

Enabling High-Speed Die-to-Die Interfaces with Wallstrip: A Study on Insertion Loss and Crosstalk Metrics with a Novel Transmission Line Structure

Stephen Newberry, Chipletz, Inc.
stephen.newberry@chipletz.com

Ching-Ping Wong, Chipletz, Inc.
ching-ping.wong@chipletz.com

Victor Kronberg, Chipletz, Inc.
victor.kronberg@chipletz.com

Abstract

This paper focuses on multi-die packages utilizing multiple chiplets, and the process of optimizing next-generation die-to-die (D2D) interconnects by addressing two critical signal integrity parameters: Insertion Loss (IL) and Crosstalk (XT). Current D2D standards vary in their specification detail: High-Bandwidth Memory (HBM) does not specify IL/XT requirements, Bunch of Wires (BoW) includes only IL, and UCIe specifies Voltage Transfer Function (VTF) IL/XT.

A novel approach based on a unique transmission line, Wallstrip, is proposed. Compared to traditional microstrip or stripline, Wallstrip redefines the signal-to-reference structure in package routing and achieves notably improved IL performance. This paper will describe the package substrate-integrated Wallstrip D2D routing and benchmark against traditional D2D routings seen on silicon interposers and bridges utilizing electromagnetic simulation as well as measurement-based correlation.

The Wallstrip 2-D channel will be presented, followed by a review of the challenges and design constraints related to return path construction. These evaluations highlight the advantages of Wallstrip and yield signal/return path design recommendations, providing key insight into signal integrity tradeoffs when scaling up to higher data rates. Together, they demonstrate Wallstrip's ability to improve signal integrity and establish it as a valuable alternative routing scheme within the D2D interconnect landscape.

Author(s) Biography

Stephen Newberry received his B.S. degree in electrical and computer engineering (ECE) from Rutgers University, New Brunswick, NJ, USA in 2015 and his M.S. degree in electrical engineering (EE) from the University of Idaho, Moscow, ID, USA in 2024. Prior to his B.S. degree, he was an Electrician's Mate with the U.S. Coast Guard. He has held various roles in hardware design and signal integrity engineering, most recently as a Signal and Power Integrity Engineer at Chipletz.

Ching-Ping Wong earned his Master of Science degree in Electrical Engineering at Mississippi State University in 1996. He started his career at Intel in RTL design engineering; then detoured into MCM package layout engineering at AIC Microelectronics before returning to Intel as I/O transistor-level circuit designer. He then worked as a microprocessor circuit designer during the height of the dot-com era at Santa Clara, California before returning as an I/O circuit designer and signal integrity engineer at Broadcom in 2002. He is currently with Chipletz, Inc, as Fellow Engineer in Signal and Power Integrity Engineering, since incorporated in 2021.

Victor Kronberg is an electrical engineer specializing in IC package design. He holds a master's degree in electrical & computer engineering (ECE) from the University of Colorado Boulder (2021) and multiple bachelor's degrees from the University of Texas at Austin (2014). Prior to joining Chipletz as an IC package design engineer, he was a field applications engineer at Siemens EDA (formerly Mentor Graphics).

Introduction

Chiplet technology is critical for achieving the computing performance necessary as the industry pushes for more powerful hardware, especially in the domain of AI/ML (Artificial Intelligence/Machine Learning) training and inference. The primary *raison d'être* of chiplet technology is to reduce costs by improving overall yields as transistor scaling slows down [1]. Traditional flip-chip die attachment has been done by attaching a single, monolithic die to an organic package substrate through a Controlled Collapse Chip Connection (C4) bump, or a smaller μ bump as depicted in Figure 1 [2]. As more features are added to the die—such as parallel and serializer/deserializer (SerDes) connectivity, wireless connectivity, optical and radio subsystems, security, compute and graphics cores including cache memory and flash—the die size will ultimately grow. Further, the entire die must be fabricated in the same process node for all features, even if all portions of the die do not require such a small process node. This integration will both increase cost and reduce yields of the die.

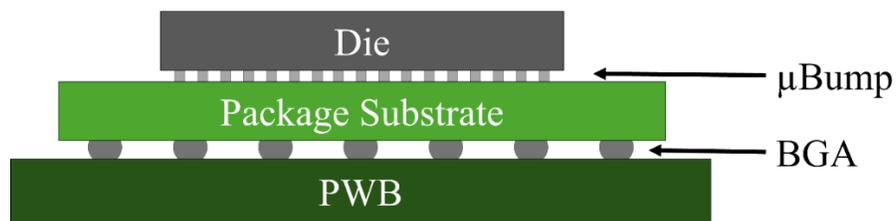


Figure 1: Flip-chip attachment of a traditional, monolithic die through an organic package substrate.

