Evaluating Vulnerability of Chiplet-Based Systems to Contactless Probing Techniques

Aleksa Deric*, Kyle Mitard†, Shahin Tajik‡ and Daniel Holcomb§
*§Department of Electrical and Computer Engineering
University of Massachusetts Amherst, Amherst, Massachusetts 01003
Email: aderic@umass.edu* dholcomb@umass.edu§
†‡Department of Electrical and Computer Engineering
Worcester Polytechnic Institute, Worcester, Massachusetts 01609
Email: kmitard@wpi.edu† stajik@wpi.edu‡

2024 IEEE International Test Conference (ITC) | 979-8-3315-2013-7/24/$31.00 ©2024 IEEE | DOI: 10.1109/ITC51657.2024.00022

***Abstract*—Driven by a need for ever-increasing chip performance, a growing number of semiconductor companies are opting for all-inclusive System-on-Chip (SoC) architectures. Increasingly, the solution adopted to minimize the impact of silicon defects on manufacturing yield of larger dies has been to split a design into multiple smaller dies called chiplets, which are then brought together on a silicon interposer. Advanced 2.5D and 3D packaging techniques that enable this kind of integration also promise increased power efficiency and opportunities for heterogeneous integration.**

However, despite their advantages, chiplets are not without issues. Disaggregating a design into multiple separate dies introduces new security threats, including the possibility of tampering with and probing exposed data lines. In this paper we evaluate the exposure of chiplets to probing by applying laser contactless probing techniques to a chiplet-based AMD/Xilinx VU9P FPGA. First, we identify and map interposer wire drivers, and show that probing them is easier compared to probing internal nodes. Lastly, we demonstrate that delay-based sensors, which can be used to protect against physical probes, are insufficient to protect against laser probing.

I. Introduction

Chiplets are separately produced silicon dies that are assembled on an interposer to form a composite system that is comparable to a monolithic integrated circuit. They allow designers to circumvent reticle limits and provide an opportunity to increase yield, leverage multiple process nodes for heterogeneous integration, and reuse IP. Driven by demand for CPUs and GPUs in data centers, by 2027 the market for chiplet-based processors is expected to reach $135 billion [1]. Currently, AMD [2], Intel [3], Nvidia [4] and Apple [5] all feature chiplet-based products in their lineup, which have initially proliferated in top of the line devices that can absorb the added overhead costs.

Unfortunately, disaggregating a system into parts opens many new attack vectors, including possibility of cross-die side-channels, introduction of hardware Trojans in the supply chain, IP piracy, die swapping, and probing of chiplet interfaces. Among these, probing is of particular concern, as signals carried on the interposer wires are likely to be important data channels that are of high-value to an attacker.

To the best of our knowledge, this paper is the first work to demonstrate contactless probing techniques in the context of chiplets. The specific contributions of this paper are:

- We experimentally demonstrate the first laser probing attack on chiplets, which are fabricated in 16nm technology and packaged with a silicon interposer in 65nm technology.
- We compare the exposure of inter-chiplet wires relative to that of on-die wires in the same technology.
- We assess the effectiveness of delay-based sensors in detecting contactless probing attacks against chiplets.

II. Background

In this section, we discuss the defining features of chiplet interconnects and cover relevant previous work in the areas of chiplet security and failure analysis. Our work covers a unique gap pertaining to probing of chiplet interfaces.

A. Chiplet Interconnect Technology

Signals between chiplets travel a distance of only a few millimeters through microbumped connections and densely packed wires of a multi-layer interposer, as depicted in Figure 2a. The wires in die-to-die connections, although much smaller than wires used in board-level connections, have more capacitance than local on-die wires routed in lower metal layers, and therefore are driven with upsized transistors; probing these upsized transistors is the focus of this paper.

Intel [6], TSMC [7], and AMD/Xilinx [8] all use their own chiplet interfaces, which provide capability for sub-pJ energy-per-bit and a 35μm to 55μm bump pitch with hundreds of wires per millimeter of shoreline to deliver hundreds of GB/s of aggregate bandwidth. Other shared features include low-resistance microbump connections to the interposer, shielded wiring and source-synchronous clocking.

