Chiplets for Automotive – Are We There Yet?

Nir Sever, *Sr. Director, Business Development*
proteanTecs, 36 Kdoshei Bagdad St., Haifa 33032, Israel; nir.sever@proteanTecs.com

***Abstract*—For decades, Automotive Electronics were based on semiconductors manufactured on mature and stable process technologies. Designs were well characterized for robustness; tight screening at the production line enforced quality; using industry-standard test methodologies such as JEDEC JESD22 and JESD47 [1] assured reliability. Functional Safety (FuSa) relies on monitoring software, system redundancy, and safety protocols. Today, Electric Vehicles (EV) and Autonomous Driving (ADAS) require using the most advanced semiconductor technologies. Reliability requirements exceed those commonly used for commercial applications, device screening becomes a challenge, and safety measures that take effect after an error has already occurred may be insufficient.**

Recently, chiplet [2] based designs are driving the most advanced semiconductors for High-Performance Computing (HPC) and AI. Is chiplet-based design ready to be adopted by the Automotive industry?

Keywords— Semiconductor, Multi-die Packages, Heterogeneous Integration (HI), Chiplets, Die-to-die (D2D), UCI Express (UCIe), Automotive, Functional Safety (FuSa), Test, Production Test Flow, Continuous Performance and Health Monitoring (CPM,CPHM), Predictive Maintenance

I. Introduction

An excellent overview of the industry motivation, as well as challenges when adopting chiplet based designs is provided in the IEEE EPS Test TC Newsletter titled "Architecting Chiplets for Product Manufacturing Test Resiliency" [3], We recommend reading the referenced work to set context for this newsletter article, as we address some challenges laid out, especially in the context of Automotive. This paper will focus on the challenges and proposed solutions related to the high bandwidth interconnect between the chiplets.

II. Chiplet Interconnect Test Challenges

The characterization tests of high-speed interconnect in monolithic designs are based on running traffic across the primary interface between the Device Under Test (DUT) chip and a dedicated protocol tester running a certification software, e.g., Keysight's P5570A PCIe 6.0 Protocol Analyzer [4]. However, such a method is not possible for the chiplet die-to-die (D2D) interconnect since after assembly, these interfaces are no longer exposed as primary IOs, and visible to the external protocol characterization equipment.

The same challenge applies to the testing of the product during production. Nowadays, the most advanced wafer probing is designed for a bump pitch of 75-45um or higher [5]. With a typical chiplet bump pitch of 45um or smaller (25um for "Advanced Package" profiles and less than 10um for 3D stacking), wafer probing is no longer possible, and even if it was, the delicate bumps likely would be damaged during the process.

After assembly, the die-to-die interface is no longer exposed on the primary IOs and, therefore, is inaccessible from the Automated Test Equipment (ATE). This inaccessibility leaves us with thousands, sometimes hundreds of thousands, of interconnect lanes that cannot be tested. We often say these D2D interconnects are "blind spots" to the ATE and are a test coverage hole. The following image illustrates that concept.

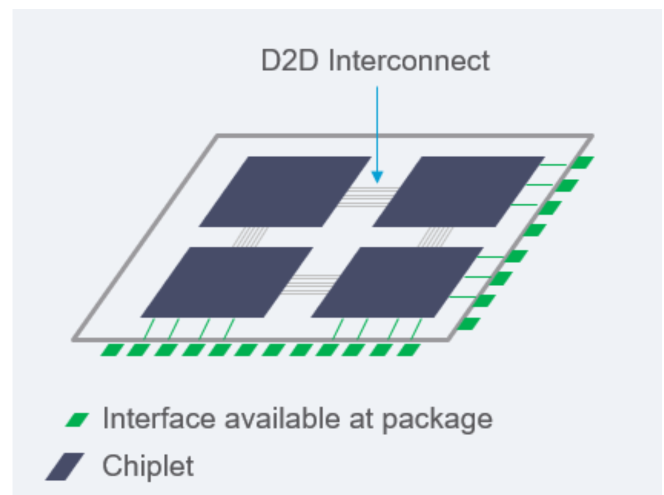

Fig. 1: D2D Interconnect is an ATE "blind spot"

As illustrated in Fig 1, only a small portion of the chiplet IOs (shown in green) are exposed as primary IOs, while most of the interconnect (grey) are only internal and hidden.