To meet the goal of improving yield and thus reducing cost, the individual functions of this monolithic die can be broken into smaller *chiplets*. This allows each chiplet to use a different process node which may be optimized for its specific function. For example, while the core processor may require a small feature size such as 5 nm, perhaps the cellular modem can still be constructed with 28 nm. By mixing process nodes, existing silicon can be reused where either schedule or technological constraints would not allow for updating that subsystem in the device. Further, a larger die size results in a higher probability of wafer defect which could cause a failure of the entire monolithic die. If each chiplet were smaller, then that one defect spot causes an overall improved yield for the total silicon area in the device [2]. Finally, it is feasible that the total silicon area required for a product exceeds the reticle size, thus making the device impractical to build as a monolithic die, even with reticle stitching. Transitioning from monolithic to chiplet architecture provides significant cost savings, including one real-world example in which the chiplet option is 41% cheaper than single-die [3].

One of the primary applications of an advanced chiplet architecture is in AI/ML accelerator devices, which require large memory bandwidths. To sustain performance in these systems, it is critical to provide ample, low-latency memory to the compute core and to ensure sufficient bandwidth such that the memory interface does not become a bottleneck to further computation. In fact, the performance of these AI workloads is often limited by the overall throughput available in the memory interface [4]. Unfortunately, shifting from monolithic to chiplet architecture necessitates on-package, die-to-die interfaces which are simply not present in a single die. As such, the Signal Integrity (SI) community must devise

methods of supporting ever-increasing on-package interface bandwidths for each successive AI accelerator device family.

Advanced Packaging Techniques

The world of semiconductor packaging has progressed leaps and bounds ahead of the chip carriers and wire-bonds of early devices. While Microminiature Multichip Modules (MCMs) have been in use for decades, even with flip-chip attachment and feature sizes as small as $10\mu\text{m}$ [5], today's AI accelerators are orders of magnitude more complex; they often feature thousands of pins at the Printed Circuit Board (PCB) interface and hundreds of thousands of pins at the die interface. Recent years have seen numerous efforts to achieve reliable, high-performance interconnects between multiple on-package dies. Some technologies such as Samsung Si-Less and TSMC InFO rely on redistribution layers (RDL) built directly onto the wafer substrate, which then connects to the package substrate via a C4 interface [2]. Others, such as TSMC's Chip-on-Wafer-on-Substrate (CoWoS), rely on an interposer which the dies attach to; this interposer is then assembled on an organic package substrate with a C4 interface as shown in Figure 2 [2].

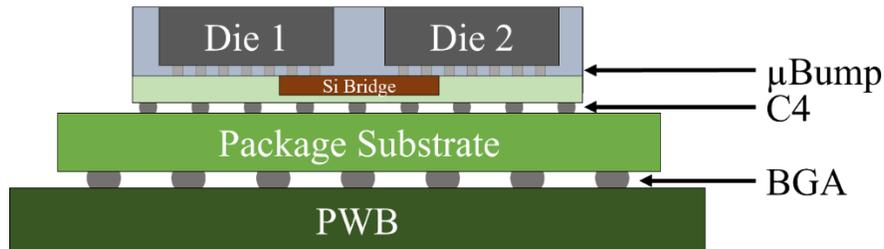


Figure 2: Device which uses an interposer to attach multiple chiplets to a single package substrate via C4 interface. This substrate is then attached to the Printed Wiring Board (PWB) through a Ball Grid Array (BGA) interface.

While innovative, this configuration requires an additional connection point between the die bumps and the BGA, introducing an impedance discontinuity for high-speed signals and a localized increase in impedance and inductance in the power distribution network (PDN). Different variations of interposer technology exist such as fully silicon (CoWoS-S), fully organic (CoWoS-R) and organic with silicon bridges (CoWoS-L) as shown in Figure 2. The electrical characteristics of the interposers present significant challenges to meeting die-to-die interface standards, especially in the area of insertion loss (IL). The maximum data transfer rate is fundamentally limited by both IL and XT in a package, and methods of reducing these two properties must be explored. As an alternative architecture, a package substrate which is capable of fully routing these tight-pitch die-to-die interfaces within its organic RDL, such as the one shown in Figure 3, could provide substantially better IL while eliminating the C4 interface for off-package signals. Though still in the development phase, this type of structure is gaining traction in the industry.

Removing the interposer layer from the package not only improves the electrical characteristics but also simplifies the assembly process. The supply chain is also streamlined because fewer vendors are required in the production of the end device.

However, new challenges arise in designing a die-to-die interface capable of reliable data transfer in a reduced layer-count configuration.

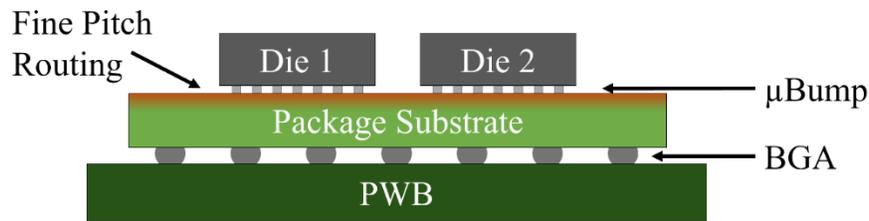


Figure 3: A device in which the chiplet dies are attached with a μ bump interface directly onto a package substrate without an interposer. The substrate integrates fine-pitch routing in its upper layers to accommodate high-density, die-to-die interconnect.