B. Related Work in Chiplet Security

With progress being made toward addressing reliability, performance and testing challenges [9], security of chiplets has been left as an afterthought. Papers have surveyed and raised awareness of possible security issues [10], yet there are few concrete examples of real word attacks on chiplets. [11] proposes a reverse engineering framework for chiplet packages and uses 3D X-ray tomography to extract information about the package architecture. Similarly, [12] discusses a SAT-based attack to reverse engineer missing connections of 2.5D split manufacturing netlists. Finally, [13] show that an inter-chiplet covert channel can be created through a shared power distribution network (PDN) that many chiplets have. Still, many of the attack vectors remain theoretical and more work

2378-2250/24/$31.00 ©2024 IEEE
DOI 10.1109/ITC51657.2024.00022

is needed to validate the threat models discussed in literature. Importantly, no works have considered the problem of laser probing attacks which is the focus of this paper.

C. Failure Analysis Applied to Security

In order to refine process parameters, advanced imaging and probing techniques have been developed. Less destructive approaches are desirable when the device-under-test (DUT) needs to be running during debugging. Near-infrared (NIR) wavelength photons can pass through silicon and are thus used to non-destructively peek through the back side of dies. The signal-to-noise ratios (SNR) of optical and e-beam probing are compared in recent work [14]. Research has also considered how to use optical probing to locate signals of interest in large devices [15]. More critically, [16] demonstrated how optical contactless techniques can be used to bypass on-die encryption and extract a full plaintext bitstream from a 28nm AMD/Xilinx FPGA. Similarly, [17] showed how thermal laser stimulation (TLS) can be used to read out secret keys from inside a battery-backed RAM (BBRAM) - the process which can even be fully automated [15]. Nonetheless, few works have employed probing against a 16nm technology, and none have evaluated probing against inter-chiplet connections, which is what our work achieves.

D. Threat Model

Our threat model assumes that the attacker has physical possession of a multi-die flip-chip device from which they wish to extract secret data. The attacker rents a contactless probing system from a failure analysis lab, with the rest of their arsenal consisting of off-the-shelf tools and commodity hardware. A chiplet system is likely to be partitioned based on functionality of individual modules. Therefore, we assume the attacker probes communication links between the chiplets.

III. Contactless Probing Techniques

In this section, we provide an overview of the contactless probing techniques employed in the paper.

A. Photon Emission

Measuring photon emission of a chip provides coarse information about the activity level of its circuits. When a complementary metal-oxide semiconductor (CMOS) gate is quiescent, that is none of its nodes are switching, there are only small leakage currents flowing through its nMOS and pMOS transistors, and hence almost no photons being emitted. However, when an input to the gate changes, large currents flow to charge or discharge individual node capacitances. The energy of accelerated hot carriers traveling through transistor channels is sometimes released as photons, which travel though the silicon backside of the integrated circuit (IC) and can be captured to construct an activity map of the device.

B. Electro-Optical Frequency Mapping (EOFM)

Electro-Optical Frequency Mapping (EOFM) is a laser probing technique in which a set of precisely controlled mirrors scan a laser beam across a device and build a detailed activity map of nodes switching. More specifically, after leaving the light source, the laser beam travels through the silicon and is reflected off the metal structures on the silicon surface. While scanning, the laser is held at each position for a period of time, and the returning beam is fed into a spectrum analyzer with a bandpass filter. Circuit nodes that switch at a particular frequency will modulate the reflected light at their switching frequency, which will then be processed by the spectrum analyzer. The measurements collected during the entire scan create a frequency-selective activity map.

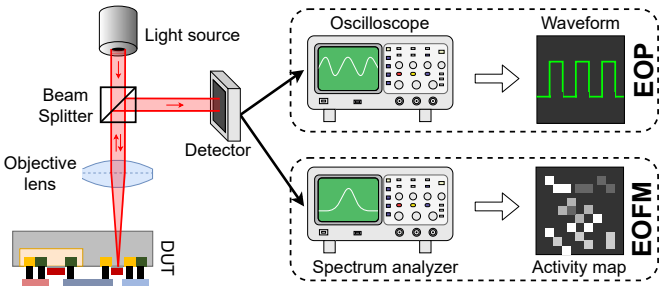

Fig. 1. Diagram illustrating the basic working principles of EOP and EOFM, as described in Sections III-C and III-B. Apart from EOFM involving scanning the beam, the two techniques differ in how the reflected light is processed.

