Microelectronics Reliability 184 (2026) 116230

Contents lists available at ScienceDirect

Microelectronics Reliability

journal homepage: www.elsevier.com/locate/microrel

Research Paper

Innovative FA hardware solution to enable system-level debug of 3D ICs

Lesly Endrinal [a,c,*], Jehan Saujauddin [b], YuanChung Ho [a], Dilbagh Singh [a], Sasi Sekaran Sundaresan [b], Vinod Kumar Kakumanu [a], Jessen Gonzalez [b], Willem D. van Driel [c], G.Q. Zhang [c]

[a] Google LLC, 6420 Sequence Drive, San Diego, CA, 92126, United States of America
[b] Google LLC, 2015 Stierlin Ct, Mountain View, CA, 94043, United States of America
[c] Department of Microelectronics, Delft University of Technology, the Netherlands

ARTICLE INFO

Keywords:
3D IC
Electrical Fault Isolation
Failure analysis
Interposer
Optical Fault Isolation
Package-on-Package
Photon Emission
SLT FA
System Level Test
System-on-a-Chip

ABSTRACT

The rapid growth of generative Artificial Intelligence (AI), automotive electrification and Internet of Things (IoT) is driving an unprecedented demand for high-performance Integrated Circuits (ICs), projected to push the semiconductor industry to a $1 trillion business by 2030 (Burkacky et al., 2022; PWC, 2025). This drove the adoption of advanced process (FinFET, Gate-All-Around, Forksheet, CFET, Backside Power Delivery Network) and 3D package technologies (chiplets, Heterogeneous Integration, Co-packaged Optics). Unfortunately, component stacking in 3D ICs such as Package-on-Package (PoP) products creates an optical barrier during Electrical Fault Isolation (EFI).

This paper presents a novel Failure Analysis (FA) hardware and sample preparation solution that enables System-Level FA on mobile System-on-a-Chip (SoC) PoP devices, where the DRAM is stacked atop the logic controller device. The primary challenge in performing EFI on the bottom die is providing direct line-of-sight (LoS) access while preserving the top DRAM functionality through extremely small interconnects called Through Interposer Vias (TIVs). For the first time, this groundbreaking hardware solution overcomes the challenge of connecting a DRAM atop the SoC silicon, utilizing the smallest possible interposer pin pitch (~210 um) and advanced DRAM card design that incorporates stringent design rules to enable up to 6.3 Gbps using a DDR training test and runs standard Android stress applications. The results of the FA hardware solution for SLT platform will be discussed, together with several FA use cases that were enabled through this innovative solution. Lastly, this paper will discuss the limitations of this hardware, together with opportunities for improvement and further work.

This novel solution paves the way for the failure analysis of 3D IC devices on a system level platform, which are utilized not only in mobile applications, but also for cloud, server, and quantum computing ICs assembled in complex packages.

1. Introduction

Failure Analysis (FA) plays a critical role in post-silicon development by collaborating with cross-functional disciplines such as design, product, test, process, and packaging to identify product weaknesses and process improvements. A successful root cause analysis is essential to accelerating yield ramp-up and ensuring the delivery of high-quality, cost-effective and reliable chips.

Failure Analysis is divided into two major phases: Electrical Fault Isolation (EFI) and Physical Failure Analysis (PFA). EFI typically utilizes Optical Fault Isolation (OFI) techniques, which are non-destructive and uses light in the near infrared (NIR) and mid-wave IR (MWIR) wavelengths to diagnose failures and isolate them. Such techniques include Photon Emission Microscopy (PEM), Thermal Analysis and Lock-in-Thermography (LIT), Dynamic Laser Stimulation (DLS), and Laser Voltage Imaging/Probing (LVX). Although advancements in Design For Test (DFT) have paved the way for logic and memory diagnostics that have a potential to predict the physical location of failures, such solutions are not always sufficient to guarantee a successful direct PFA result. The goal of EFI is to further isolate the failure and narrow down the Area of Interest (AOI) to a smaller region, to increase the chances of success with PFA. Once the failure is isolated to a specific cell or net, PFA

* Corresponding author at: Department of Microelectronics, Delft University of Technology, the Netherlands.
E-mail address: l.z.endrinal@tudelft.nl (L. Endrinal).

https://doi.org/10.1016/j.microrel.2026.116230
Received 1 May 2026; Received in revised form 15 June 2026; Accepted 30 June 2026
Available online 3 July 2026

is performed by delayering the device and performing a thorough inspection and electrical measurement down to the transistor level to identify the failure mechanism. Since PFA is a destructive approach, it is critical to have an accurate EFI localization for a successful PFA.

Due to advancements in DFT, most of the failures in the chip can be screened out using an Automated Test Equipment (ATE) that runs a comprehensive set of structural DFT test patterns, also known as Automated Test Pattern Generation (ATPG) vectors. Using ATPG and fault simulations, the location of the failure can be predicted using ATPG diagnostics. In addition, specific DFT hardware such as Memory Built-in Self-Test (MBIST) and Logic Built-in Self-Test (LBIST) are integrated into the chip design to facilitate automated self-testing and diagnosis from within the IC. Similar to ATPG, these MBIST and LBIST test patterns are run using an ATE. In addition to the structural and parametric tests, a typical ATE test program would also incorporate a set of functional tests to cover the functionality of the different IPs within the chip. While ATE tests are successful in screening out defects with more than 90% coverage, some failures still go undetected during Final Test (FT) and get screened at System Level Test (SLT). Hence, system level testing is a very critical component of device testing, to ensure that the product is defect-free and performs according to the specification.

2. Problem statement

The advent of advanced technologies fueled by AI and data servers, high performance computing, vehicle electrification and Internet of Things (IoT) has accelerated the adoption of 3D IC, chiplets, Co-Packaged Optics (CPO), and heterogeneous integration (HI) package technologies. The semiconductor industry is expected to grow to $1 trillion by 2030 [1,2]. In mobile System-on-a-Chip (SoCs), processors usually come in a 3D stack where the bottom die is typically a logic controller and a Low Power Dual Density Rate (LPDDR) DRAM IC is stacked on top of it. Refer to Fig. 1.

The latest DRAMs are mostly developed using LPDDR5 and LPDDR6 technologies. LPDDR5 was introduced in 2019 [3] and provides data rates of up to 6.4Gbps with improved energy efficiency using Dynamic Voltage Scaling (DVS) [4,5]. In 2021, an upgraded version, LPDDR5X, promised improved performance with up to a 10.67 Gbps data rate [5,6]. LPDDR5/5× remains the most commonly used DRAM in today's mobile SoCs. In July 2025, JEDEC published the new LPDDR6 specification [7], which is expected to deliver up to a 14.4 Gbps data rate [5,7] to meet the high compute demands of mobile SoCs and IoT devices.

Most of the advanced digital EFI is carried out using ATE-based test platforms that utilize structural test patterns such as ATPG and MBIST vectors that focus on the bottom logic die. Such configurations do not require the DRAM to be present, and hence, OFI techniques such as PEM, LIT, DLS and LVX are performed on the SoC by grinding off the DRAM and exposing the backside silicon of the logic die to the optics of the FA tool. An illustration of this setup is shown in Fig. 2.

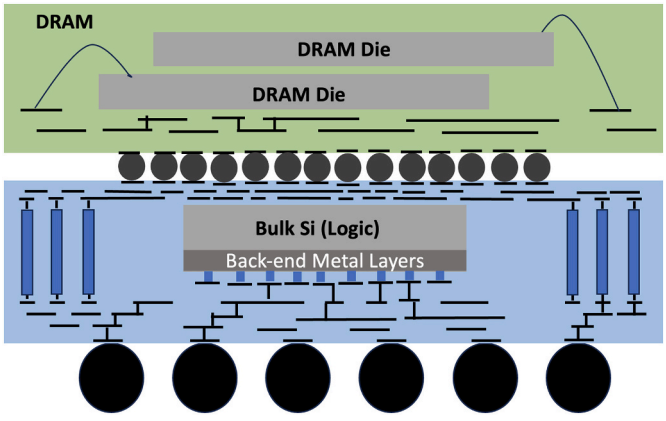

Fig. 1. Cross-section diagram of a typical SoC device, with the logic die at the bottom and the DRAM chip at the top.

Unfortunately, not all failures can be screened by the automated test equipment (ATE) alone. We have observed failures that only manifest themselves during System-Level Testing (SLT), even when these devices have completely passed ATE screening. In such cases, it is crucial to enable EFI on the SLT platform to find the root cause of the failures, which requires the DRAM to be intact to execute instructions and facilitate Android-based testing.

The main challenge in performing OFI at system level is achieving full optical access to the logic die at the bottom, while simultaneously preserving system-level functionality through the use of the top DDR. Fig. 3 illustrates this challenge, as the DDR in the diagram fully obstructs the optical path to the logic die.