Die-to-Die Interface Standards

Many different electrical interfaces are available to connect one chiplet die to another such as JEDEC’s High Bandwidth Memory (HBM), Open Compute Project’s (OCP) Bunch of Wires (BoW), Optical Internetworking Forum’s (OIF), Extra Short Reach (XSR), the CHIPS Alliance’s Advanced Interface Bus (AIB) and the Universal Chiplet Interconnect Express (UCIe). Of these, the latter has emerged as one of the most popular [6].

Table 1: Comparison of Common Die-to-Die Interfaces

Standard	Maximum Data Rate	Trace Configuration	Modulation Scheme	Data Rate per Conductor
HBM4 [7]	8 Gb/s	Single-Ended	NRZ	8 Gb/s
BoW [8]	32 Gb/s	Single-Ended	NRZ	32 Gb/s
AIB [9]	6.4 Gb/s	Single-Ended	NRZ	6.4 Gb/s
UCIe [10]	64 Gb/s	Single-Ended	NRZ	64 Gb/s
CEI-112G-XSR-PAM4 [11]	112 Gb/s	Differential (2-conductor)	PAM-4	56 Gb/s

Within the UCIe standard, many different variables allow the designer to optimize for whichever form factor is most critical, whether that is signal density, data rate, or some combination in order to achieve the desired throughput. For example, a Standard Package supports up to 32 data bits for a 25mm reach while an Advanced Package supports up to 64 data bits with a 2mm reach. Since each different package has recommended bump maps, the specification provides the maximum die edge bandwidth density in GB/s/mm of beachfront width [10, Sec. 1.4].

While other investigations of UCIe interfaces have shown better crosstalk performance with silicon [12] due to smaller trace and via dimensions, we will instead show the advantages of using thicker organic dielectrics. In this paper, we will present Wallstrip: a novel transmission line structure which preserves optimal IL and XT performance in die-to-die interfaces when routing on layer-constrained substrates. We will compare Wallstrip to a more traditional silicon bridge channel and evaluate the critical metrics required for a UCIe interface in a pre-layout, transmission-line simulation while ignoring the effects of transition vias and die-attach bumps. We will show how a Wallstrip channel can enable a

die-to-die interface to operate at a higher data rate, or for longer channels than a similar silicon bridge interface. Finally, we introduce the Bandwidth Density per Routing Area (BDRA) measure which quantifies the throughput in a chiplet package while considering the overall route area across all redistribution layers.

Formulation

Wallstrip Transmission Line

A myriad of transmission line structures have been devised to meet ever-changing system requirements. In fact, Wadell's canonical reference on the topic [13] includes nine coaxial, sixteen coplanar or parallel, ten microstrip and seven different stripline structures within its repertoire of single-ended transmission lines. We propose an additional structure: Wallstrip. The Wallstrip transmission line as visualized in Figure 4 is neither fully coplanar nor fully stripline due to the large height-to-width aspect ratio of the individual conductors. Wallstrip is distinctly different than a hicroplanar transmission line [14] due to the fact that only one ground reference is necessary.

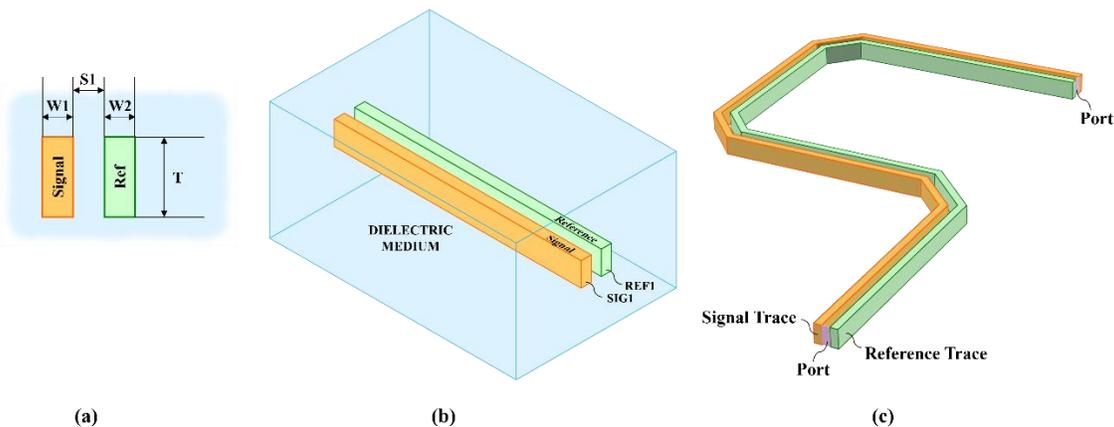


Figure 4: Wallstrip transmission line structure including (a) dimensioned cross-section, (b) isometric view of straight section and (c) example of routed section including trace bends and simulation ports.