C. Electro-Optical Probing (EOP)

Electro-Optical Probing (EOP) is a Laser Voltage Probing (LVP) technique that works by focusing light on a single point of interest and measuring the intensity of the reflected beam over time. The degree to which the returning light's intensity is altered depends on carrier concentrations in the area of interest i.e., electrical signal at the targeted spot. Hence by feeding the returning beam into a photodetector and measuring its intensity, the voltage level of the node being probed can be determined. Multiple measurement repetitions must be collected and averaged to obtain a signal with an acceptable signal-to-noise ratio. To align the repeated measurements, a trigger reference synchronized to the target waveform must be provided. Figure 1 illustrates EOFM and EOP.

IV. Experimental Setup

In this section, we introduce our experimental setup, including the chiplet-based FPGA that is our target, the failure analysis microscope used, and the board preparation process.

A. Chiplet FPGA Platform

We conduct experiments using the VCU118 development board, which features a VU9P FPGA (xcvu9p-flga2104-2L-e). The VU9P is a high-capacity UltraScale+ FPGA consisting of three identical 16nm chiplets sitting on a 65nm interposer in a flip-chip configuration with 17,280 wires running between neighboring chiplets. The wires begin and terminate with special Laguna registers, organized into groupings of six to make a Laguna site, four of which comprise a Laguna tile.

B. Board Preparation

Optical probing requires direct access to the backside silicon of the FPGA, so the board must be prepared to expose this surface. Before probing we remove the large copper heatsink and then carefully break the Integrated Heat Spreader (IHS)

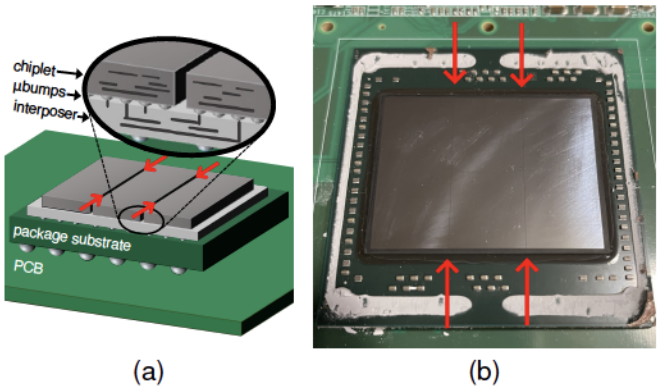

Fig. 2. (a) Diagram of chiplet-to-chiplet connections through interposer, and (b) corresponding photograph of delidded AMD/Xilinx chiplet-based VU9P FPGA used in evaluation. The red arrows in each picture annotate the chiplet boundaries.

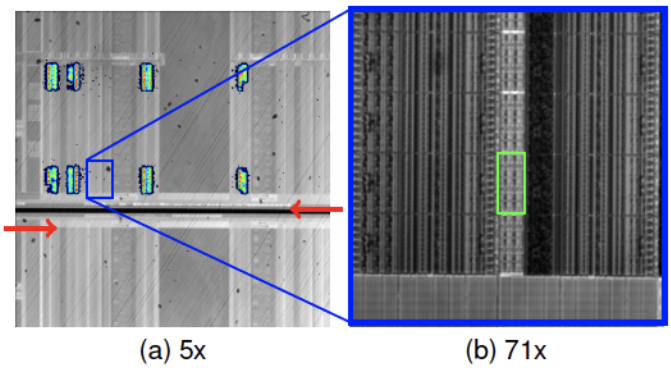

Fig. 3. Photon emission highlights circuits with high switching activity, such as groups of ring oscillators in (a), which are used as guideposts to locate circuits of interests like a Laguna column in (b). A single Laguna tile is outlined in green. The red arrows annotate the boundary between two chiplets.

loose using a chisel. Finally, we remove a thin white coat of thermal paste covering the chiplet dies by wiping it away with a 99.9% isopropyl alcohol solution. The chiplet backsides revealed by this preparation are shown in Figure 2b.