The standard method of testing these D2D interconnects involves setting the interface into a "loopback mode" and running a BIST test. This method is useful for identifying hard faults such as opens or hard shorts because such faults will cause a functional error that the BIST response checker catches. However, some of the known assembly-related defects can cause a degradation of performance (data eye closure and jitter) and may still pass the BIST Pass/Fail test. Additional limitations of BIST are that the traffic pattern is not

"real", i.e., not necessarily representing the expected traffic patterns and workload, that loopback is predetermined, and can miss interactions between specific lanes. We consider such cases of undetected faults as "walking wounded" devices. Such defects tend to degrade faster and cause premature failure. To summarize, BIST-based Pass/Fail test solutions, are insufficient for high availability and mission-critical applications, especially Automotive, and better methods are needed.

In the next section, we present the concept of in-chip monitoring and lane grading to address this problem and provide the quality and reliability assurance mechanism needed by Chiplet-based Automotive System In Package (SIP) vendors.

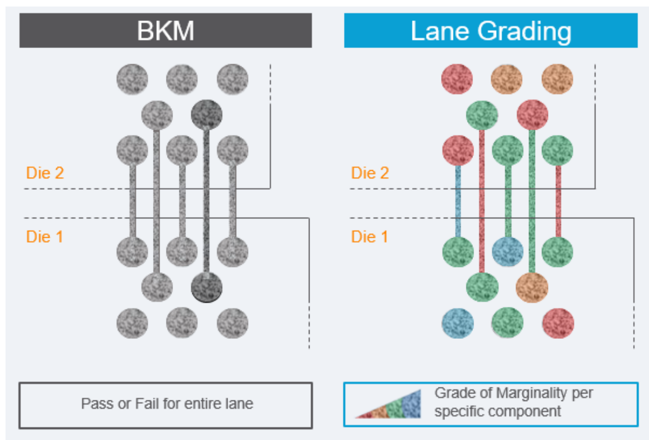

Fig. 2: Parametric Lane Grading: best-known methods vs. in-chip monitoring

The rest of the document will discuss the three main challenges identified for assuring high quality and reliability Chiplet-based SIP:

1. High visibility for accurate characterization of the interface will be presented in sections III and IV
2. High quality manufacturing and DPPM reduction is described in section V
3. Lifetime performance and health monitoring concept is described in section VI

III. Parametric Lane Grading of High Bandwidth D2D Interconnect

In-chip monitoring of high bandwidth, parallel interconnect is based on inserting a measurement circuit on every interface lane. Such a measurement circuit must be small enough to fit inside the bump array area along with the transceiver circuit, must not affect the signal quality and performance, and must consume just a small amount of power compared to the per-bit energy of that interface. Such a measurement element can be inserted on every lane of the interface to provide 100% lane coverage and must be able to operate in test and in mission modes. The parametric measurements must be granular enough to show, for example, the worst-case lane performance and averages for detecting outliers. The next illustration shows an example of such fine-resolution measurement of the data eye:

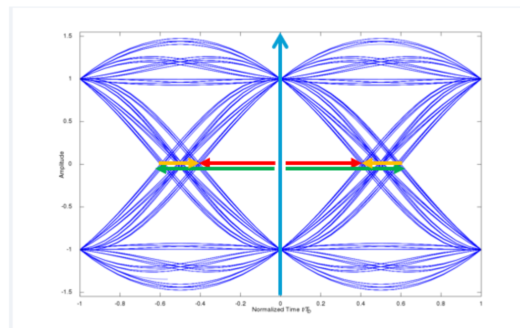

Fig. 3: Data eye width measurements of D2D interconnect [6]

In this example, the measurements represent the maximum, minimum, and jitter of data crossing to clock and clock to data in the time domain. The image shows one of the clock phases, but all the clock phases must be used for the measurements. Due to the large amount of data measured, data analysis must be done by hardware inside the chip, and only "meaningful" information, should be kept for later use.