Wallstrip is best used in extremely dense, space-constrained layouts to provide massively parallel data transfer where it is impossible or infeasible to include reference planes for the return path. It takes advantage of the fact that these traces are necessarily narrow yet tall to provide adequate return path in adjacent conductors. The reference is defined only as the individual conductor directly adjacent to the signal conductor, and a dedicated return plane above or below the signal is not necessary; furthermore, a return plane would change the Wallstrip line to be a micro- or stripline structure. We have developed Wallstrip out of a realization that dense, multi-conductor die-to-die interfaces can take advantage of thicker dielectric layers on a chiplet package developed with our unique technology. In fact, this transmission line structure is not nearly as capable if the dielectric thickness between metal layers is equal to or less than the metal thickness. The evolution of Wallstrip is shown in Figure 5 where we begin with a traditional microstrip, then cut the reference plane to match the width of the signal conductor. Finally, the two-conductor structure is rotated to be on a single plane within the design. Wallstrip is suitable (and, in fact, required) for use in a lithographic build-up process wherein the RDL metal thickness (T) is at least 1.5 times as

large as the conductor width ($W1$ and $W2$). Furthermore, the dielectric thickness to adjacent layers must be greater than the metal thickness itself.

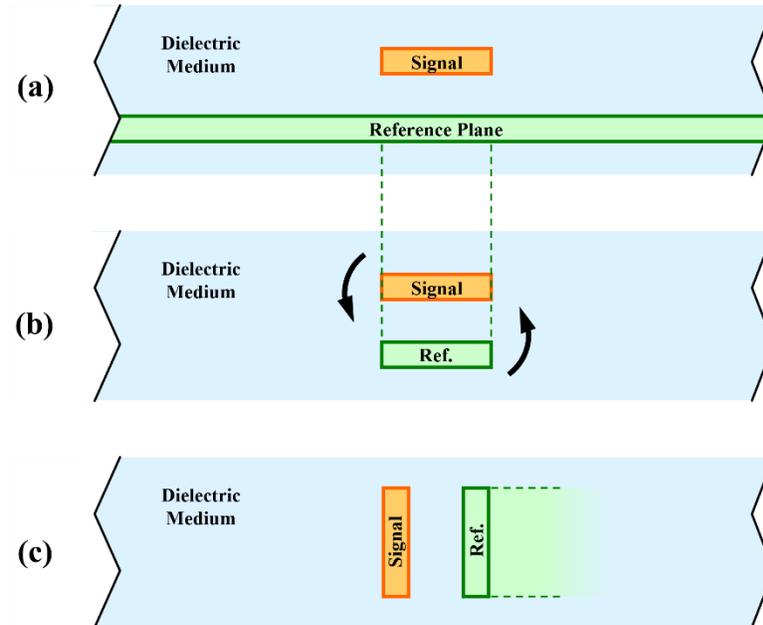


Figure 5: Evolution from microstrip to Wallstrip beginning with (a) traditional microstrip, (b) cutting the reference plane and (c) rotating the entire structure. Note that the reference trace may be of variable width when compared to the signal trace.

Return Current Study

In developing Wallstrip, we first investigate a standard microstrip and its return current distribution. The purpose of this effort is to determine what percentage of the return current density is present in the return plane within a width which is equal to the signal conductor. We begin with a traditional fine-pitch microstrip embedded within a dielectric medium as depicted in Figure 6. This trace is relatively narrow with thin dielectric, resulting in a single-ended characteristic impedance (Z_0) of 54.8Ω at 10 GHz.

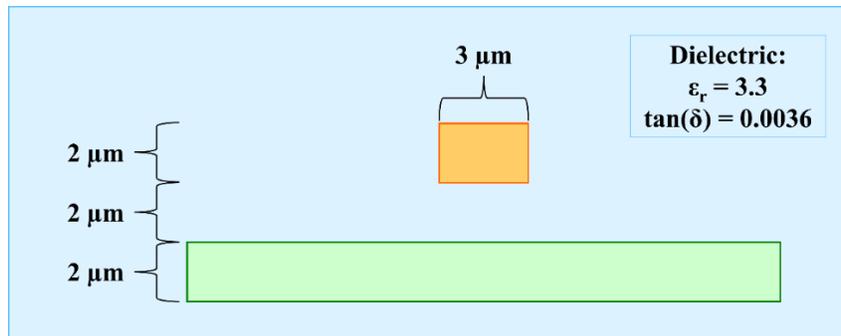


Figure 6: Microstrip geometry as used to investigate return current across the return plane.