C. Emission and Laser Scanning Microscope

Probing experiments are performed using the Hamamatsu PHEMOS-X Emission and Laser Scanning Microscope. The microscope features an InGaAs photon emission camera and a 1.3μm wavelength laser for optical probing [18]. A set of galvanometric mirrors guides the laser while the reflected light is collected using one of the four objective lenses available: 5x/0.14 NA, 20x/0.4 NA, 50x/0.76 NA and 71x/0.86 NA [18].

V. Laser Probing of Interconnect Drivers

The large size of the transistors that drive die-to-die connections makes them an easy target for probing. These transistors are upsized to supply the current necessary to drive the load of the microbump and the long interposer wire. In the following section, we locate a group of drivers and compare their probing exposure to the equivalent logic function implemented using on-chip fabric registers.

A. Localizing Circuits-of-Interest

1) Using photon emission to navigate the chip: The assembled silicon package measures 35mm×26mm, making navigating to find features of interest challenging. For this purpose, we flash the FPGA with a bitstream with individually controllable blocks of oscillators custom placed around the chip to serve as guideposts for learning the physical position of on-die circuits-of-interest. Figure 3a shows a photon emission view of active ring oscillators. We steer the camera to the edge of the chip and to the vicinity of a specific Laguna tile of interest by turning specific groups of oscillators on and off while monitoring photon emission.

2) Using EOFM to pinpoint individual cells: Once in the area of interest, we use EOFM to pinpoint the location of individual registers and drivers. To compare their visibility, we create a test circuit that drives a 100MHz square wave into 32 fabric registers and 24 Laguna registers. Figure 4 shows the resulting activity map under 20x magnification. Although fewer in number, the Laguna drivers appear much larger than the slice registers, making them uniquely exposed to probing. Notable is also that each spot on the left side of Figure 4 represents not only one, but two slice registers.

Analogous to how we located specific Laguna tiles, we further identify specific Laguna drivers within a Laguna tile by setting individual registers to toggle. The reconstructed layout of Laguna sites and individual drivers are visible in Figure 5b and Figure 5a respectively.

The entire process of starting from an unknown chip to locating points of interest takes a couple of hours. However, once located, the unique structure of Laguna drivers can be visually identified without EOFM. In Figure 3b, a Laguna column can be seen with Laguna tiles separated vertically by white bars. Similarly identifying individual elements in the FPGA fabric is not possible due to their small feature size.

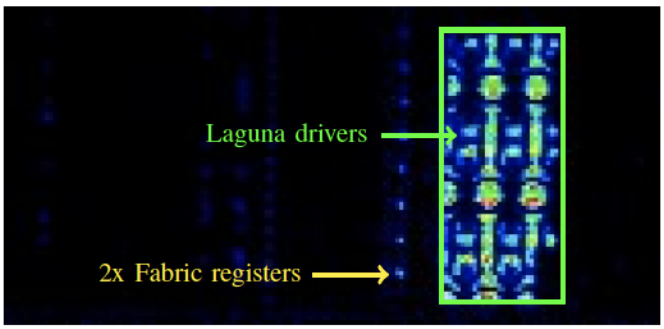

Fig. 4. EOFM of 32 fabric flip-flops and 24 Laguna registers toggling at 100MHz under 20x magnification.

B. Reading waveforms with EOP

Once we locate individual elements with EOFM, we use EOP to read out waveforms from the transistors. While this is also achievable in fabric, locating the points to probe requires much more effort as longer, higher-resolution scans are needed. Figure 6 shows the results of probing fabric (in blue) and Laguna registers (in red) at 5 and 100 integrations.

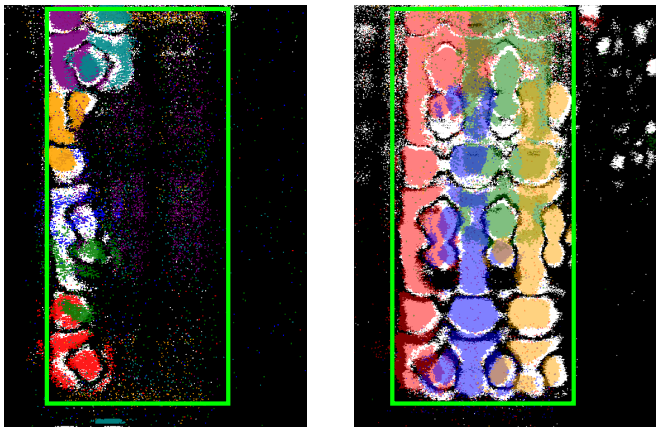

Fig. 5. By individually toggling groups of Laguna registers in different configurations, we can use EOFM to make out the layout of individual Laguna sites and drivers. In (a), we highlight the areas that light up when six registers in a single Laguna site are toggled. In (b) we redo the same for each site.

