

Interconnect Technologies for Terabit-per-second Die-to-Die Interfaces

Behzad Dehlaghi¹, Rudy Beerkens², Davide Tonietto², and Anthony Chan Carusone¹

¹Dept. of Electrical & Computer Engineering, University of Toronto, Toronto, Canada

²Huawei Canada Research Centre, HiSilicon Division, Ottawa, Canada

tony.chan.carusone@isl.utoronto.ca

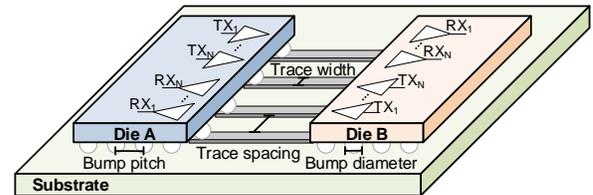
Abstract — Seamless package-level integration of multiple dies for high-performance computing and networking requires broadband dense die-to-die interconnect. Organic packaging substrates offer lower cost and lower loss interconnect, whereas silicon interposers offer higher density interconnect. In this work, a silicon interposer is fabricated in a relatively inexpensive 0.35 μm CMOS technology as an alternative to conventional organic or silicon interposer substrates. Flip-chip assembly technologies such as solder and gold-stud bumping are discussed. Measured eye diagrams at 16.4 Gb/s and bathtub curves at 20 Gb/s show the impact of assembly and bumping technology on the link performance. Considering signal integrity issues such as inter-symbol interference (ISI) and crosstalk, the maximum achievable aggregate bit rate is estimated for different interconnect lengths.

Index Terms — Bandwidth density, crosstalk, die-to-die communication, organic substrates, silicon interposers.

I. INTRODUCTION

Package-level integration can enable high aggregate bandwidth within heterogeneous multi-chip systems. Applications that require broadband interfaces between two dies in different technologies (e.g. CPU and memory, digital and high-speed/analog, etc.) stand to benefit particularly. Fig. 1 illustrates two dies that are flip-chip mounted onto a silicon or organic interposer (referred to generically as a substrate). The representative dimensions for interconnect trace width and spacing are also shown in Fig. 1. The dies may be attached to the substrate with either C4/ μC4 /gold-stud bumps or copper pillars. Finer pitch is currently achievable using silicon interposer substrates partly because their thermal expansion coefficient is well-matched to silicon dies, thereby obviating mechanical stresses. Typical bump diameter and pitch dimensions are also shown in Fig. 1. Although a higher density of interconnects/bumps is achievable on silicon interposers [1] rather than their organic counterparts [2], the interconnects on silicon interposers have more frequency-dependent insertion loss. This can limit the achievable bit rate, depending on the required interconnect length.

This paper explores the trade-offs in different packaging solutions between interconnect length, area, bandwidth, and cost. A link model is developed based on measured results of a silicon interposer prototype in a low-cost 0.35 μm CMOS technology. Discontinuities in die-to-die communication are



Interposer (Substrate)	Trace Width	Trace Spacing	Bump Pitch	Bump Diameter
Silicon	3 μm	9 μm	40 μm	26 μm
Organic	30 μm	75 μm	150 μm	100 μm

Fig. 1. Die-to-die communication over an organic/silicon interposer with representative trace and bump dimensions.

discussed and measurement results are provided for different die attachment technologies. The remainder of the paper is organized as follows. Section II describes the link model including the interconnects and attachment bumps. The link model is validated with measurements of a prototype interposer. The signal integrity impairments along with the maximum achievable bandwidth for each packaging solution are discussed in Section III and conclusions are drawn in Section IV.

II. LINK MODEL

Fig. 2 shows the link model used in this work to compare the three interposer technologies in Fig. 3. The transmitter and receiver are substituted with their equivalent termination impedances equal to the characteristic impedance of the interconnect. The bump and via inductance and capacitance are calculated from theoretical equations. The interconnects

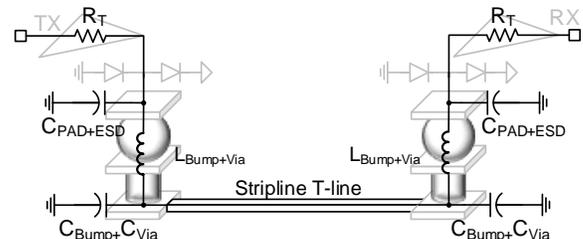


Fig. 2. Link model including pad+ESD capacitance and bump+via lump model.