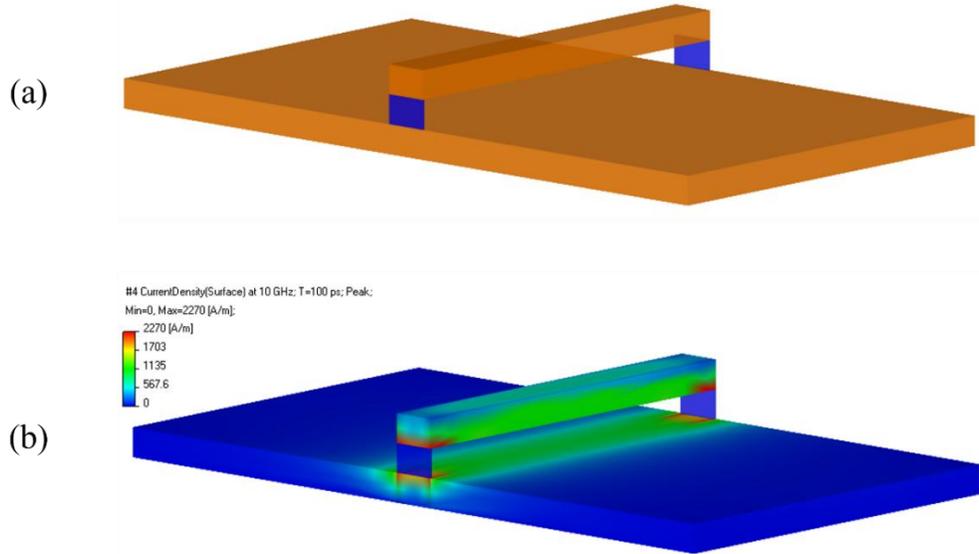


Figure 7: Isometric view of the microstrip line: (a) geometry and (b) field plot of the surface current density magnitude.

Performing a full-wave simulation of this trace, we note the magnitude of the surface current on the return path in Figure 7. Only a small portion of the surface current is present at a distance far from the area directly underneath the signal trace. A numerical analysis shows that 26.2% of the total surface current density is within the region of the return path which sits directly underneath the signal trace, as depicted in Figure 8. Next, we cut the return path to match the width of the signal trace and perform an identical simulation shown in Figure 9.

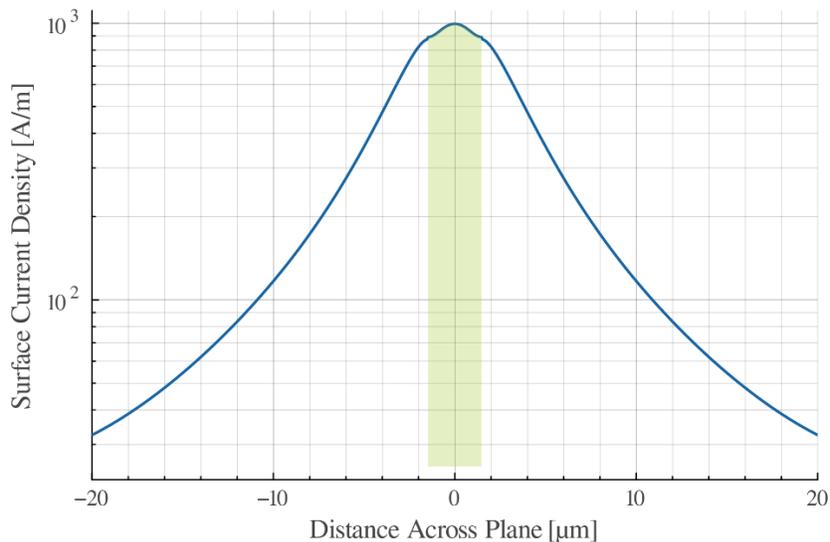


Figure 8: Surface current density in the return path of a microstrip trace with the y-axis plotted in a logarithmic scale. The shaded area represents the current density directly underneath the signal trace.

A common rule-of-thumb recommends that the return path of a signal is at least three times as wide as the signal trace width [15, App. A]. Since we are deliberately violating this rule,

we know that there will be changes to the characteristic impedance of the trace. Z_0 is a function of the per-unit-length inductance and capacitance [16] and thus we expect the impedance to rise when reducing the return path width. This is confirmed as the Z_0 of the trace rises to 68.9Ω when no other geometry is modified. When we compare the field plots of the two cases, it is apparent that the charge density is much higher along the edges of the return path in the latter case. This phenomenon is already well known in signal traces [17, Sec. 3.4.5] but now is showing up in the return path.

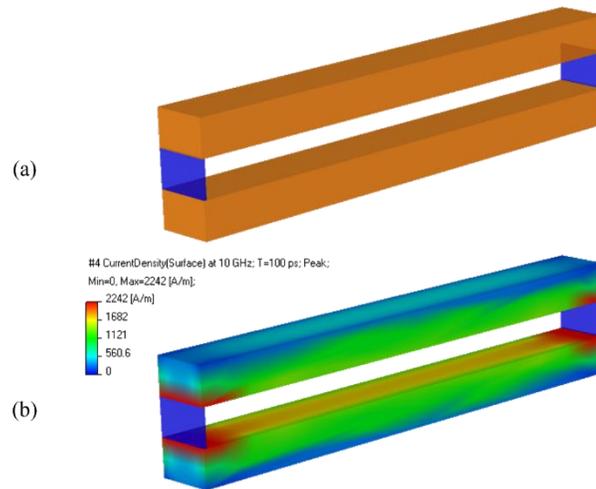


Figure 9: Isometric view of the microstrip line after the return path has been reduced to match the signal trace width: (a) geometry and (b) field plot of the surface current density magnitude.

Material Property Identification

We establish a model-to-hardware correlation for this work using a baseline test vehicle, codenamed Mini Falcon (Figure 10). Although this vehicle employs a different dielectric material set than that intended for the final substrate design, it enables a well-defined correlation flow between simulation and hardware. The measured data from this baseline structure, combined with material property datasets obtained through cavity-resonator methods on bulk samples, provides confidence in the dielectric parameters used for simulation. In addition, the metal-processing steps for the baseline vehicle closely match those expected in the final design, ensuring that the correlation accurately reflects the behavior of the fine-feature geometries required for Wallstrip routing. While the dielectric properties in the Wallstrip simulations differ from those of the baseline test vehicle, the demonstrated agreement between simulation and measurement confirms that the analysis tool reliably predicts the electrical behavior of the small feature sizes relevant to the final organic-substrate implementation.

