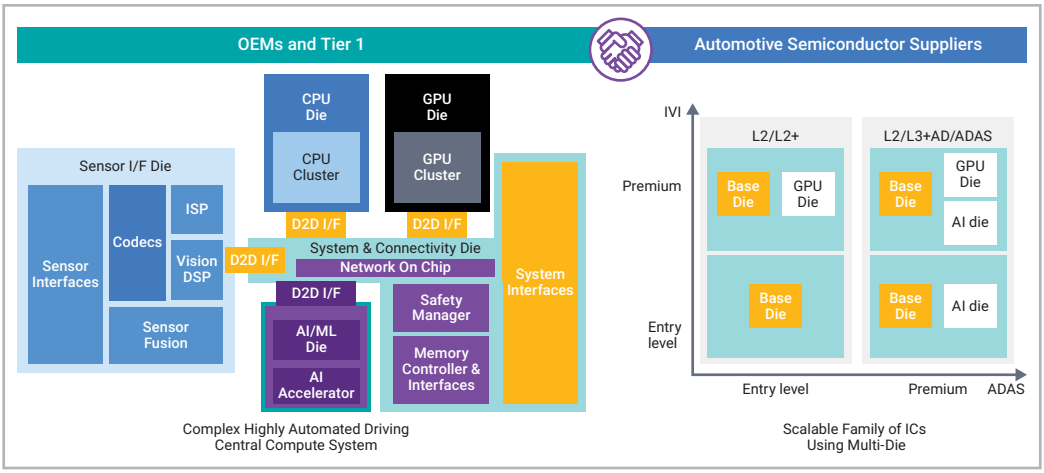

Figure 4: Evolution of automotive electronics with multi-die design

Synopsys Solutions for Automotive Multi-Die Design

Synopsys provides a complete automotive solution for multi-die designs, silicon-proven with numerous leading-edge chip manufacturers. It starts with Synopsys Platform Architect for Multi-Die, a standards-based performance and power analysis tool for early architecture exploration of multi-die designs. It enables “what if” scenarios and satisfies all requirements for the architecture phase.

The implementation phase is handled by Synopsys 3DIC Compiler, which provides a unified exploration-to-signoff platform for multi-die/advanced package co-design and co-optimization with multiphysics analysis.

Synopsys also provides a ASIL B and D compliant, automotive-grade IP portfolio to help designers accelerate the requirements specification, design, implementation, integration, verification, validation and configuration of their SoC level functional safety.

This rich portfolio enables a wide range of possible multi-die configurations, all compliant with relevant industry standards and initiatives. In the case of UCIe, the Automotive Working Group is considering appropriate extensions, including those for functional safety. The IMEC semiconductor research center is driving the Automotive Chiplet Program (ACP), Bosch is leading the Automotive Chiplet System Initiative, and the "Advanced SoC Research for Automotive" group has been established in Japan. Such organizations, along with standards such as ISO 26262, IEEE 1838, IEEE 1149.1, ISO/SAE 21434, and IEEE 2851, are helping to ensure that multi-die design for automotive applications satisfies all requirements for reliability, safety, and security.

SLM IP and the other components of the Synopsys SLM family of products provide particular value for the automotive market, as illustrated in Figure 5. This includes sophisticated ECC solutions to mitigate soft errors for critical on-chip SRAM/register file types of memories. At the silicon level, manufacturing support provides chips with quality as high as ten defective parts per million (DPPM) with a faster ramp to full production for shorter TTM and earlier new product introduction (NPI). Telemetry data provides insights for failure diagnosis and digital virtualization for future enhancements. At the system level, SLM feedback enables optimization tuning to manage voltage, temperature, and path margins. This improves functionality and use for CPU, GPU, and zonal clusters.

Within a single vehicle, continuous mission-mode monitoring detects aging and helps assess remaining useful life. In addition, predictive AI stress detection algorithms provide actionable insights for identifying anomalous versus systemic events. Across a fleet of vehicles, consolidation of feedback provides early warnings of safety or degradation issues. This enables faster remediation via OTA updates, lowering service costs and avoiding potential recalls.