3. SLT FA hardware requirements

Beyond the fundamental requirement for line-of-sight (LoS) optical access, the FA hardware solution for system-level testing must also satisfy the operational specifications of the OFI tools. Specifically, the PCB and FA socket design must take the optical lens system into consideration. An example of a lens configuration is shown in Fig. 4.

The FA hardware and application board requirements include:

- **Full die exposure:** The fixture aperture must fully expose the silicon backside of the logic die to facilitate unobstructed OFI, while maintaining secure contact with the TIVs to ensure electrical connectivity between the logic die and the DRAM.
- **Application Board Keep-out Zones**: A typical OFI system utilizes a turret configuration for lens selection, as shown in Fig. 4. To prevent mechanical interference or damage between board components and the objective lenses, a keep-out zone with diameter “X” must be established around the Device under Test (DUT) center. However, on high-speed boards, this requirement often conflicts with the need to place decoupling capacitors and other critical components in close proximity to the SoC to maintain signal integrity.
- **Component height**: If the keep-zones cannot be strictly met, the hard requirement is to avoid any component height that sticks out above the top surface of the DUT within the keep-out zone diameter “X”. This means that any connectors, jumpers or components with high profile will need to be removed or modified. Straight connectors will need to be replaced with 90-degree connectors to avoid potential damage to the lens system.
- **FA Socket Lid working distance**: Every air-gap objective lens has a specific working distance (WD), which is the physical clearance

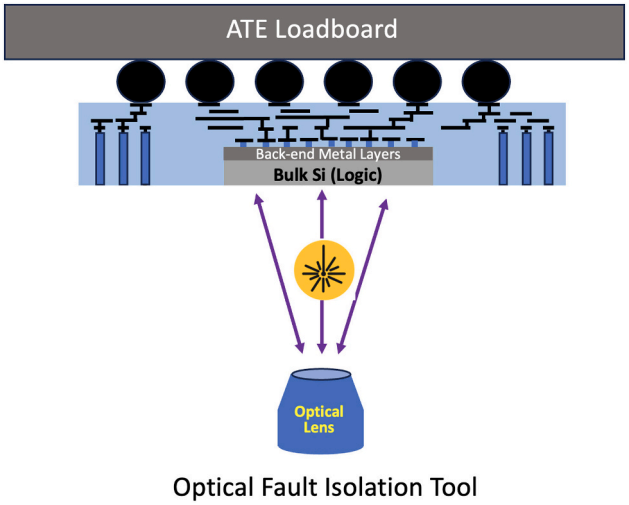

Fig. 2. Example of an ATE-based OFI setup.

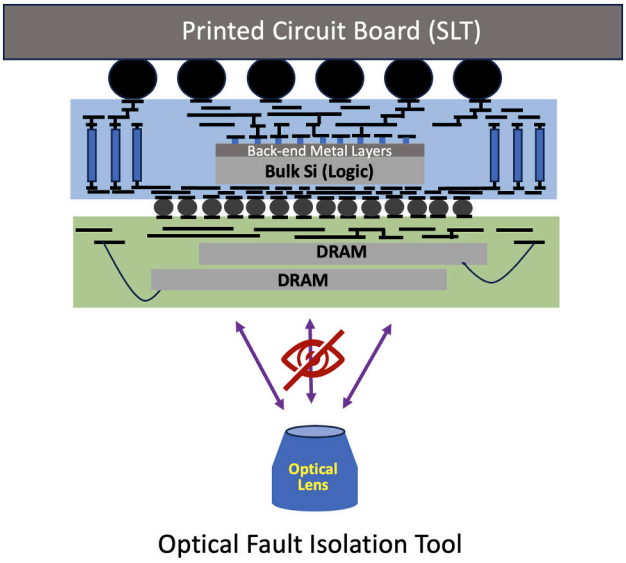

Fig. 3. Example of a system-level test (SLT) setup.

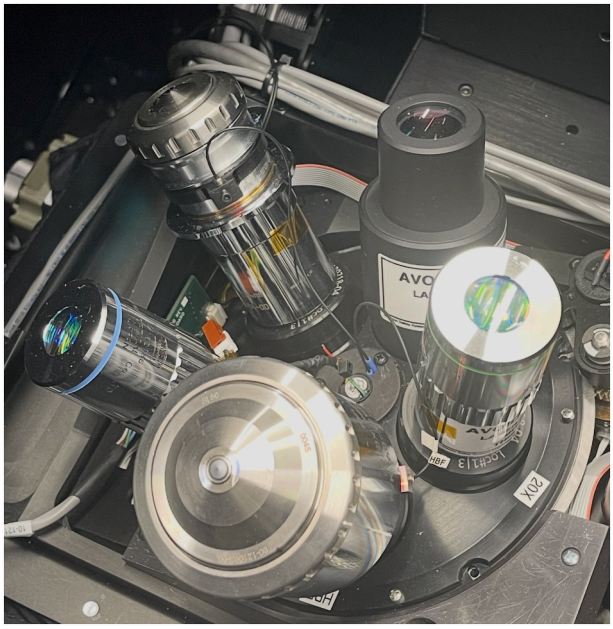

Fig. 4. Example of an OFI lens configuration.

between the front of the lens and the DUT. This typically ranges from 10 to 20 mm. Consequently, careful consideration is required when designing FA socket lids to ensure they securely mount the DUT without obstructing the optical path or interfering with the lens's focal range. An example of this setup is shown in Fig. 5.

- **Solid Immersion Lens (SIL) access**: SIL is an objective lens with a high refractive index used to increase the numerical aperture (NA) of the optical system to improve spatial resolution. The SIL is placed directly in contact with the silicon surface to minimize the refractive index mismatch between the sample and surrounding medium (typically air), essentially improving the light collection of imaging systems [8]. To ensure proper SIL contact, any obstruction around the die must be removed. Hence, Flat Top Socket (FTS) designs and PCB interposers called "carrier" or "coupon" boards are developed [9] to provide open optical access to the silicon backside. Such a

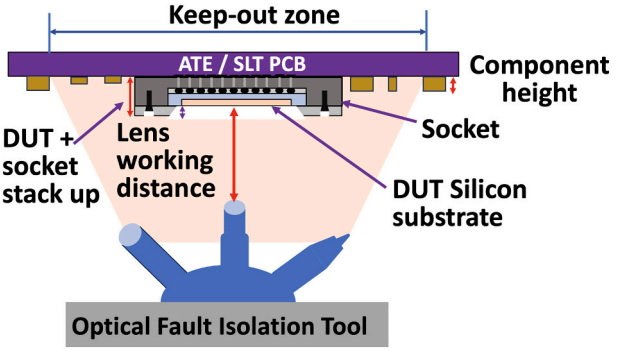

Fig. 5. FA hardware using lidded socket.

setup is shown in Fig. 6A, while an illustration of a SIL in contact with a device on a PCB interposer is shown in Fig. 6B.

- **Other requirements**: Aside from keep-out zones and lens working distance requirements, other critical design constraints for hardware development include:
  1. SLT board docking hardware design and standardized mounting to ensure mechanical stability when interfaced with the OFI tool.
  2. Connectors, headers, and switches on the DUT side must be avoided and relocated to the other side of the board to prevent mechanical interference with the lenses.
  3. If manual switches are required, they must be moved to the opposite side of the PCB or replaced with remote-controlled interfaces, such as wired switches, to allow for operation without disturbing the physical setup during OFI.
  4. LEDs and any external light sources must be avoided and removed to facilitate Photon Emissions. Consequently, board designs can incorporate a bypass switch for these components as an alternative.
  5. Provisions for GPIO access points must be included in the board design to provide synchronization and triggering signals required for LVX and DLS.

4. SLT FA hardware solution

To enable FA on the bottom logic die, it is necessary to establish a direct line-of-sight (LoS) optical access while DRAM remains connected. This section introduces a novel hardware architecture and an innovative sample preparation technique designed to enable system-level FA debug.

4.1. Flat top socket and PCB interposer

Production SLT boards typically employ pocketed sockets with lids to secure the device during testing. However, these lids obstruct the optical path, making them unsuitable for OFI. To address this, a "Flat Top Socket (FTS)" [10] is used. Flat top sockets are widely utilized for FA applications on the ATE-based systems. An example is shown in Fig. 7. In this configuration, the DUT is soldered onto a PCB interposer, often called a "carrier" or "coupon" board, which is then fastened to the FTS with screws. This setup eliminates the need for a lid, though care must be taken to ensure that no surrounding components exceed the height of the DUT, to avoid mechanical interference with the lenses. To secure the device further, an underfill material is applied to provide mechanical rigidity. During assembly, appropriate solder reflow conditions must be followed to avoid package warpage.