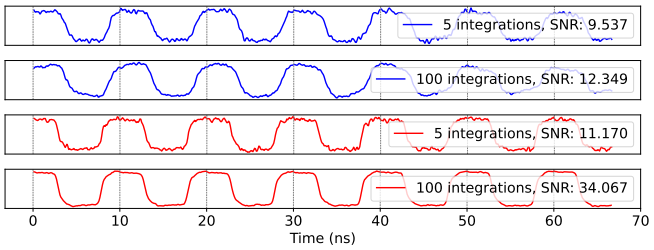

Fig. 6. Probed waveforms of fabric (blue) and Laguna (red) registers using 5 and 100 integrations.

For both traces, the quality of the signal increases with the number of integrations, however, the fabric waveform is markedly worse as evident from its noisier flat sections and lower SNR.

VI. Delay-Based Sensing and Mitigation

A natural way to protect against probing involves monitoring propagation delays of signals one wishes to protect, because the delay of a probed wire increases due to the increase in capacitance or temperature. In this section, we quantify the effect of contactless probing on wire delay.

A. Differential Sensing Circuits

Probing a specific spot on a die involves pointing a laser at the area of interest. The 1.3μm wavelength light source in the PHEMOS-X does not induce faults, but the point being probed experiences localized heating that leads to an increase in propagation delay. The temperature of the chip also varies due to environmental factors, and the drift in delay over time due to it might make it hard to discern the exact impact of probing. For this reason, we do not consider absolute delays of wires, but differential delays. That is, we record delays from two wires placed far apart, such that the probe only affects the first wire while environmental factors affect both. Accordingly, we define differential delay as the difference between the probed wire delay and the delay of the control wire.

B. Phase-Sweeping Sensor

We implement the same delay measurement mechanism detailed in [19], which measures the propagation delay of an interposer wire using two registers at the boundary of two chiplets. Using the dynamic phase shifting capability of the mixed-mode clock manager (MMCM) we then sweep the phase of the clock driving the receiving register while the transmitting register sends a known pattern of rising and falling edges that is used to check for correctness of the sampled values. Depending on the phase relationship between the transmitting and receiving clocks, values sampled by the receiving register will either be correct or incorrect. Initially, assuming a large enough phase difference, no values will be sampled incorrectly. However, as the phase difference is reduced (in increments of 14.286 ps) and the setup time of the circuit is violated, with an increasing probability the receiving register will sample values incorrectly. The phase difference (in ps) at which the transmitted rising edge is sampled correctly 50% of the time is taken as the propagation delay of the wire.

C. Delay Change

Figure 7 shows the output of the phase-sweeping sensor while its wire is probed at 100% laser power. The laser is toggled every two minutes, starting with it off. When the laser is off, the delay difference is -39.090 ps (the difference is negative as the control wire is slower than the probed wire). Once the laser is turned on, the delay of the probed wire increases, resulting in an average differential delay of -38.298 ps, a 0.792 ps jump. Figure 8b additionally shows the delay change of the same sensor at 75%, 50% and 25% power.

D. Obstacles to Reliable Detection

Our measurements show that the effects of probing can be observed by delay sensors, though conclusively detecting that a wire is being probed remains an open challenge as the delay at most increases by 0.8ps and does so steadily over a period of one second as shown in Figure 8a. The challenge is then to distinguish whether this increase is due to probing or other potential factors, such as temperature changes due to logic switching. Furthermore, Figure 8b shows a linear relationship between laser power and delay change, meaning the adversary could further reduce the laser power, making any delay changes even harder to distinguish from noise.