The following sections provide examples of how such data can be used for more accurate characterization and reliability testing, production testing, and outlier detection.

IV. Lane Grading Use for Characterization

As explained before, since D2D lanes are not exposed as primary IOs, in-chip monitoring can be used to characterize them for verifying the design robustness after assembly. Please consider the following image:

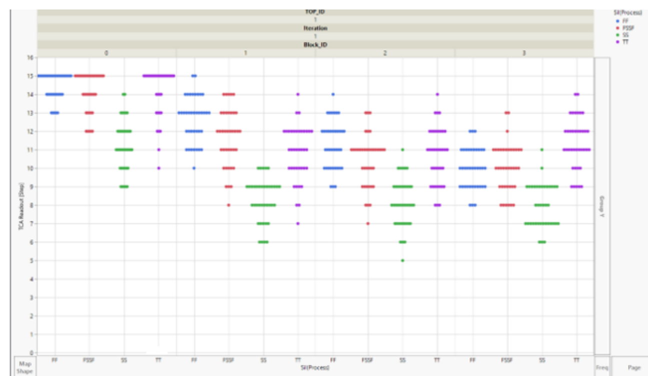

Fig. 4: Interface characterization using lane monitoring data

In this example (Fig 4), four corner samples are characterized across voltages and temperatures. Each color represents one corner sample; each dot is one lane worst-case measurement across a test sequence. The higher the reading, the more "eye-opening" occurs. Here, all samples show very good eye-opening, with only the Slow-Slow corner (in green) showing slightly less performance, as was expected. Interface characterization equipment makers can use such

measurements to test, characterize, and certify D2D interfaces. proteanTecs and Global Unichip (GUC) jointly presented additional characterization challenges and results in a paper titled “GUC GLink™ Test Chip Uses In-Chip Monitoring and Deep Data Analytics for High Bandwidth Die-to-Die Characterization” [7].

V. Outlier Detection in Production

To achieve high-quality production screening, testing the final assembled product aims to ensure all lanes perform as expected. Most parallel interfaces, such as HBM2, HBM3 [8], and UCIe [9], have redundant lanes defined in the standard, which can be used at the production line to replace broken or marginal lanes. The following image shows the test results of a large manufacturing batch that uses lane grading and outlier detection algorithms.

Fig. 5: Lane outlier detection

In this example (Fig 5), higher equals worse. Here, two packaged units were identified as outliers and marked as Unit #1 and Unit #2. Let’s look into these two units.

Fig. 6: Unit #1 Lane outlier detection

When drilling-down to Unit #1 (see Fig 6), we can easily see that two outlier lanes showed significantly worse (higher) measurements than the rest of the population. We can then locate them and see that they are two adjacent lanes in the same lane group. Since this lane group had two redundant lanes, they were used to replace the outlier lanes, and this unit could be shipped.

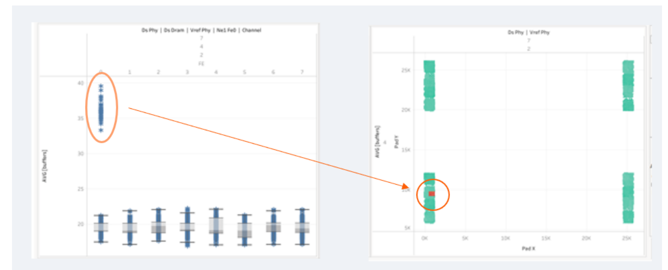

Fig. 7: Unit #2 Lane outlier detection

Drilling-down to Unit #2 (Fig 7) shows that many lanes are flagged as underperforming. When locating these lanes on the device, we can see they are all in the same area. Further failure analysis revealed a slight delamination that caused a parametric shift of these lanes. So, discarding this device is the only option.

In both cases, the units passed all their acceptance tests and would have shipped as good devices, and likely cause errors under stressful conditions if not flagged by using in-chip parametric lane grading.