Fig. 8 shows an example of an SLT application board, which is mounted on a custom-designed docking hardware. In this illustration, the production socket was replaced with a Flat Top Socket. The DUT is mounted on a PCB interposer, which is fastened to the FTS with a set of 6 screws.

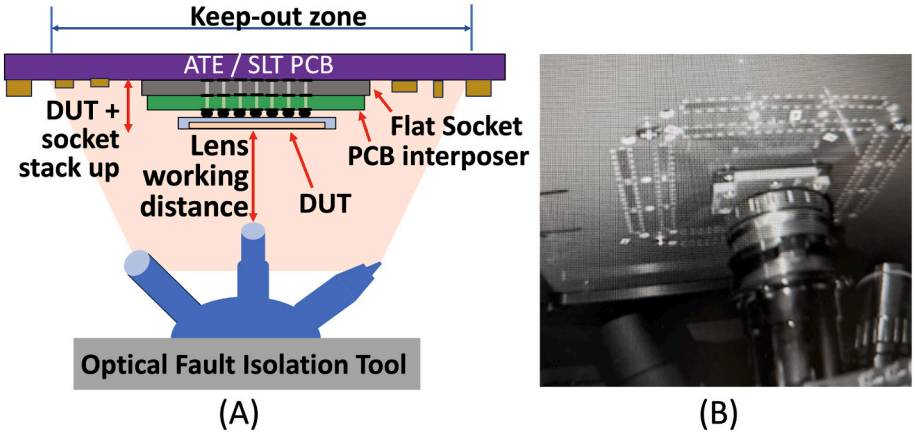

Fig. 6. (A) Flat top socket with PCB interposer for SIL access (B) SIL in contact with a DUT mounted on PCB interposer.

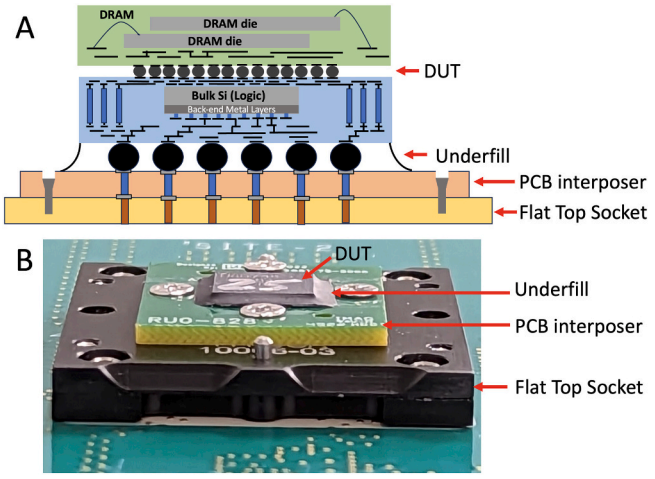

Fig. 7. (A) Cross-section diagram of a DUT mounted on PCB interposer and flat top socket (B) image of a DUT mounted on PCB interposer and FTS assembly.

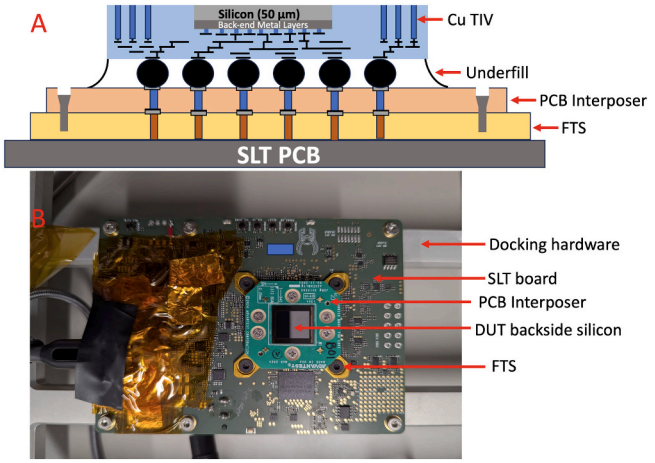

Fig. 8. (A) Cross-section diagram showing exposed DUT backside silicon and copper TIVs after grinding off the overlying DRAM and RDL layers. (B) Top view image of the polished DUT mounted on the PCB interposer and FTS.

4.2. Sample polishing and preparation

Once the device is mounted onto the PCB interposer, the DRAM and Redistribution Layers (RDL) overlying the logic die are removed. The silicon is then polished to a Residual Silicon Thickness (RST) of approximately 50 µm to facilitate imaging at the backside using OFI tools. At this thickness, the bulk silicon is exposed and is transparent to near-infrared (NIR) and mid-wave (MWIR) wavelengths of 1 to 3 µm and 3 to 5 µm, respectively. This same RST is required when performing fault isolation on ATE platforms.

It is important to note that when the silicon is polished to 50 µm, the DRAM and all overlying RDL layers connecting it to the logic die are removed. These connections must be regenerated later to interface a new DRAM positioned alongside the logic die. As illustrated in Fig. 8, this process exposes both the silicon backside and the copper Through-Interposer Vias (TIVs). Because copper is highly susceptible to rapid oxidation, a specialized sample preparation procedure was developed to protect the exposed copper contacts. This method involves coating the copper with Electroless Nickel Electroless Palladium Immersion Gold (ENEPIG), which is an advanced plating technique and the industry standard for surface finishing. ENEPIG deposits three distinct metal layers consisting of nickel, palladium and gold [9]. Aside from preventing TIV oxidation, ENEPIG offers superior corrosion resistance, enhanced high-frequency response, lower contact resistance and extended shelf life [11]. This process is illustrated in Fig. 9.

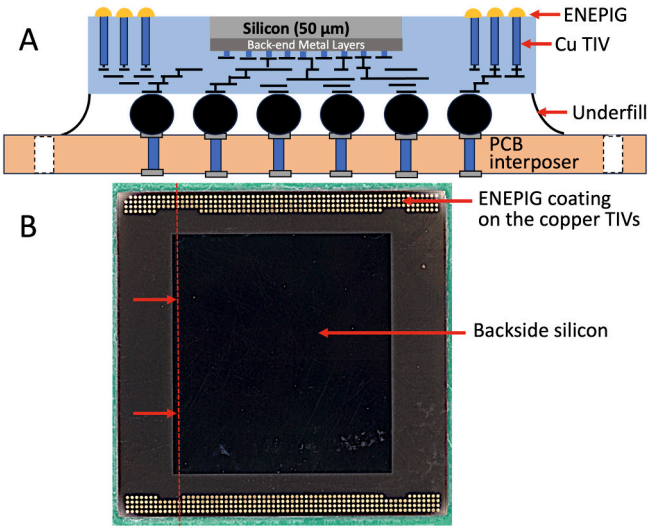

Fig. 9. (A) Cross-section diagram showing ENEPIG finish on top of the copper TIVs after polishing. (B) Image of a device subjected to ENEPIG coating.

4.3. DRAM interposer and DRAM card

An interposer is an interface that facilitates the translation from a high-density connection to a lower-density connection. This component is typically utilized to adapt connections with a fine pitch to a coarser pitch to allow for easier routing.

This novel hardware solution employs a specialized DRAM interposer (Fig. 10) as the interface between the DUT TIVs and the DRAM card. In this setup, the DUT is mounted on a PCB interposer, also referred to as a carrier or coupon board. The DRAM interposer connects the TIV contact pads of the DUT to the contact pads of the DRAM card through needle-like probe tips (Fig. 11). This hardware is unique due to its capability to contact a TIV pitch of approximately 210 µm, which is not currently known to be commercially available. The DRAM card itself holds the DRAM and serves to regenerate the RDL that were previously removed during the silicon polishing process. By utilizing this configuration, the DRAM is mounted off-center relative to the chip, providing full optical access to the logic die while maintaining system-level functionality.

One of the primary challenges in the design and manufacture of the DRAM interposer is providing a sufficient aperture for the silicon die without compromising the mechanical contact between the needle-like probes and the TIVs. Furthermore, because the TIVs are situated very close to the die edge with a very fine contact pitch (~210 µm), mechanical simulations are required to ensure consistent electrical contact while maintaining sufficient opening for full die backside imaging during OFI. Alignment pins between the DUT, the DRAM interposer, and the DRAM card must be precisely engineered to ensure full alignment across all layers. Due to the variability in the mechanical polishing and ENEPIG coating thicknesses, appropriate tolerances must be integrated into the design and the selection of the pogo pins. The dimensions and material properties of the pogo pins must also be considered for high bandwidth performance.

