Collective die-to-wafer bonding enabling low-loss evanescent coupling for optically interconnected System-on-Wafer

Pengfei Xu, Junwen He*, Koen Kennes, Anton Dvoretskii, Arnita Podpod, Guy Lepage, Negin Golshani, Rafal Magdziak, Swetanshu Bipul, Dieter Bode, Peter Verheyen, Maumita Chakrabarti, Dimitrios Velenis, Andy Miller, Yoojin Ban, Filippo Ferraro, Joris Van Campenhout

*imec, Kapeldreef 75, 3000, Leuven, Belgium - * now with Huawei Technologies Co., Ltd.*

e-mail: pengfei.xu@imec.be

Abstract:

We present a collective PIC die-to-wafer dielectric bonding process, enabling SiN waveguide-based die-to-wafer evanescent couplers with insertion losses of 0.36 ±0.18dB at 1310nm wavelengths, paving the way to optically interconnected 300mm wafer-scale multi-chip compute systems.

1. Introduction

The accelerating deployment of artificial intelligence and machine learning has led to an increasing demand for high-bandwidth, low-latency, and low-power interconnects between compute chips (XPU) and high-bandwidth memory (HBM) in multi-node AI/ML clusters [1,2]. By leveraging low-loss, high-density on-wafer waveguides, wafer-level optical interconnects (WL-OI) are an emerging technology to interconnect several tens of XPU's or HBM's with superior bandwidth density, energy efficiency and latency compared to fiber-coupled links, and with much longer reach compared to on-wafer electrical interconnects [3]. In future scaled-out implementations of WL-OIs, it may be beneficial from a cost and yield perspective to partition the optical interconnect in active PIC dies including modulators and photodetectors on the one hand, and a passive optical interconnect wafer including long-range routing waveguides and fiber coupling interfaces on the other hand (see Fig. 1a). However, such partitioning requires the availability of low-loss, high-yield optical coupling interfaces between the active PIC dies and the optical interconnect wafer.

In this paper, we present a high-precision collective die-to-wafer (CoD2W) dielectric bonding process, enabling low-loss SiN waveguide-based evanescent coupling between the assembled PIC dies and the 300mm-sized optical interconnect wafer (see Fig. 1b). Tapered SiN-waveguides, specifically designed for efficient and alignment-error tolerant, broadband evanescent coupling, are shown to yield insertion losses consistently below 0.5dB in the O-band, for evanescent couplers (EVC) as short as 0.5mm.

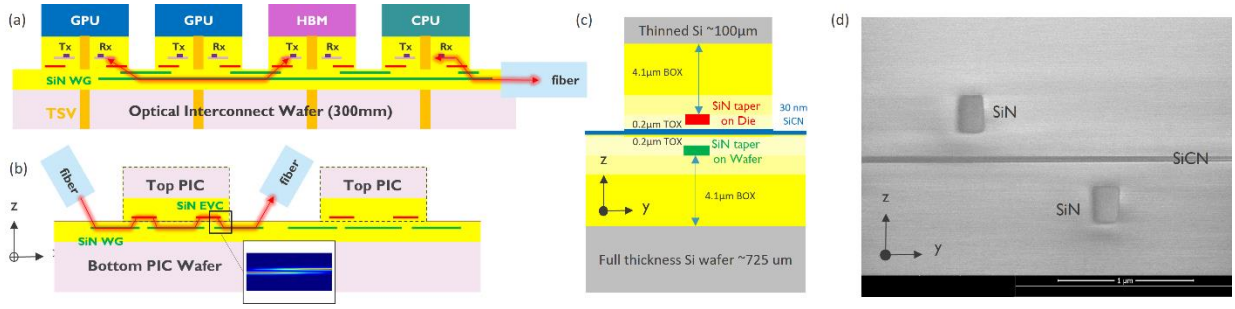

Fig. 1. (a) Schematic view of an optically interconnected wafer-scale compute system, leveraging EVCs for low-loss coupling between the active PICs and passive interconnect wafer, (b) simplified CoD2W assembled photonic system reported in this paper, (c) detailed schematic cross section of the assembly and EVC stack, and (d) cross-section TEM image of the CoD2W bonded SiN-based EVC.