The test substrate includes many structures such as an IEEE P370 compliant 2X-Thru [18], single-ended and differential transmission lines, three different types of resonant structures (Beatty [19], T- and ring resonators) as well as a resistance measurement structure. We perform scattering parameter measurements of the various structures, which are then post-processed to extract the material property information. Ultimately, the material in our test substrate will not be the final material used in products, thus we are using this step to ensure

our simulation environment can capture the behavior of these small feature-size transmission line structures.

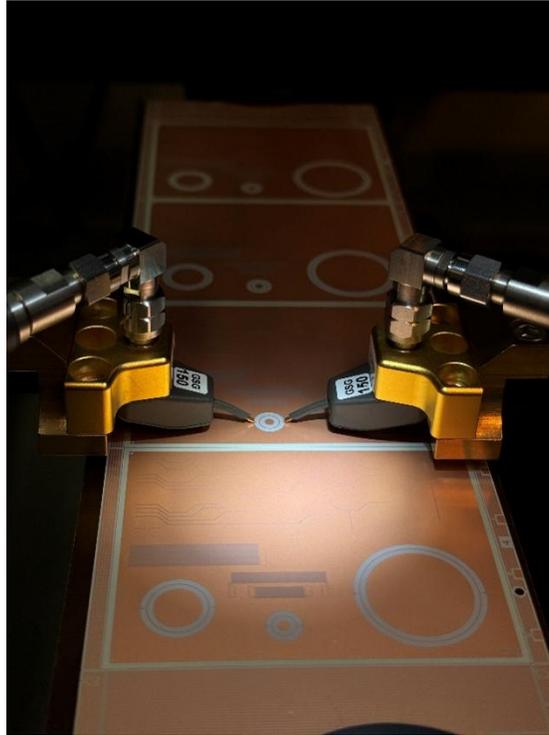


Figure 10: Test substrate, codename Mini-Falcon, measured for material property extraction and simulation/measurement correlation.

The S-parameters are measured for the 2X-Thru, the Beatty resonator, transmission lines with lengths of 1mm, 5mm and 20mm, and three different T-resonator structures for a frequency range of 100 kHz to 20 GHz. The S-parameter quality is verified for all measurements as depicted in Figure 11 where all measurements show acceptable passivity per [18, Sec. 7.3]. Also described in [18] is a list of Fixture Electrical Requirements (FERs) which should be used to evaluate the quality of the test structures. Figure 11 depicts the impedance variation of FER5 at the midpoint of the 2X-Thru. While the T-resonators present a large discontinuity which affects the TDR at the mid-point, all of the transmission line structures and the Beatty resonator structure show excellent agreement with the 2X-Thru. All measurements easily meet the Class A limit of 2.5% impedance variation; in fact, the worst-case impedance variation is only 0.19%. The measurements also meet the -10 dB Class A FER1 insertion loss requirement up to 20 GHz and meet FER2's return loss requirement at Class B up to 15 GHz. Ultimately, these measurements are of sufficient quality to proceed with material property extraction and simulation tool validation.