Fig. 12 shows the SLT application board, with the production socket replaced with the FA hardware. This hardware stack up consists of the following: (1) Flat Top Socket, which is mounted on the SLT board, (2) PCB interposer, (3) DUT mounted on the PCB interposer and polished to 50 µm, with TIVs coated with ENEPIG, (4) DRAM interposer, (5) DRAM card. It is important to consider the total stack up height of the DRAM interposer and the DRAM card, which sits on top of the DUT. As mentioned in Section 3, the hardware must be designed such that it satisfies the working distance requirements of the optical lenses. The average working distance for high magnification lenses (50× and 100×) is approximately 10 mm. Since the optical lenses need to focus on the die silicon surface, careful consideration is required to ensure that the stack

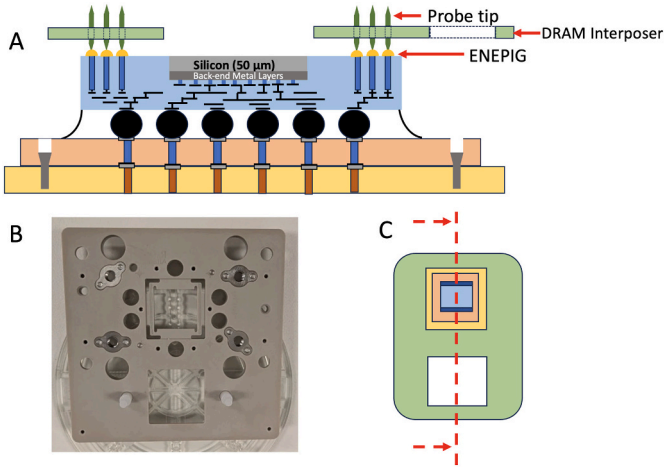

Fig. 10. (A) DRAM interposer cross-section diagram. (B) Image of a DRAM interposer. (C) Top-view diagram of the DRAM interposer.

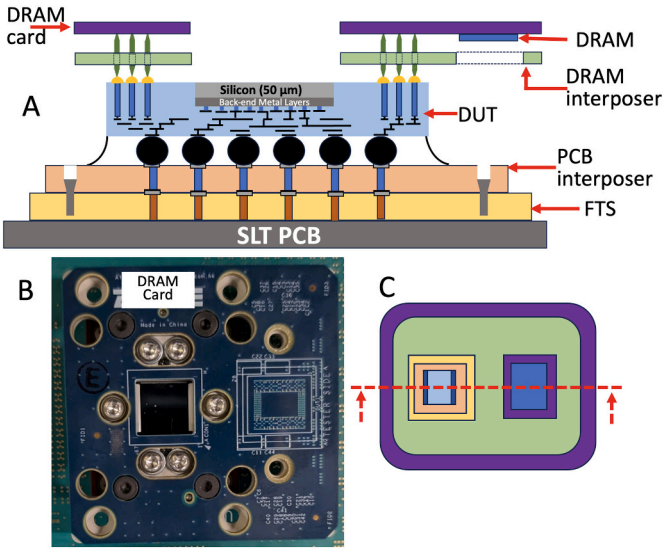

Fig. 11. (A) Cross-section diagram of an SLT FA hardware stackup. (B) Top view image showing DRAM card. (C) Top view-diagram showing the FTS, PCB interposer, DUT, DRAM interposer and DRAM card.

up of the DRAM interposer and DRAM card don't exceed the working distance of the air gap lenses. To ensure a safe working distance and accommodate potential SLT board tilt during docking, the total height of the DRAM interposer and DRAM card stack was limited to 6 mm.

As shown in Fig. 12, the DRAM card utilizes a single DRAM chip, identical to the original component used on the device. The primary challenge of this configuration is that the DRAM channels are not physically symmetrical; consequently, routing cannot be fully optimized due to the DRAM's positioning.

Originally, the DRAM was stacked directly atop the logic die, ensuring symmetrical routing across all channels. Moving the DRAM to one side of the logic die introduces signal path mismatches and increased path lengths on specific channels. While placing two DRAMs, one on each side, could resolve these issues, doing so would deviate from standard architecture and require significant additional effort from the software team to generate and validate new scripts. To minimize development overhead, the single-DRAM solution was chosen, considering a compromise in the DRAM communication bandwidth. The best practices for DRAM card design will be discussed in detail at the end of this section.

4.4. Hybrid ATE/SLT flat top socket

A primary challenge of the SLT platform is that it lacks the built-in continuity testing found in ATE systems to verify individual pin connectivity. This absence makes it difficult to determine whether a device failure is caused by functional defects or simple contact issues. Furthermore, the process of solder-mounting a device to the SLT carrier board can introduce solder defects that may lead to catastrophic failures unrelated to the actual failures. Because SLT sockets usually have smaller footprints and more stringent requirements than ATE hardware, these carrier boards cannot be tested using standard ATE flat top sockets. To resolve this, a hybrid flat top socket was designed to allow for the ATE testing of the SLT carrier boards. As shown in Fig. 13, this design features a top FTS interface compatible with the SLT carrier board and a bottom interface that mates directly with the ATE loadboard, enabling continuity checks on the board assembly. Using the same socket, the DRAM interposer and DRAM card can be attached to perform ATE-based DRAM testing, further increasing the confidence level in the hardware setup for the SLT platform.

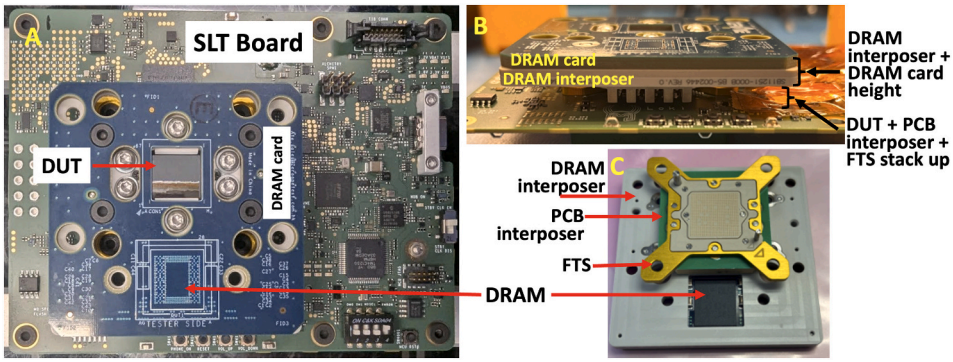

Fig. 12. (A) SLT FA hardware top view (B) SLT FA hardware side view (C) SLT FA hardware bottom view.

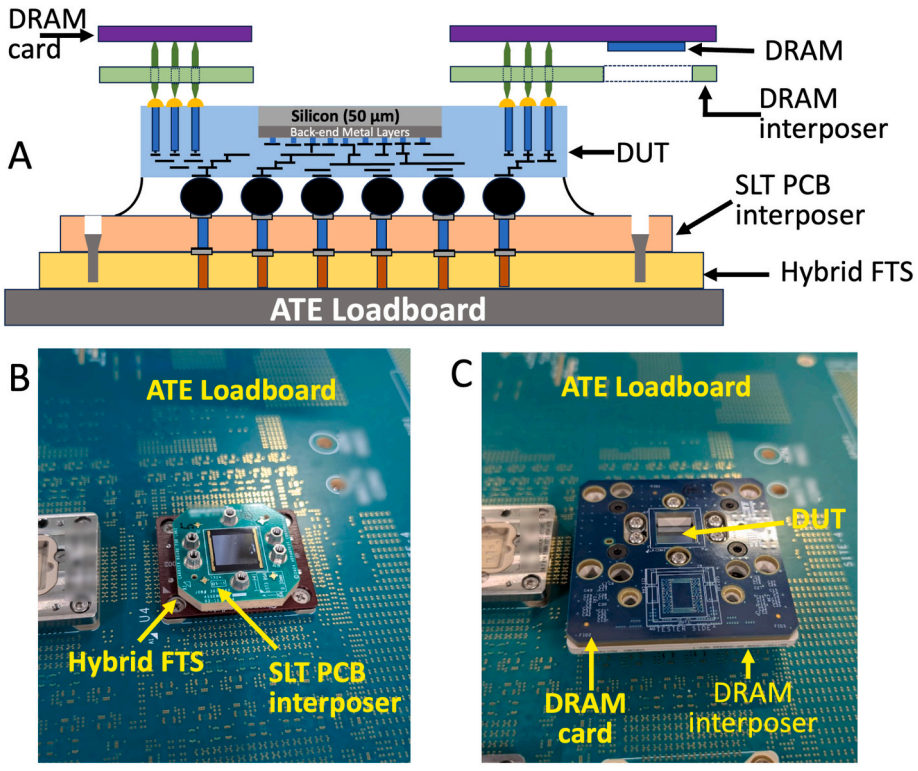

Fig. 13. (A) Cross-section diagram of the Hybrid FTS interfacing between the SLT PCB interposer and the ATE loadboard for continuity testing. (B) Top view image showing the Hybrid FTS mounted on an ATE loadboard to test the SLT PCB interposer with DUT. (C) Full FA hardware stack using the hybrid FTS to interface to the ATE and run DDR tests.