E. Mitigating Probing at Chiplet Interfaces

A limitation of contactless probing is that many traces need to be integrated to reconstruct the signal of interest. A simple and cheap mechanism to reduce probing vulnerability is for the transmitting chiplet to include a true random number generator (TRNG) that generates a one-time-pad to XOR with the outgoing data stream. The random pad is then transmitted with the data and the receiving chiplet can decode the incoming signals by combining the incoming data lines with the pad stream. As the TRNG never generates the same stream of values, even if the system is forced into replaying its data,

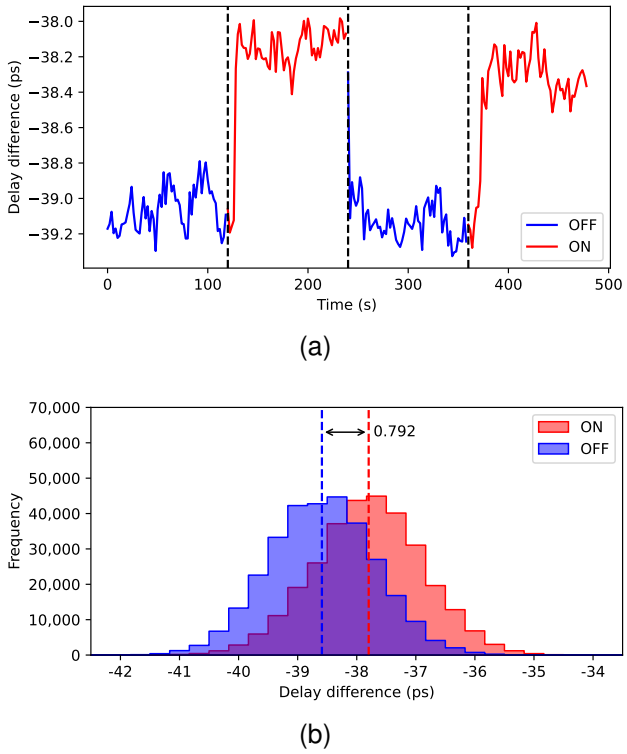

Fig. 7. Effects of probing a Laguna driver on interposer wire delay measured using the phase-shifting sensor. (a) shows the moving average of delay with a two-second window. (b) shows differential delay distributions when the laser is on/off

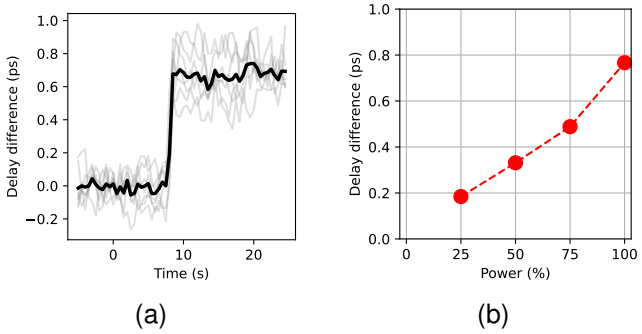

Fig. 8. (a) Change in differential delay after turning on the laser probe at 100% power (80mW). The solid black line represents the mean of all measurements. (b) Delay change of a wire is observed to have a roughly linear relationship to laser power.

like after resetting, the values on the interconnect wires will be random, preventing an adversary from being able to integrate multiple waveforms and extract the actual data.

VII. Conclusion

As SoC designs become increasingly common due to their performance and power benefits, chiplets are expected to continue to play an important role in the industry. Apart from packaging and test challenges, chiplets also bring unique security challenges. Our work shows that their large drivers that power inter-chiplet wires make them an easy and attractive target for probing. We additionally attempt to detect contactless probing using a cross-die delay sensor. Our findings indicate that while delay sensors are capable of measuring the small delay changes due to contactless probing, reliably distinguishing probing from environmental fluctuation remains a challenge. In the meantime, techniques like simple masking can help protect against probing attacks.

Acknowledgment

This research was supported in part by National Science Foundation Grants CNS-1902532 and CNS-2150123, and by a Research and Development (R&D) grant from the Massachusetts Technology Collaborative.