2. Wafer fabrication, collective die-to-wafer assembly and SiN EVC design

The work presented in this paper has been carried out in imec's 300mm CMOS pilot line. The 300mm wafer processing starts with a 4.1µm thick silicon oxide deposition, which serves as a bottom cladding layer and reduces the interaction of the EVC mode with the Si substrate. Next, a 400nm-thick PECVD SiN layer is deposited and patterned using 193nm immersion lithography. After the SiN layer is patterned, an oxide top cladding layer is deposited and subsequently planarized, targeting 200nm of remaining oxide on top of the SiN. A thin SiCN layer is then deposited to enhance the CoD2W bonding strength. Subsequently, the assembly process starts by flipping and bonding the wafer to a first temporary carrier, to thin the Si substrate down to 100µm. Next, the 7x7 mm^2 sized

dummy "active" PIC dies are singulated using blade dicing. The PIC dies are then placed and aligned using primary alignment marks to a second temporary glass carrier. Finally, the PIC dies on the glass carrier are collectively transferred to the bottom PIC wafer and once again aligned, using secondary overlay markers. A schematic view of the resulting cross section is shown in Fig 1(c). Fig1(d) shows a TEM image, illustrating the two SiN EVC couplers separated by the 400-nm-thick oxide cladding and 30-nm-thick SiCN dielectric bonding layer.

Multiple design approaches exist to optimize the adiabatic SiN EVC's for realizing broadband low-loss coupling and robustness against alignment errors. In this work, we have adopted the "FAQUAD" (fast quasi-adiabatic) method [4], resulting in a typical EVC profile as depicted in Fig. 2c, tapering the SiN waveguide (WG) width from the nominal 710nm width (single-mode waveguide) down to a minimal 130nm width in the EVC region. EVC lengths of 0.5mm, 1mm and 1.5mm have been implemented to explore the trade-off between coupler footprint, coupling performance and robustness against misalignment. Of the 100 PICs bonded to the wafer as shown in Fig. 2a, 51 of them contain the FAQUAD EVCs, which we will report on in the remainder of this paper. The other 49 bonded PICs contain alternative EVC designs that will be reported elsewhere. To extract the EVC loss, three test structures with varying number of EVC transitions (0, 6 and 18) are implemented and measured at wafer-scale using SiN fiber grating couplers (Fig. 2d,f), using TE polarized laser light in the O-band. A bivariate linear fit is applied to decouple SiN EVC from SiN WG loss (Fig. 2f-i).