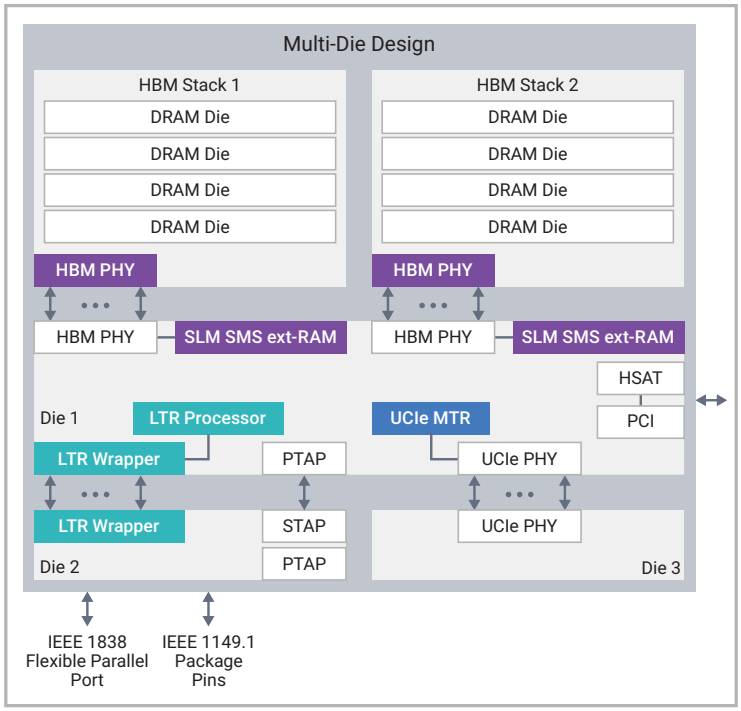

Figure 5: Value of Synopsys SLM for automotive applications

Summary

Automotive has always been one of the most demanding markets for semiconductor devices. Evolution in the size and complexity of automotive chips, driven by features such as IVI, AI cockpit, and autonomous driving, has dramatically increased the design requirements. Functional safety considerations and industry standards complicate the picture even more. It is not surprising when developers struggle to satisfy all the competing technical and time to market demands using traditional monolithic SoCs. For this reason, multi-die design is rapidly gaining favor for automotive ADAS and IVI applications. Incorporating advanced test and SLM increases chip yield, improves operation in the field, and predicts failures before they can disable vehicles or compromise safety.

Migrating to multi-die design has its own challenges, and the Synopsys EDA and IP products makes this process much smoother. When coupled with the company's expertise in multi-die design and experience they help users create production products and realize first-time silicon success. The resulting chips, systems, and vehicles have better PPA, improved reliability, more scalability, and a much greater chance of winning in the market. Figure 6 summarizes just some of the benefits to automotive manufacturers and their customers. Choosing Synopsys as their multi-die design partner makes all these benefits possible.

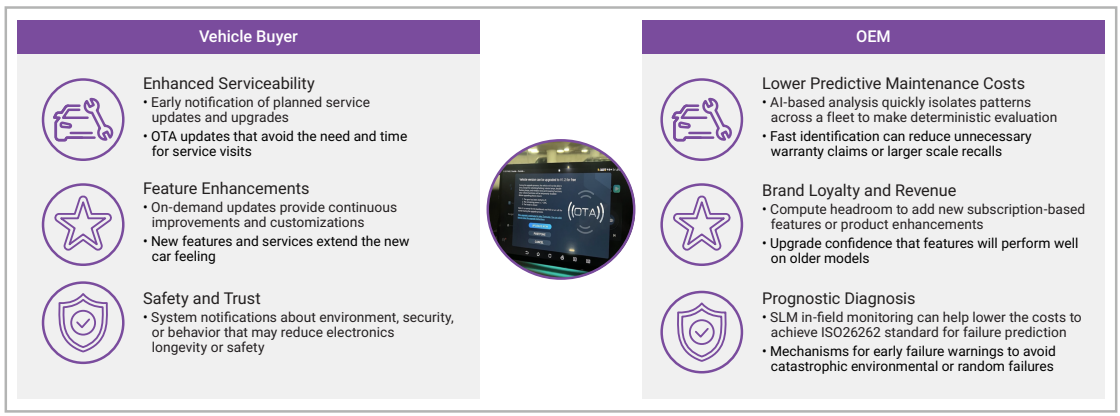

Figure 6: Benefits of Synopsys multi-die design for automotive

12/04/25.SNPS1797283967-Multi-Die-for-Automotive-WP.