References

[1] T. Hackenberg and J. Lorenz, "Chiplet Market Update," 2023.
[2] "AMD Unveils World's Most Advanced Gaming Graphics Cards, Built on Groundbreaking AMD RDNA 3 Architecture with Chiplet Design," 11 2022, Investor Relations: AMD Press Release.
[3] "Meteor Lake Architecture Overview," 9 2023, Intel Tech Tour: Malaysia + Meteor Lake Tech Day.
[4] "NVIDIA Opens NVLink for Custom Silicon Integration," 3 2022.
[5] "Apple unveils M1 Ultra, the world's most powerful chip for a personal computer," 3 2022.
[6] D. Kehlet, "Accelerating innovation through a standard chiplet interface: The advanced interface bus (AIB)," *Intel White Paper*, 2017.
[7] M.-S. Lin, T.-C. Huang, C.-C. Tsai, K.-H. Tam, C.-H. Hsieh, T. Chen, W.-H. Huang, J. Hu, Y.-C. Chen, S. K. Goel, C.-M. Fu, S. Rusu, C.-C. Li, S.-Y. Yang, M. Wong, S.-C. Yang, and F. Lee, "A 7nm 4GHz Arm-core-based CoWoS chiplet design for high performance computing," in *2019 Symposium on VLSI Circuits*, 2019, pp. C28–C29.
[8] K. Saban, "Xilinx Stacked Silicon Interconnect Technology Delivers Breakthrough FPGA Capacity, Bandwidth, and Power Efficiency," *White Paper: Virtex-7 FPGAs*, 2012.
[9] M. Hutner, R. Sethuram, B. Vinnakota, D. Armstrong, and A. Copperhall, "Special session: Test challenges in a chiplet marketplace," in *2020 IEEE 38th VLSI Test Symposium (VTS)*. IEEE, 2020, pp. 1–12.
[10] N. Vashistha, M. L. Rahman, M. S. U. Haque, A. Uddin, M. S. U. I. Sami, A. M. Shuo, P. Calzada, F. Farahmandi, N. Asadizanjani, F. Rahman, and M. Tehranipoor, "ToSHI-Towards Secure Heterogeneous Integration: Security Risks, Threat Assessment, and Assurance," *Cryptology ePrint Archive*, 2022.
[11] M. S. M. Khan, C. Xi, M. S. U. Haque, M. M. Tehranipoor, and N. Asadizanjani, "Exploring Advanced Packaging Technologies for Reverse Engineering a System-in-Package (SiP)," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 2023.
[12] W.-C. Wang, Y. Wu, and P. Gupta, "Reverse engineering for 2.5-D split manufactured ICs," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 39, no. 10, pp. 3128–3133, 2019.
[13] I. Giechaskiel, K. Rasmussen, and J. Szefer, "Reading Between the Dies: Cross-SLR Covert Channels on Multi-Tenant Cloud FPGAs," in *IEEE International Conference on Computer Design (ICCD)*, 2019, pp. 1–10.
[14] E. Amini, T. Kiyan, L. Renkes, T. Krachenfels, C. Boit, J.-P. Seifert, J. Jatzkowski, F. Altmann, S. Brand, and S. Tajik, "Electrons Vs. Photons: Assessment of Circuit's Activity Requirements for E-Beam and Optical Probing Attacks," in *Proceedings of the 49th International Symposium for Testing and Failure Analysis*, 2023, pp. 339–345.
[15] T. Krachenfels, T. Kiyan, S. Tajik, and J.-P. Seifert, "Automatic extraction of secrets from the transistor jungle using Laser-Assisted Side-Channel attacks," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 627–644.
[16] S. Tajik, H. Lohrke, J.-P. Seifert, and C. Boit, "On the power of optical contactless probing: Attacking bitstream encryption of FPGAs," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '17, 2017, p. 1661–1674.
[17] H. Lohrke, S. Tajik, T. Krachenfels, C. Boit, and J.-P. Seifert, "Key extraction using thermal laser stimulation: A case study on xilinx ultrascale FPGAs," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 573–595, 2018.
[18] "Emission microscope C15765-01 PHEMOS-X," 6 2022.
[19] A. Deric and D. Holcomb, "Know Time to Die–Integrity Checking for Zero Trust Chiplet-based Systems Using Between-Die Delay PUFs," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 391–412, 2022.