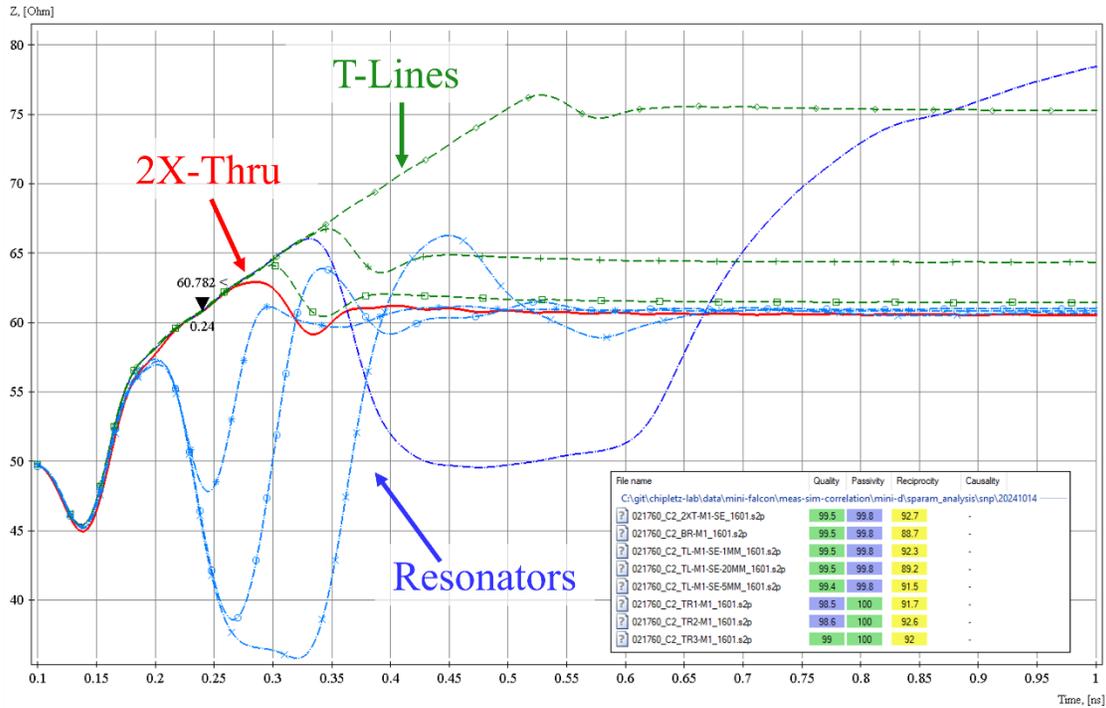


Figure 11: FER5 verification of measured S-parameters from test vehicle.

Upon validating that the measurement quality is sufficient, the generalized modal S-parameters (GMS-parameters) are computed for the transmission lines as described in [20]. By modifying the critical material properties used in simulation, a high-quality match is made between the simulated and measured data. Figure 12 shows the phase delay of the transmission line in ns/mm as well as the attenuation of the transmission line in dB/mm after fitting all of the material properties. For this test substrate, the conductor model used a Huray-Bracken roughness model [21] with RF of 8.33 and SR of 0.03 μm , corresponding to an R_z roughness factor of 0.5 μm . The microstrip dielectric uses an infinite-pole model [17] with ϵ_r of 2.6 and $\tan|\delta|$ of 0.005 at 10 GHz. This test substrate uses trace widths which are slightly larger than the minimum traces required for the die-to-die interface, but ultimately the measurement effort proves that simulations of sub-10 μm trace widths accurately represent real-world measured data. This affords confidence in the results of further simulation, even when dimensions and material properties may change slightly.

While the simulations which seek to quantify the effect of the Wallstrip transmission line on an organic dielectric do use different dielectric properties than those which were measured here, we have thus shown that the simulation tool used to perform these analyses is well capable of faithfully representing the behavior of the small feature sizes required to route Wallstrip.

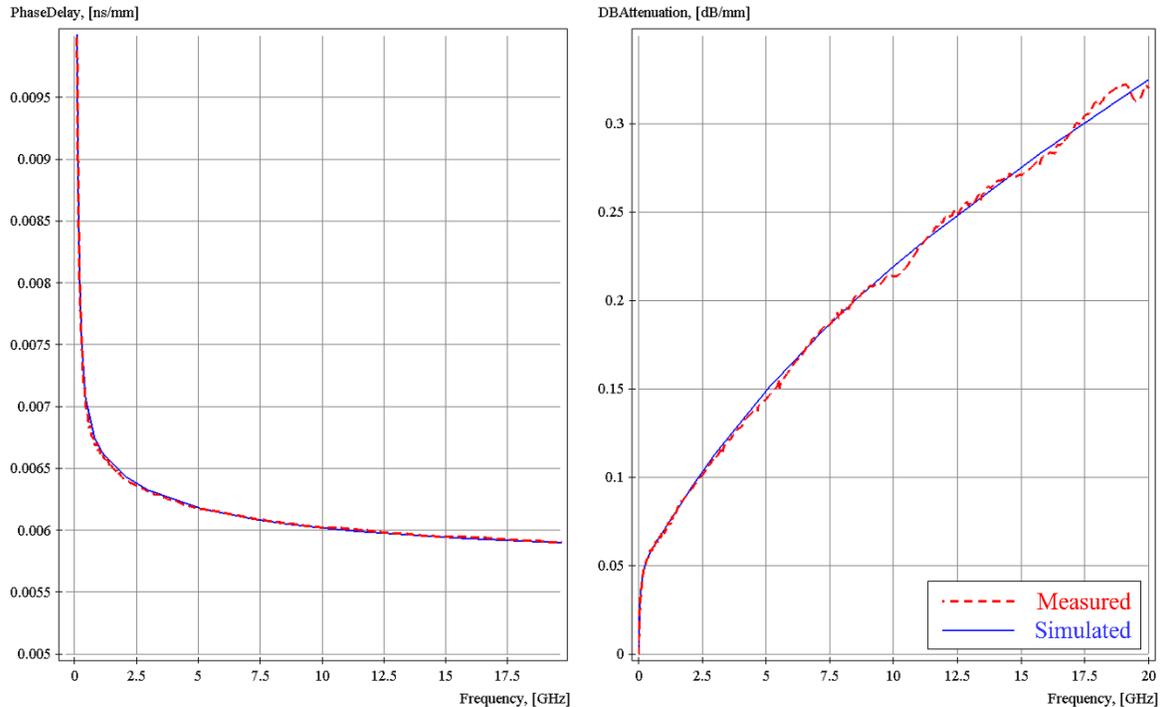


Figure 12: Transmission line properties of the measured GMS-parameters and simulated lines after fitting.