		Organic Int.	Si Int.	Low-cost Si Int.
ϵ_r		3.2	3.9	3.9
$\tan\delta$		0.007	0.001	0.001
σ_M		4.83×10^7	4.83×10^7	3.2×10^7
H		$33 \mu\text{m}$	$3 \mu\text{m}$	$1 \mu\text{m}$
T_S		$16 \mu\text{m}$	$3 \mu\text{m}$	$0.64 \mu\text{m}$
T_{GB}/T_{GT}		$16/16 \mu\text{m}$	$1/1 \mu\text{m}$	$0.64/0.92 \mu\text{m}$

Fig. 3. The stripline configuration and parameters in different interposers.

are modeled using a 2D electromagnetic field solver. In the following subsections, the details of the presented model along with some experimental results are presented.

A. Interconnect Model

In [3], it was shown that assuming a constant interconnect pitch in silicon interposers, single-ended signaling without ground shielding offers lower crosstalk and insertion loss compared to either differential signaling or single-ended signaling with ground shielding. Therefore, all the interconnects in this work are single-ended. In all three interposers, the interconnects are built using the 2nd metal layer from the top in a stripline configuration which has better immunity to crosstalk. Fig. 3 shows the stack-up of the three packaging substrates along with their electrical and geometrical parameters.

A silicon interposer was fabricated in TSMC 0.35 μm CMOS technology. Fig. 4a shows the die photo of the fabricated silicon interposer. There are two 4.2 mm and 6 mm interconnects on the low-cost silicon substrate. Fig. 4b and 4c show the simulated vs. measured results of the insertion loss and return loss respectively. The DC loss of the interconnects (due to large series resistance of thin traces) leads to poor return loss at frequencies below 5 GHz. There is a good correlation between the simulated and measured results which validate the accuracy of the models. Moreover, the simulated results of the silicon interposer described in Fig. 3 are verified with the measured results published in [4].

B. Bump and Via Model

In die-to-die communication, the C4/ μC4 bumps, copper

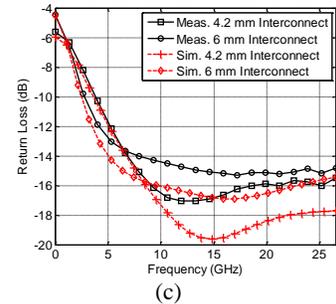
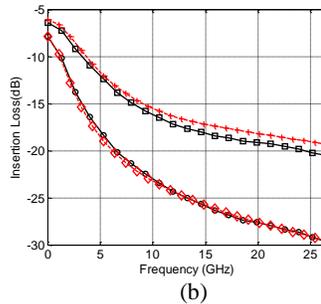
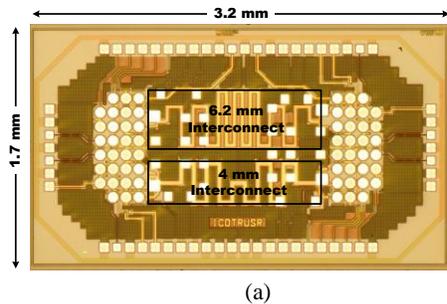


Fig. 4. (a) Die photo of silicon interposer; measured vs. simulated results of 4.2 mm and 6 mm interconnects (b) insertion loss (c) return loss.

TABLE I
DIFFERENT METHODS OF DIE-SUBSTRATE ATTACHMENT

Bump	C4	μC4	Cu Pillar
Diameter	$100 \mu\text{m}$	$70 \mu\text{m}$	$26 \mu\text{m}$
Pitch	$150 \mu\text{m}$	$100 \mu\text{m}$	$40 \mu\text{m}$
$L_{\text{bump}}+L_{\text{via}}$	$11 \text{ pH} + 2 \text{ pH}$	$5 \text{ pH} + 0$	9 pH
$C_{\text{bump}}+C_{\text{via}}$	$7 \text{ fF} + 4 \text{ fF}$	$2 \text{ fF} + \sim 0$	$3 \text{ fF} + \sim 0$
C_C	3.5 fF	2.5 fF	4 fF
$C_{\text{PAD+ESD}}$	150 fF	120 fF	75 fF