4.5. Best practices for DRAM card design

Designing a high performance DRAM card requires special attention to Signal Integrity (SI) and Power Integrity (PI) to ensure robust data transfer and stable voltage rails. The following best practices detail the main considerations for stackup, routing and power distribution.

1. PCB Stackup and Material Selection

   A critical part of a high speed interposer lies in its substrate composition.
   - Material Properties: Select high-frequency laminates characterized by a low Dielectric Constant (Dk) and a low Loss Tangent (Df). These properties minimize propagation delay and signal attenuation which are critical for maintaining eye-opening at high data rates.
   - Layer Composition: The stackup must feature an adequate layer count to provide enough layers for high speed signals with a continuous reference plane as well as dedicated layers for voltage rails.
   - Routing Topology: High-speed signals should utilize stripline routing rather than microstrip. Stripline provides a more uniform dielectric environment, better impedance control and superior electromagnetic interference (EMI) shielding.
   - Via Selection: In order to mitigate the parasitic effects of via stubs, the design should prioritize blind and buried vias.
2. High-Speed Signal Routing
   - Impedance Control: Calculate trace widths and spacing specifically to meet the target characteristic impedance, typically 40 to 50 Ω SE and 80 to 100 Ω differential for DDR interfaces.
   - Continuous reference: Ensure all high speed traces are routed over solid, uninterrupted reference planes. When a layer transition is unavoidable, place a GND stitching via immediately close to the signal via transition to provide a low impedance return path and minimize crosstalk.

- Crosstalk Mitigation: To minimize coupling maintain a spacing of 2 to 3× the height to the reference plane.
- Skew: Follow the length matching requirements, reducing skew between clock and data lines is critical for maximizing timing margins during interface training.

3. Power Distribution Network (PDN)

A robust PDN ensures that the DRAM receives clean power with minimal fluctuations.

- Utilize dedicated power planes for critical rails to handle high transient currents while meeting the low ripple requirements.
- Use enough amount of vias for power and ground connection to minimize DC drop and reduce loop inductance.
- Place decoupling capacitors as close to DRAM bumps to improve power delivery during high-speed switching.

5. Results

This section evaluates LPDDR5x performance metrics captured via the novel FA hardware solution, demonstrating its efficacy across diverse system-level platforms. The focus of this section is to highlight the varying complexities of test methodologies that can be performed using this novel hardware to leverage the full suite of FA tools like Photon Emissions, PEM etc. The below test classification framework follows a bottom-up hierarchy that transitions from hardware-centric validation at the bare metal layer to concurrency and boot-flow testing in the Little Kernel, eventually scaling to high-level service and application-layer verification in the Android environment. While traditional LPDDR failure analysis often relies on ATE-based platforms or augmented testers [12], this study showcases real-world system test cases that leverage Android-based mission-mode applications for performance qualification and benchmarking.

Table 1 showcases the correlation between increasing data rates and software stack complexity, where higher frequencies reveal instabilities that lower-level hardware training might initially mask. The novel FA hardware solution (shown in Fig. 12) is fully validated for all layers of SW complexities at 820 Mbps.

For performance qualification and benchmarking at speeds exceeding 1076 Mbps, the platform is currently limited to low-level kernel and training functions; at these higher frequencies, complex system-level stresses expose underlying signal integrity bottlenecks that prevent stable operation of the full Android stack. Since Android boot represents the most rigorous system stress test—and given that the majority of power sequencing, calibration, and training protocols occur at the 820 Mbps boot frequency—this provides an ideal baseline for tailoring specialized FA hardware solutions toward critical system debugging.

To evaluate platform stability at performance limit, “stressapptest” was deployed in an aggressive, multi-threaded configuration (-A -M 1024 -m 4 -i 2 -f ./neon -s 60) to simulate intensive mission-mode conditions. This methodology utilized a 1024 MB memory allocation distributed across four memory-copy and two pattern-inversion threads, alongside a concurrent disk I/O thread, to maximize pressure on the memory controller and I/O bus. By standardizing these 60-s stress iterations, the study effectively identified the specific data rate thresholds where high-speed signaling limitations lead to system-level failures.

Fig. 14A & B shows the on chip eye monitor based eye diagrams at 820Mbps for 8 DQ bits & 1 DMI per channel across all four PHY channels. The horizontal axis represents the calibrated WR/RD delay in steps, while the vertical axis tracks VREF offsets. Despite the extended trace length inherent to the DRAM path, the plots reveal robust signal integrity with wide open eyes and healthy voltage/timing margins, confirming stable data sampling at the boot frequency. The eye height and width values were purposely obscured due to proprietary reasons and ensured that eye margins comply with system requirements and JEDEC Standards, preventing false test failures induced by setup.

The progressive degradation of signal integrity is illustrated in Fig. 15, which compares the per-bit eye diagrams for Rank0 Byte0 Bit8 across the 820 Mbps, 1076 Mbps, and 1344 Mbps data rates. As the eye opening narrows at higher frequencies, the diminishing timing and voltage margins directly correlate with the failures observed in the high-level software stress tests detailed in Table 1. Consequently, these findings highlight a critical opportunity for future platform design to enhance signal margins and hardware-software synchronization, ultimately enabling full-stack functionality at higher data rates.

One critical point to note is that since the SLT FA hardware can run the Android “stressapptest” at the minimum data rate of 820Mbps, this means that it can run any test, with the caveat that the device may not be stressed enough due to limitations in the DDR frequency. To ensure that the developed interposer does not alter or mask the device fail behavior, an SLT FA debug workflow was developed and described in Fig. 19, Section 7.

6. FA case studies

This section will present some of the FA case studies which utilized this novel FA hardware solution for system-level platforms, with the goal of identifying defects or design issues in the bottom SoC die, package interconnect and integration issues with an external DRAM or software issues. It is important to note that most of the SoC bottom die failures could be screened by standard ATE-based ATPG test vectors. However, such tests cannot detect 100% of failures, and hence some of these defects get screened out as SLT rejects. This percentage is very low, and

Table 1
LPDDR5x stability matrix across system complexity layers and data rates.

Data Rate (Mbps)	DDR Training	Bare Metal	LK DDR	LK Memtester	Android Stressapp
820	✓	✓	✓	✓	✓
1076	✓	✓	✓	✗	✗
1344	✓	Reduced Eye	✓	✗	✗
1688	✓	Shmoo hole	✓	✗	✗
1996	✓	Shmoo hole	✓	✗	✗
6296	✓	✗	✗	✗	✗