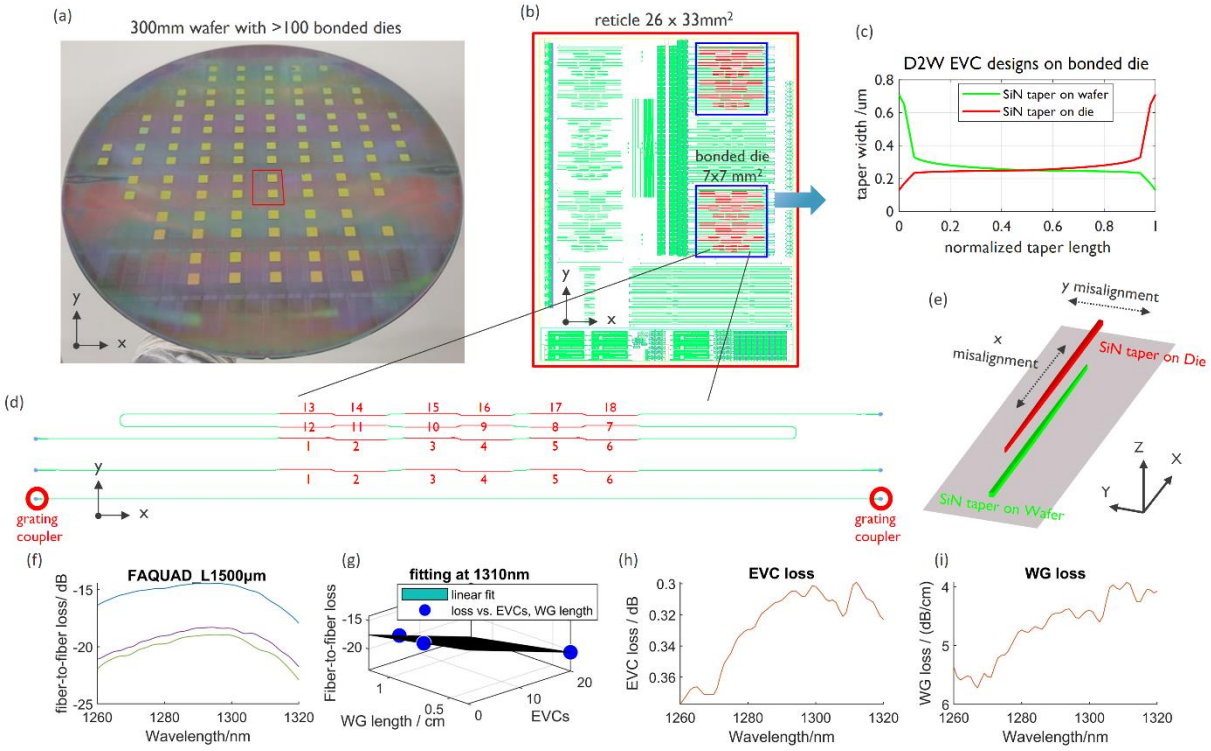

Fig. 2. (a) Photograph of a 300mm optical interconnect wafer with 100 assembled PIC dies, (b) overall layout of the bottom wafer and top PIC, (c) "FAQUAD" taper profile used for the SiN waveguide EVC, (d) EVC loss test macro, (e) x-y alignment definition, (f) typical measured fiber-to-fiber transmission spectrum, (g) bivariate fitting procedure to extract EVC and WG loss, (h) typical extracted EVC and WG loss spectra.

3. Wafer-scale measurement and analysis of SiN EVC loss

Using the test structures described above, we carried out wafer-scale measurements of the EVCs of varying lengths. Fig. 3a shows the resulting EVC loss spectra across 60nm of the O-band. At 1310nm wavelength, the majority of EVCs have extracted insertion losses below 0.5dB (see Fig. 3b), with mean ±3σ values of 0.36 ±0.18dB, 0.37 ±0.24dB, and 0.32 ±0.15dB for the 1.5mm-, 1mm- and 0.5mm-long EVCs respectively. At shorter wavelengths, several dies show a higher EVC loss, which occurs mostly on PIC dies with larger lateral (y) misalignment. It should be noted that the SiN WG propagation loss on the reported wafer (and top PIC dies) is relatively large at 4-6dB/cm due to a processing issue. Part of this excess propagation loss is also embedded in the extracted EVC loss, and by reducing this propagation loss in future experiments, we expect to be able to reduce EVC loss by 0.1-0.2dB.

In terms of overall optical yield, the 1.5-mm long EVC design performs best at 75.5%, followed by the 1-mm and the 0.5-mm long EVC at 68% and 57% respectively. Imperfect yield obtained in this initial development run is caused by several factors, including die loss during the CoD2W assembly process, undesired voids (mostly appearing at the die edges), and lateral misalignment (y). The 0.5mm long EVCs are particularly sensitive to the latter, as illustrated in the left panel of Fig. 3c, showing 7 dies with non-functional EVCs having a lateral misalignment (y) beyond 1µm. The 1.5mm long EVC can tolerate up to 1.5µm lateral misalignment, as shown in the right panel of Fig. 3c. As expected, longitudinal misalignment (x) has a much less pronounced effect on coupling loss. All the factors driving yield loss are being addressed by optimizing the CoD2W bonding process and results will be reported in the future.