pillars and/or μvias /vias (excluding core vias in build-up organic packaging substrates) are the primary sources of discontinuity. The C4 bumps are most commonly used commercially for die attachment on organic substrates. Silicon interposers may take advantage of the fact their coefficient of thermal expansion matches that of the silicon dies and use smaller bumps such as μC4 bumps or copper pillars. The bumps and vias can be modeled with lumped elements [5]. The inductance of the bump/via can be approximated as

$$L = \frac{\mu h}{2\pi} \left[\ln \left(\frac{h}{r} + \sqrt{\left(\frac{h}{r} \right)^2 + 1} \right) + \frac{r}{h} - \sqrt{\left(\frac{r}{h} \right)^2 + 1} \right], \quad (1)$$

while its capacitance is

$$C = \frac{2.82 \epsilon_0 \epsilon_r \cdot r \cdot h}{D - 2r}, \quad (2)$$

and its coupling capacitance from a neighboring bump/via is

$$C_C = \frac{\pi \epsilon_0 \epsilon_r h}{\ln \left(\frac{p}{2r} + \sqrt{\frac{p^2}{4r^2} - 1} \right)}, \quad (3)$$

where r is the bump/via radius, h is the bump/via height, D is the diameter of the antipad, and p is pitch from the neighboring bump/via. Table I summarizes different types of bumps with their lumped element models. Parasitics L_{via} and C_{via} for μC4 bumps and copper pillars are negligible due to the use of extremely small vias (typically $0.5 \mu\text{m}$ in width/length) in silicon interposers.

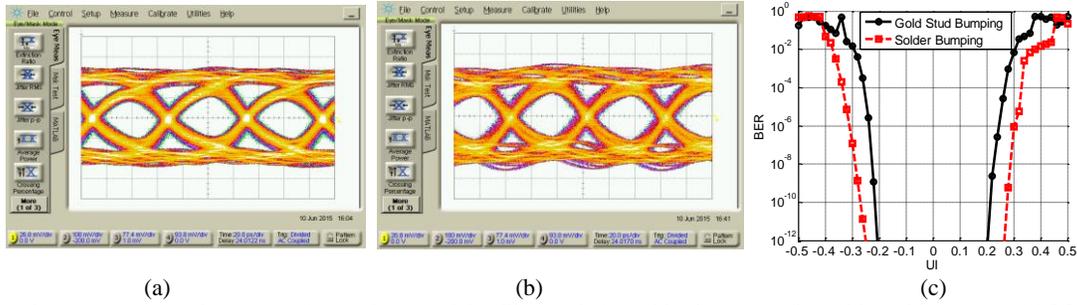


Fig. 5. Measured transmitter eye diagrams after package at 16.4 Gb/s with (a) solder bumping (b) gold-stud bumping; (c) Measured received bathtub curve at 20 Gb/s over 2.5 mm interconnect on low-cost interposer with gold-stud and solder bumping.

The parasitics associated with the die-attachments in Table I are negligible compared to pad+ESD capacitance, however imperfections during flip-chip assembly can increase them. For example, gold-stud bumping allows for geometries similar to μ C4 bumps using gold in place of the solder balls. Hence, the performance of gold-stud bumps should be practically the same as that of μ C4 bumps. However, unlike flip-chip assembly with μ C4's, gold-stud bumps are not self-aligning which can lead to an increase in parasitics. Fig. 5a and b show measured transmitter eye diagrams 16.4 Gb/s after package using the transceiver described in [6] and attached to a silicon interposer with 4 dB loss using solder bumping and gold-stud bumping respectively. The package with gold-stud bumping appears to have more discontinuities which results in degradation in signal integrity. Fig. 5c shows the received bathtub curves measured using the transceiver in [6] over a 10.7 dB loss channel at 20 Gb/s that indicates the degradation in signal integrity using gold-stud bumping.

III. MAXIMUM ACHIEVABLE BIT RATE

To calculate the maximum achievable bit rate for each packaging solution, signal integrity impairments such as channel loss and crosstalk need to be taken into account.

DC channel loss results in reduced signal swings at the receiver while frequency dependent channel loss causes ISI. As the channel loss increases, sophisticated equalization circuits are required to compensate for it, but the power overhead of these circuits are not tolerable due to tight power budgets in die-to-die communication. In [6], we used a passive equalizer to compensate for 12.9 dB loss at 8.2 GHz (7.1 dB relative to DC) with no extra power consumption. Therefore, when calculating the maximum achievable bandwidth, it is assumed that up to 10 dB of relative to DC loss can be compensated using a passive equalizer.