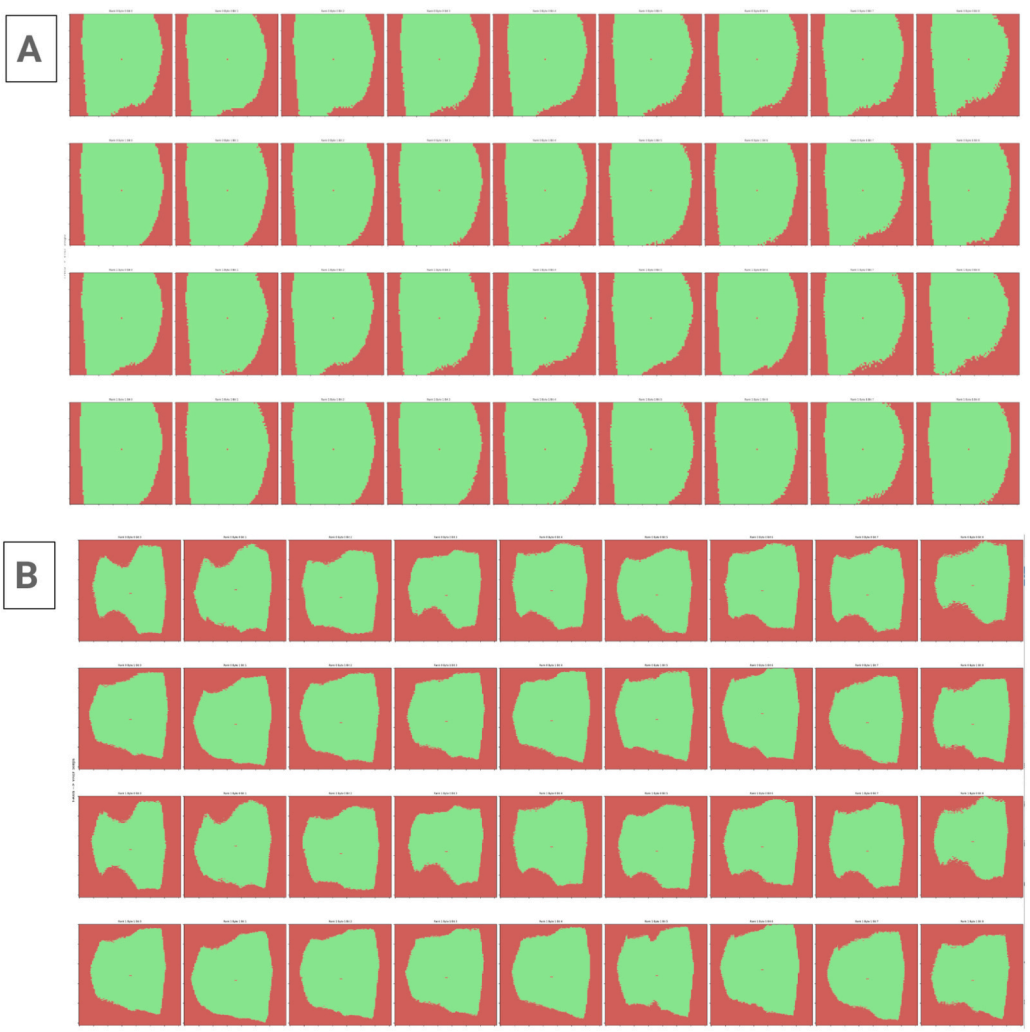

Fig. 14. Eye diagrams using bare metal tests - (A) 820 Mbps per bit Write DQ Eye for 4 channels of a PHY. (B) 820 Mbps per bit Read DQ Eye for 4 channels of a PHY.

hence, most of the SLT failures encountered were focused on the mission mode, which involves the expected operation of the SoC but in certain voltage and frequency conditions. For the most part, the failures we have encountered involved unexpected leakage or thermal dissipation during normal operation. The focus of this section is to highlight the key capabilities of the novel SLT FA hardware, in addition to already available ATE-based FA solutions. The details of the issue will not be discussed; Rather, the paper will focus on the methodologies that enabled root cause analysis on the system level.

6.1. DDR PHY leakage analysis

In this case study, the goal was to identify if there were blocks in the DDR interface that were inadvertently turned on and causing leakage using a specific test sequence. This analysis was performed on the SLT application board, using a Flat Top Socket where the PCB interposer containing the DUT was mounted to. The DUT was polished to 50 µm to facilitate Photon Emissions (PEM) from the backside of the chip. An image of the SLT setup taken from within the OFI tool is shown in Fig. 16A. The PEM results in Fig. 16B eventually proved that the test sequence is correct and the emitting transistors point to circuitries that were expected to leak current according to the Vt (transistor threshold voltage) flavor.

6.2. Thermal analysis and mapping

Aside from identifying defects, the OFI tool such as the Thermal Camera can be used to characterize the thermal performance of devices and test sequences. Fig. 17A shows an SLT application board using the novel FA hardware which consists of the flat top socket, the PCB interposer (carrier/coupon board), polished DUT with ENEPIG coating, DRAM interposer and DRAM card. Fig. 17B and C shows the different hot spots that were detected using different sequences, which provide an idea of which blocks are activated and drawing heat during the test sequence. Finally, Fig. 17D shows the thermal profile which can be obtained from any location within the die, since the entire die is fully accessible through the aperture in the DRAM interposer and DRAM card.

6.3. Quiescent and standby current measurements

IDDQ test is the gold standard for checking the quiescent current of the device using the ATE. However, to assess system-level performance, a similar test may be run on the SLT platform, and using the Photon Emission tool, the location of the leakage may be identified and correlated with different IP blocks. Fig. 18A shows the PEM setup on the SLT FA hardware. The PEM images in Fig. 18B point to an IP block that was expected to be ON during the test. No other IP blocks in the region of interest were found to emit, confirming that the test setup was good.

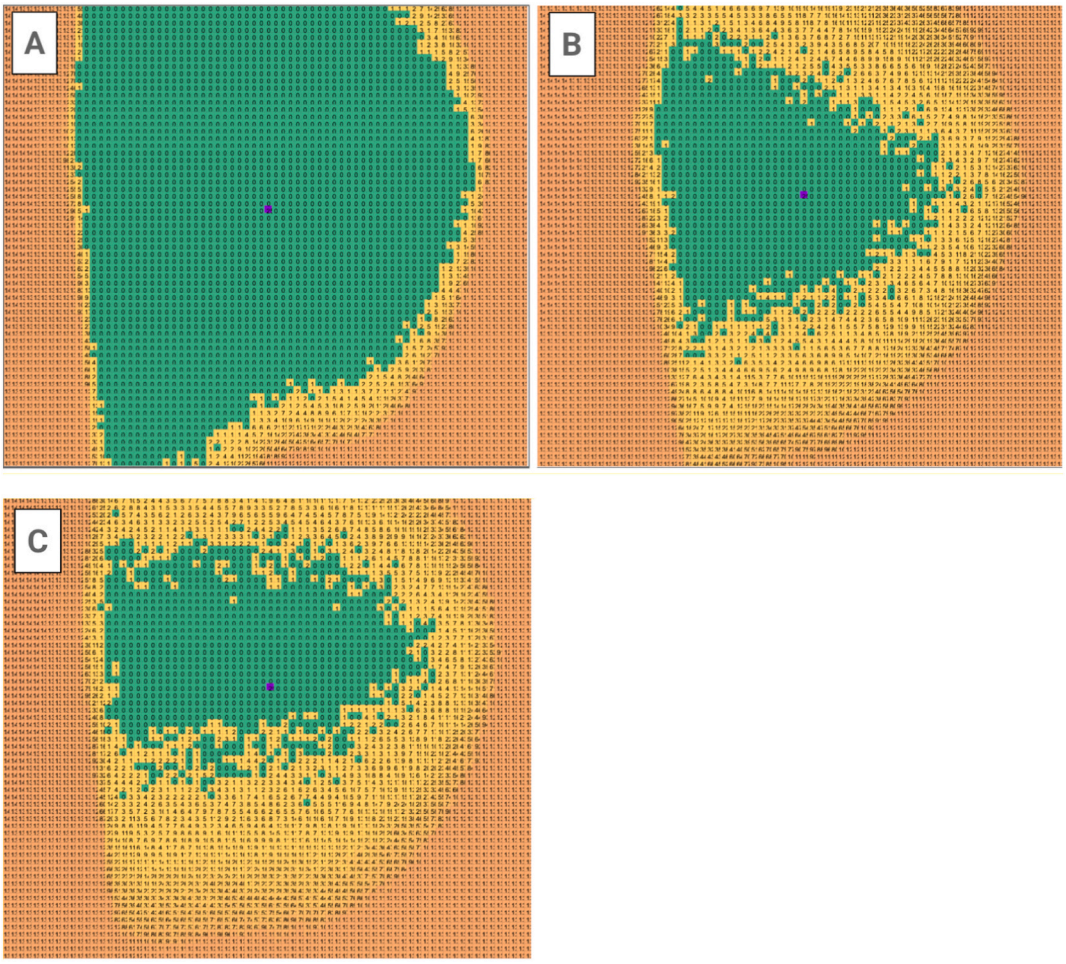

Fig. 15. Rank 0 Byte0 Bit8 eye margin comparisons for data rates, (A) 820 Mbps (B) 1076 Mbps and (C) 1344 Mbps.

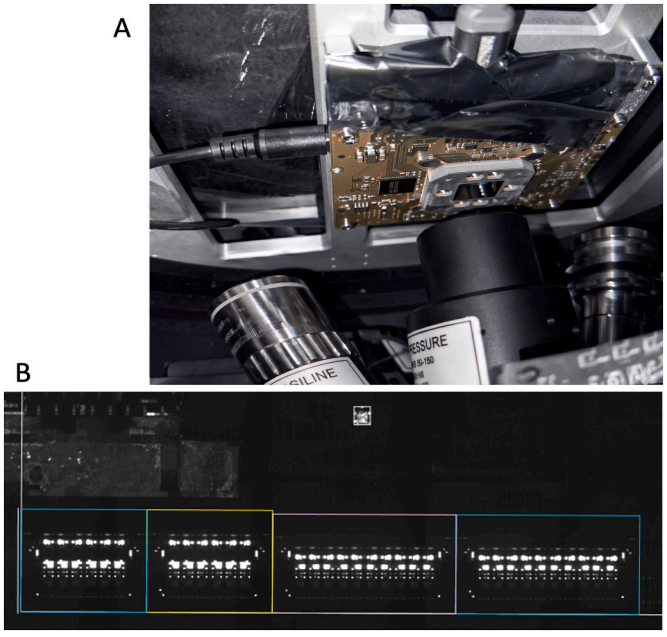

Fig. 16. (A) SLT FA hardware setup. (B) Photon emission results on the DDR interphase.