VI. In-Mission Performance and Health Monitoring

For mission-critical applications such as Automotive, the assurance of quality and reliability does not end at production, and functional safety (FuSa) standards require that the device detect and enter a safe state under fault occurrence when used in the field. New safety concepts suggest adding continuous performance and health monitoring in mission mode and using it to enable predictive maintenance. Such capabilities are now suggested by the “Predictive Maintenance” efforts of ISO26262 [10] and IEEE 1856 [11,12]. UCIe V1.1 that was just announced now includes lane monitoring registers in the standard, allowing software tools to read and analyze the data during mission mode.

The following image is a reference application that continuously reads all the lane measurements from the device, identifies the worst lane, and compare it to a spec limit and to a calculated margin limit for flagging marginal lanes before they reach the point of failing.

Fig. 8: Continuous health monitoring in mission mode

This screenshot (fig.8) is from a running system that periodically reads measurement data. The grey curves represent reading from each lane; the black curve (lowest) is the lane identified as the most marginal. Marked is a

"degradation event" purposely injected to demonstrate worse lane crossing the "control threshold" to trigger a Diagnosis flag and reduce the Health Indicator, a heuristic indicator of the system health. This event could trigger redundant lane swapping in the field or scheduling of reactive maintenance. The same method could also be used during reliability stress testing of the system to continuously monitor and predict Failure Rate (FR) and Time To Failure (TTF).

VII. Conclusions

On his LinkedIn post [13] from Dec 2022, François Piednoël, UCIe Member and Distinguished Chief mSoC Standard Architect at Mercedes-Benz Research & Development North America, Inc., said: "If your product portfolio does not include chiplets in 2024, you are behind the followers ... You are ... late".

Chiplet-based architectures promise to change how we design complex semiconductor products, and Automotive is no exception. We are already seeing Automotive and chiplet interconnect standardization groups proposing novel methods to ensure that quality and reliability requirements are met with complex and advanced node designs, including Chiplets.

In-chip monitoring and parametric lane grading is the technology needed to address Automotive industry's main concerns when considering Chiplets by:

- Powering accurate and deep-data-based characterization and qualification testing
- Enabling high-precision outlier detection and redundant lane swapping during Mass Production (MP)
- Maximizing availability and reliability by Continuous Performance and Health Monitoring (CPM, CPHM) applications and Predictive Maintenance

VIII. Acknowledgment

The author would like to thank Intel's Abram Detofsky for his vision and leadership, GUC's Igor Elkanovich for his guidance and openness to share our joint work, and proteanTecs's Andrea, Eyal, Michael and Tamar for their valuable and insightful comments and suggestions.

IX. References

[1] STRESS-TEST-DRIVEN QUALIFICATION OF INTEGRATED CIRCUITS | JEDEC

[2] https://en.wikipedia.org/wiki/Chiplet

[3] Resiliency_IEEE_EPS_TC_Newsletter_230222_detofsky_final

[4] PCI EXPRESS Protocol Solutions | Keysight

[5] Altius Probe Cards - Large I/O Logic Device Testing | FormFactor Inc..

[6] Eye illustration by Blair Bonnett - Own work, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=3198001

[7] GUC GLink™ Test Chip Uses In-Chip Monitoring and Deep Data Analytics for High Bandwidth Die-to-Die Characterization (proteantecs.com)

[8] Standards & Documents Search | JEDEC

[9] https://www.uciexpress.org/specification

[10] ISO/DTR 9839 - Road vehicles — Application of predictive maintenance to hardware with ISO 26262-5

[11] 1856-2017 - IEEE Standard Framework for Prognostics and Health Management of Electronic Systems | IEEE Standard | IEEE Xplore

[12] Functional Safety Standards Committee: Results and Perspectives (computer.org)

[13] https://www.linkedin.com/posts/francoispiednoel_intel-chip-research-pushes-power-efficiency-activity-7005765216558870528-Poao?utm_source=share&utm_medium=member_desktop