Coupling between the interconnects introduces crosstalk which determines the minimum required spacing between them. This limits the number of parallel interconnects within a fixed area on a single layer of the substrate. Fig. 6 shows simulation results of near-end and far-end crosstalk for a 5-

mm interconnect on the organic substrate. The crosstalk decreases up to a channel spacing of 75 μ m, and it does not change significantly after that. To see the impact of crosstalk on the link performance, simulated bathtub curves for different channel spacings are shown in Fig. 7. It can be seen that the link performance is not degraded by crosstalk once the channel spacing is more than 75 μ m. Based on these simulation results, once the difference between the insertion loss and crosstalk is more than 30 dB from DC to $f_{bit}/2$, the effect of crosstalk on link performance is negligible. This is used as the criteria for choosing the spacing between traces in the maximum bandwidth.

All links are single-ended. The maximum bit rate for each lane is limited to 30 Gb/s to ensure low-power transceivers are practical. Each trace width is set such that their characteristic impedance is 50 Ω at high frequencies. The absolute value of insertion loss at $f_{bit}/2$ is limited to -20 dB to ensure sufficient DC swing at the receiver. Finally, the maximum achievable bandwidth for organic, silicon, and

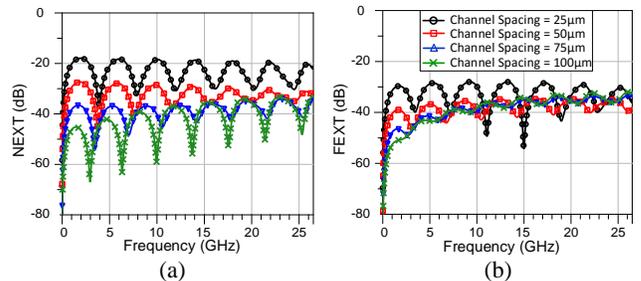


Fig. 6. Crosstalk simulation results for different 5-mm interconnect spacings on the organic interposer (a) near-end crosstalk (NEXT) (b) far-end crosstalk (FEXT).

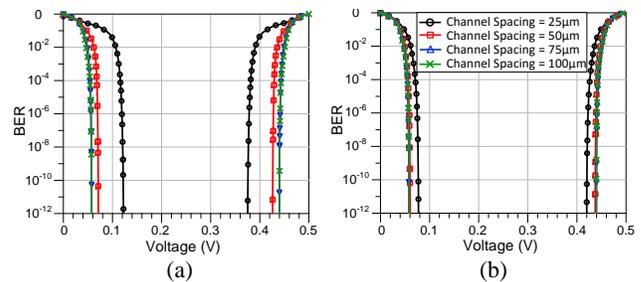


Fig. 7. Simulated bathtub curves for different 5-mm interconnect spacings on the organic interposer with (a) NEXT (b) FEXT.

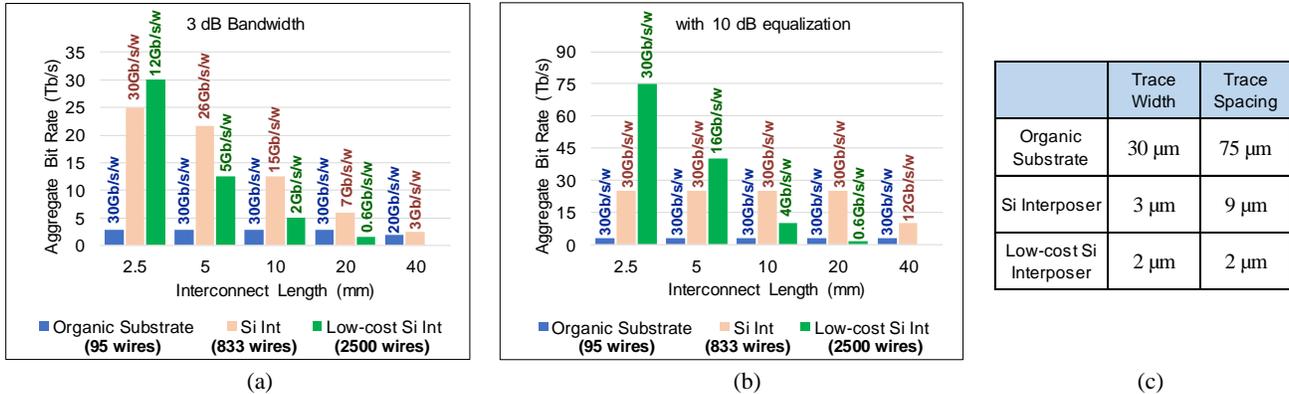


Fig. 8. Maximum aggregate bit rate for three substrates at different interconnect lengths assuming 1 cm total chip edge and (a) 3 dB bandwidth (b) with 10 dB equalization; (c) trace width and spacing for different substrates.

low-cost silicon interposers can be calculated.