7. Limitations and future work

Due to the thickness of the DRAM interposer and DRAM card stack, the optical imaging of the die backside using a Solid Immersion Lens (which requires direct-to-silicon contact) is very limited. This implies that most of the optical imaging will be done using air gap lenses, due to the depth of the aperture (approximately 6 mm). While most of the OFI techniques are still applicable using air gap lenses, the main limitation will be on Laser Voltage Probing and Imaging (LVX) techniques, which requires the SIL lens for operation. Nevertheless, a lot of valuable information can still be obtained using Photon Emissions, Thermal Imaging/Mapping and Dynamic Laser Stimulation.

Another limitation of this hardware, as earlier mentioned, is on the DDR data rate, which was caused by the asymmetrical placement of the DDR to the side of the logic die. The goal of the hardware is to identify failures with blocks like CPU, GPU other than DDR which can be run with limited DDR speed. Typically all the tests can be run at their maximum operating speed, it is the bandwidth that gets limited due to lower DDR speed and prevents us from running industry standards tests like Geekbench which requires DDR also to operate at Max speed. The DRAM plays a critical role in the overall SoC operation, optimizing performance and latency through high speed DDR transactions. The DDR holds the data for the next set of instructions that are being executed by the processors. The previous sections showed that the SLT FA interposer can successfully run a full Android “stressapptest” at a lower speed of 820 Mbps, which means that it can execute any functional test at slower DDR speeds. The issue here is not which tests can be run, but rather, if the failures can be replicated and if the hardware

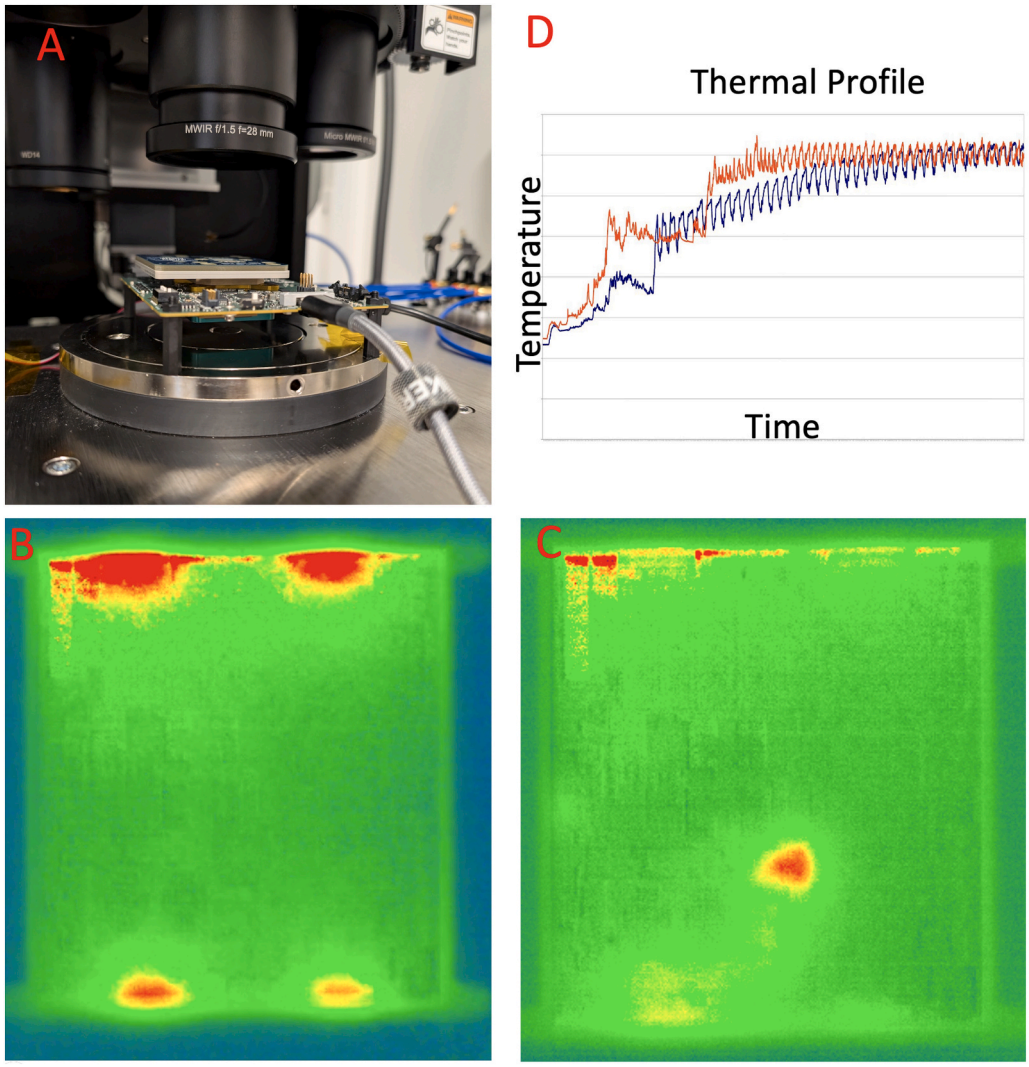

Fig. 17. (A) Thermal measurement setup using SLT application board. (B) Thermal mapping showing hotspots on DDR interface when DRAM test is run. (C) Thermal hot spot observed on expected IP block. (D) Thermal profile ran on a certain location in the chip using 2 different test sequences (red and blue).

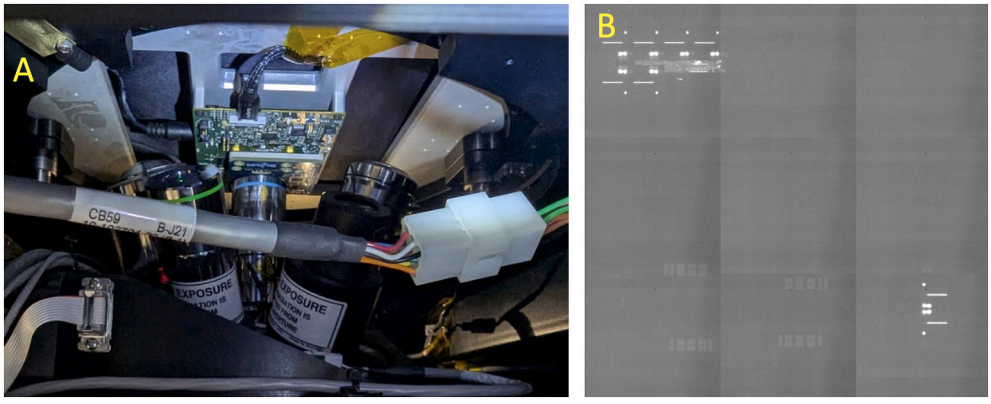

Fig. 18. (A) SLT FA hardware setup on the PEM tool. (B) PEM results showing leakage on an IP block which was expected to be ON during the test.

would introduce other failures as a result of the asymmetrical placement of the DDR. This hardware is effective in identifying hard SLT failures (that fails across all voltage and frequency) caused by defects that were not effectively screened by the ATPG tests (test escapes). However, the limitation of this hardware is when the failures occur as a result of a performance degradation, where the device has to be stressed to the maximum voltage or frequency condition, or using very heavy workloads. As the DDR runs slower, it is not certain whether the test conditions at reduced frequencies will stress the device hard enough for the failures to be replicated. Furthermore, in some cases, some tests may be

fine running slower but encounter random behavior especially during performance and concurrency use cases due to contention in cross-domain traffic (e.g., CPU-DDR-GPU operations, etc.). This limitation in the hardware can introduce new failures, which are unwanted. Without thorough investigation, it may be challenging to discern the actual failure from the collateral effects of the FA hardware.

To ensure that the developed interposer does not alter or mask the device fail behavior, an SLT FA debug workflow was developed and described in Fig. 19. In the proposed workflow, it is important to establish a baseline by first identifying the failing voltage and frequency for the device. Using that information, an FA debug script is generated with reduced DDR speeds. As a sanity check, a reference device is first run to evaluate this FA debug script. This sanity checked is performed on both fully packaged and sample-prepped reference devices. Once this FA script passes, it is then used to validate the failure initially observed in the reject device. It is very important to do this step prior to sample preparation to establish a baseline. Once data is captured on the failing device, this device can now be sample-prepped. Post sample preparation checks are done to validate if the failure still exists, prior to any EFI work. Using this workflow and by documenting the results in every step, any failures introduced by the FA hardware can be effectively identified and discerned from the original failure.