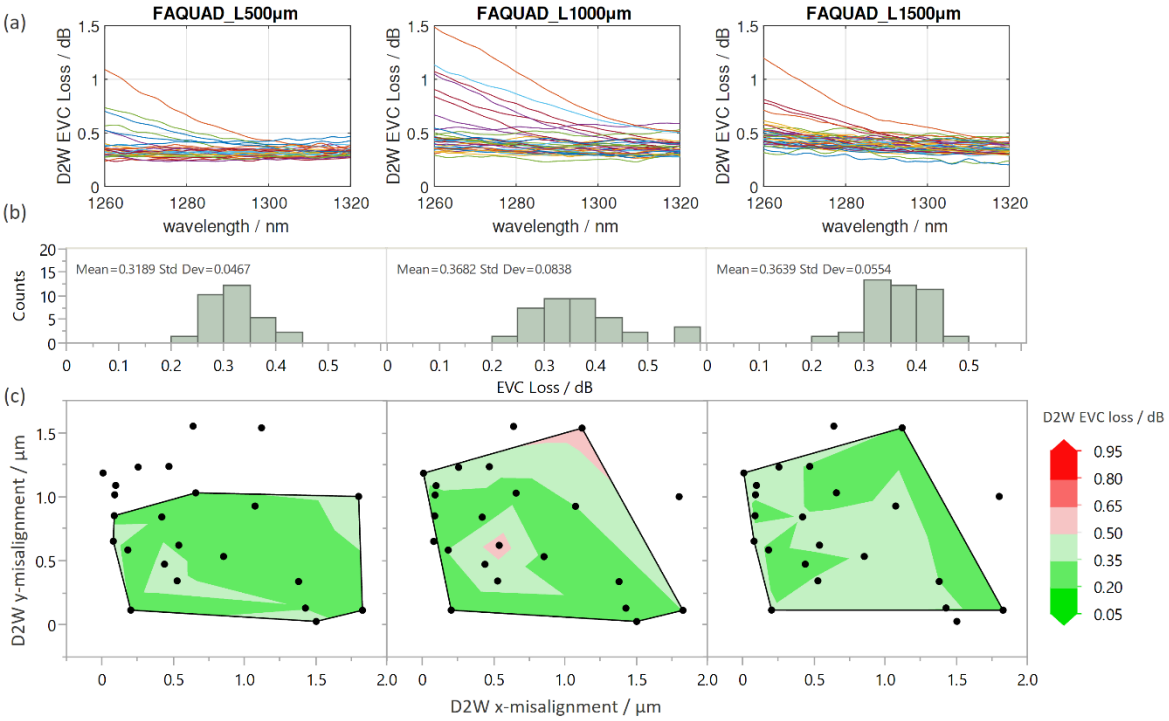

Fig. 3. Wafer-scale EVC loss measurement vs. coupler length (a) loss spectra, (b) loss statistics at 1310nm wavelength, and (c) loss vs. x- and y- misalignment. The data points without color represent failing EVCs.

4. Discussion and conclusion

In this paper, we demonstrate high-precision CoD2W assembly with dielectric bonding as a promising approach to realize low-loss evanescent coupling (<0.5dB) between top PIC dies bonded to a bottom PIC wafer with encouraging initial yield (up to 76%), paving the way to future active/passive partitioned wafer-scale optical interconnect fabrics. By combining the demonstrated EVCs with copper-based hybrid bonds and through-silicon via's for electrical connectivity and ultra-low loss (SiN) waveguides in the optical interconnect wafer for long-range links, we envisage the realization of future wafer-scale multi-XPU compute systems with unparalleled optical interconnect bandwidth, energy efficiency and latency.

Acknowledgement
This work was supported by imec's industry-affiliation R&D program "Optical I/O".

References

[1] H. Liu et al., "Lightwave Fabrics: At-Scale Optical Circuit Switching for Datacenter and Machine Learning Systems," in Proceedings of the ACM SIGCOMM 2023 Conference, 2023, pp. 499–515.
[2] M. Khani et al., "SiP-ML: High-Bandwidth Optical Network Interconnects for Machine Learning Training," in Proceedings of the 2021 ACM SIGCOMM 2021 Conference, 2021, pp. 657–675
[3] S. Pal et al, "Architecting Waferscale Processors - A GPU Case Study," 2019 IEEE International Symposium on High Performance Computer Architecture (HPCA), Washington, DC, USA, 2019, pp. 250-263
[4] D. Siriani and Jean-Luc Tambasco, "Adiabatic guided wave optics – a toolbox of generalized design and optimization methods," Opt. Express 29, 3243-3257 (2021)