Fig. 8 shows the maximum aggregate bit rate achievable assuming a 1-cm wide bus of interconnect routed on a single layer of the substrate within 1 cm for different interconnect lengths ranging from 2.5 mm to 40 mm. The number of wires in the low-cost silicon interposer is 25 times more than organic substrates. This is because of the thin metal and dielectric layers in the low-cost silicon interposer allows for a smaller wiring pitch without introducing significant crosstalk. Without any equalization, it is assumed that insertion loss must be limited to only 3 dB (relative to DC) to ensure acceptable signal integrity. Over distances 2.5 mm or less, over 10 Gb/s/w is achievable on the low-cost silicon interposer. It is therefore preferable to the conventional silicon interposer whose coarser wiring pitch would necessitate impractically high per-wire data rates (>30 Gb/s/w) in order to achieve comparable aggregate bandwidths. However, once the channel length is increased to 5 mm or more, the high frequency-dependent and DC losses of the low-cost interposer limit its performance and make a conventional interposer preferable. Organic substrates offer performance comparable to the silicon interposer for interconnect 40 mm or longer, as shown in Fig. 8a, in spite of the far fewer number of available wires, due to their much lower loss.

As shown in Fig. 8b, once a passive equalizer is used and 10 dB insertion loss (relative to DC) is tolerable, the low-cost silicon interposer achieves the best aggregate bandwidth over 2.5 mm and 5 mm interconnects, and the silicon interposers have better aggregate bandwidth than organic substrates at all other interconnect lengths.

IV. CONCLUSION

Silicon and organic substrates have been discussed as suitable candidates for high-density die-to-die communication. An alternative low-cost silicon interposer

has been introduced and measured results using a 0.35- μm CMOS prototype have been shown. A link model based on the measured interconnects and lumped elements for bumps and vias has been used to evaluate the signal integrity of the different substrate technologies. Finally, the maximum aggregate bit rate for each technology has been found for different interconnect lengths. These results can be used to choose a suitable substrate technology for different applications based on the required interconnect lengths, aggregate bandwidth, and cost.

ACKNOWLEDGMENT

The authors would like to acknowledge CMC for chip fabrication and provision of CAD tools and equipment.

REFERENCES

- [1] T. O. Dickson, Y. Liu, S. V. Rylov, B. Dang and et. al., "An 8x 10-Gb/s Source-Synchronous I/O System Based on High-Density Silicon Carrier Interconnects," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 4, pp. 884-896, April 2012.
- [2] J. W. Poulton, W. J. Dally, C. X. and e. al., "A 0.54 pJ/b 20 Gb/s Ground-Referenced Single-Ended Short-Reach Serial Link in 28 nm CMOS for Advanced Packaging Applications," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 12, pp. 3206-3218, December 2013.
- [3] A. C. Carusone, B. Dehlaghi, R. Beerkens and D. Tonietto, "Ultra-Short-Reach Interconnects for Package-Level Integration," in *Optical Interconnects Conference (OIC)*, San Diego, May 2016.
- [4] X. Gu, L. Turlapati, D. Bing, and et. al., "High-Density Silicon Carrier Transmission Line Design for Chip-to-Chip Interconnects," in *IEEE 20th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*, San Jose, October 2011.
- [5] N. Pham, B. Mutnury, E. Matoglu, M. Cases and D. N. De Araujo, "Package Model for Efficient Simulation, Design, and Characterization of High Performance Electronic Systems," in *IEEE Workshop on Signal Propagation on Interconnects*, Berlin, 2006.
- [6] B. Dehlaghi and A. C. Carusone, "A 20 Gb/s 0.3 pJ/b single-ended die-to-die transceiver in 28 nm-SOI CMOS," in *Custom Integrated Circuits Conference (CICC) 2015*, San Jose, 2015.