For potential enhancements, one might consider reducing the DRAM interposer and DRAM card stack up. Based on our experience, this has been specifically challenging due to the stringent routing requirements. Furthermore, some optimizations in the screw placement may be

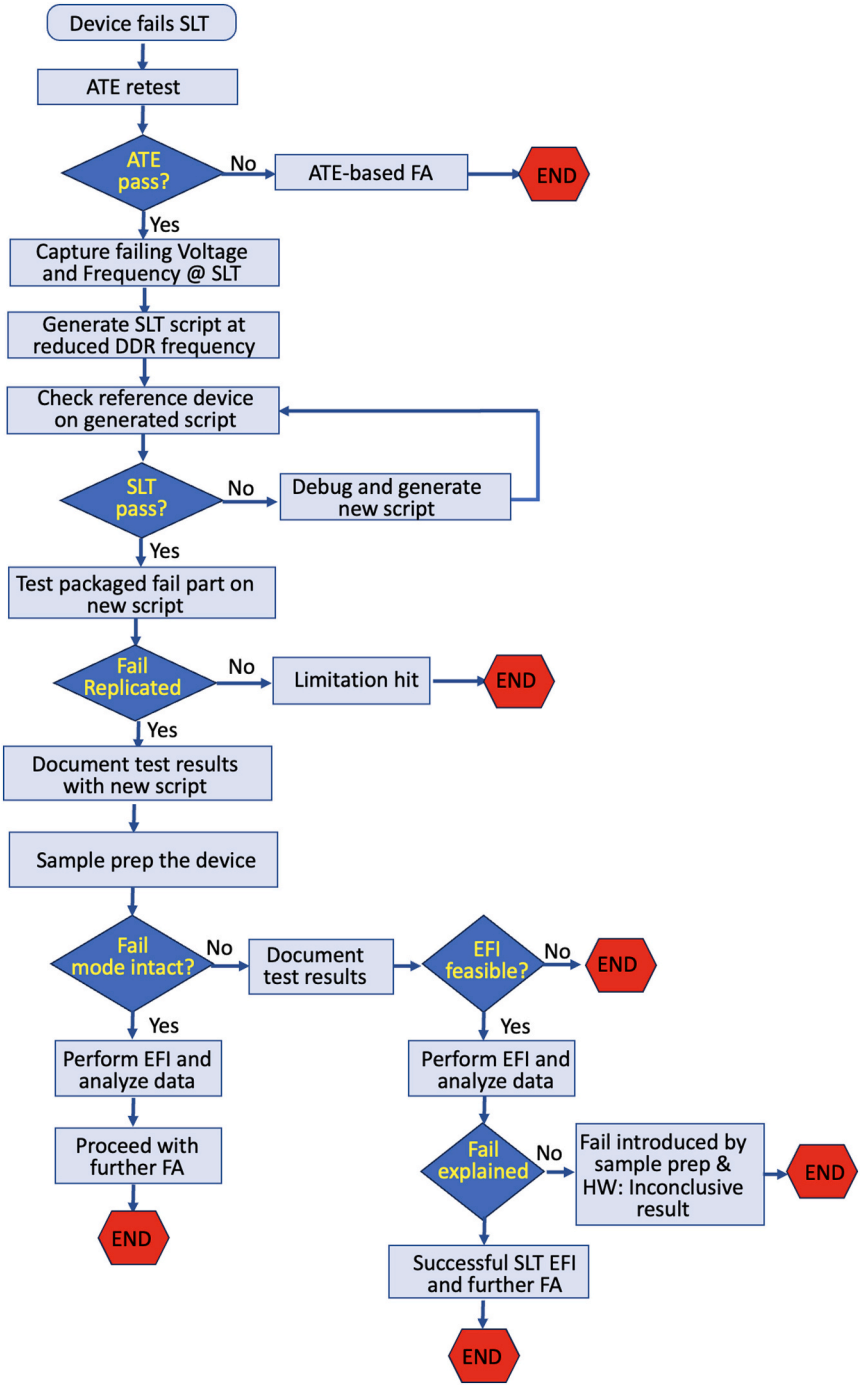

Fig. 19. Workflow for SLT-based electrical fault isolation.

considered, in favor of a more optimized routing. Another potential improvement may utilize two separate DRAMs, each situated on both sides of the logic die for optimal routing. The only caveat is that this would require additional software development and it may be difficult to compare and correlate failures from an actual device with a single DRAM.

8. Conclusion

This paper presented a novel FA hardware design that enables Optical Fault Isolation at the system-level platform. Different components of our groundbreaking SLT FA hardware are presented, together with the caveats and the limitations. This paper also presented the best practices for DRAM card design and optimization. Our first characterization results on the first prototype were promising – achieving full functionality up to the Android application. There were some limitations with the DRAM operating frequency, but since the primary goal is to enable OFI on the bottom logic die, majority of the tests can still be run at slower DDR data rates, making this hardware solution a significant breakthrough in 3D IC debug of SoCs in system level platforms.

CRediT authorship contribution statement

Lesly Endrinal: Conceptualization, research and review of materials/documents, writing of the paper
Willem van Driel: Adviser, suggestions on the contents, review of the paper
Guo Qi Zhang: Adviser, suggestions on the contents, review of the paper
Jehan Saujauddin: Research and providing data
YuanChung Ho: Research and providing data
Dilbagh Singh: Research and providing data
Sasi Sekaran Sundaresan: Research and providing data, writing report
Vinod Kumar Kakumanu: Research and providing data, writing report
Jessen Gonzalez: Research and providing data, writing report.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

The authors would like to thank Sajjad Pagarkar and David Greenlaw for their patience in waiting for this hardware to finally be developed, and for giving us the opportunity to write this paper. We also would like to thank the Advantest Team (Robert Lee, Daniel Lam, Gianluca Lombardi, Cymeyna Mcfarland, Kalyani Dosala, Shaan Patel). Lastly, we would like to acknowledge the support of the PDTE PE and SLT teams (Jagadeesh Gownipalli, Mrudul Sharma, Raghunandan Nagesh, Shakiruddin Mohammed, Aditya Valdas, Chandrasekhar Pullela, Parth Patel, Simrit Saroa), Package Design Team (Terence Cheung, CK Kim, Ruby Lai, Tian Qiu, Robin Lien), Rama Thotakura, Mirko Damiani and Nicolas Gagnon.

Data availability

No data was used for the research described in the article.

References

[1] O. Burkacky, J. Dragon, N. Lehmann, The Semiconductor Decade: A Trillion Dollar Industry, McKinsey&Company, April 2022.
[2] PWC, Global semiconductor industry outlook 2026. https://www.pwc.com/gx/en/industries/technology/semiconductor-and-beyond.html, 2025.
[3] JEDEC Standard (JESD209-5), Low Power Data Rate 5 (LPDDR5), Feb 2019.
[4] Biwin Storage Technology, LPDDR6 vs LPDDR5: what's the difference?. https://www.biwintechnology.com/lpddr5-vs-lpddr6/#:~:text=Its%20base%20speed%20starts%20at,economic%20performance%20from%20the%20systems, Jan 2026.
[5] S. Knowlton, LPDDR6 vs. LPDDR5 and LPDDR5X: what's the difference?. https://www.synopsys.com/blogs/chip-design/lpddr6-vs-lpddr5x-lpddr5-differences.html, Feb 2026.
[6] JEDEC Standard (JESD209-5B), Low Power Data Rate 5 (LPDDR5X), June 2021.
[7] JEDEC Standard (JESD209-6), Low Power Data Rate 5 (LPDDR5X), June 2021.
[8] L. Endrinal, A practical tutorial on ATE-based electrical fault isolation of digital SoCs using photon emission and laser voltage imaging/probing, in: 2025 International Symposium on Testing and Failure Analysis. Pasadena, California, Nov 2025.
[9] M. Sufyan, ENEPIG: the gold standard in surface finishing technology. https://www.wevolver.com/article/enepig-the-gold-standard-in-surface-finishing-technology, Sep 2024.
[10] Robson Technologies Inc, Flat top socket convertible. https://www.testfixtures.com/wp-content/uploads/2021/05/RTI_FTSCtestSockets-2.pdf, 2021.
[11] Wonderful PCB, Which PCB surface finish is better: ENIG or ENIPIG?. https://www.wonderfulpcb.com/blog/enig-vs-enepig-pcb-surface-finish/#:~:text=ENEPIG%20stands%20for%20Electroless%20Nickel,copper%20is%20prepared%20for%20plating, 2024.
[12] C.-H. Kye, et al., A 10Gb/s/pin DQS and WCK built-out tester for LPDDR5 DRAM test, in: 2022 IEEE Asian Solid-State Circuits Conference (A-SSCC), Taipei, Taiwan, 2022, pp. 6–8, https://doi.org/10.1109/A-SSCC56115.2022.9980646.