UCIe Transmitter Model

Channel simulations require a realistic transmitter model to properly represent the non-linear effects of a CMOS push-pull driver, similar to that in [10, Sec. 5.3.1]. At present, no UCIe-specific IBIS or IBIS-AMI models are openly available, thus as part of this effort we have developed a unique, parameterized SPICE model representing a full-featured UCIe compliant driver with all electrical parameters in [10].

The transmitter model is designed to implement the -2.2 dB of de-emphasis specified in the UCIe standard, although as this work is intended to directly compare only the channel interconnect, none of the analysis results presented herein will utilize this de-emphasis equalization feature. The architecture and circuit diagram of the driver is presented in Figure 13 where the separate de-emphasis delay path is clearly shown.

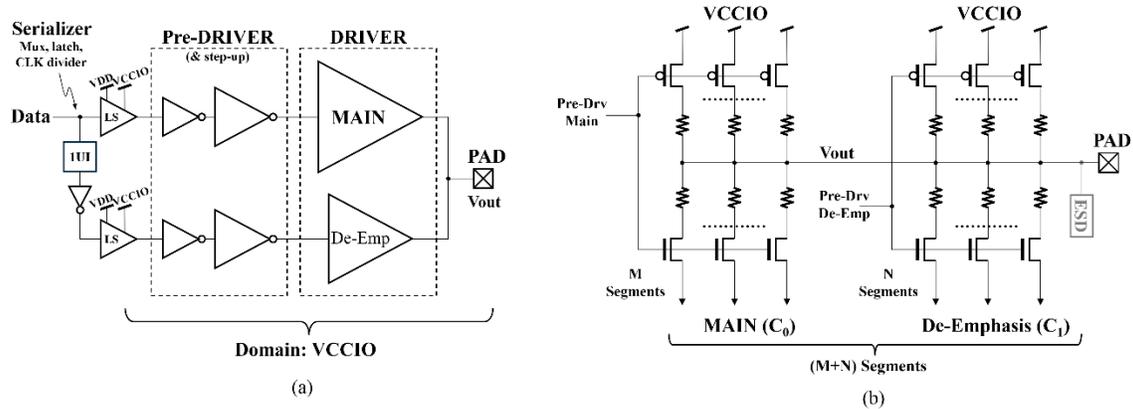


Figure 13: Block diagram (a) and schematic diagram (b) of the custom UCIE transmitter as designed for this effort.

The transmitter model is capable of varying its output impedance in accordance with the specification, but all simulation results are with the nominal 25Ω driver output impedance. The importance of including this output impedance is further depicted in Figure 14 where its effects on the driver current and output voltage can be clearly seen. One of the most critical features of this driver as compared to an idealized voltage source is the effect of non-linearities: these present as both a deviation from the ideal output impedance curve and asymmetry in the voltage crossing. This effect is evident in the eye diagrams and is shown in Figure 15. Without including this non-ideal effect, all channel simulation results would be overly optimistic and not representative of a realistic CMOS driver.

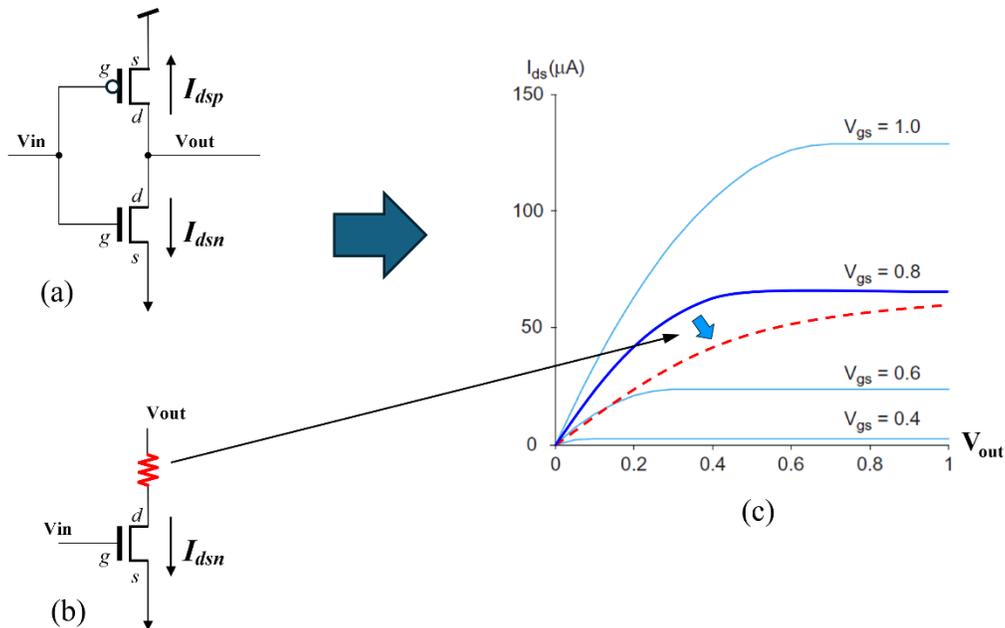


Figure 14: Illustration of an ideal push-pull CMOS driver (a), the addition of an $R_{DS(on)}$ property (b) and its effect on the output impedance of the driver (c) by plotting the drain-source current as it relates to the output voltage [22].

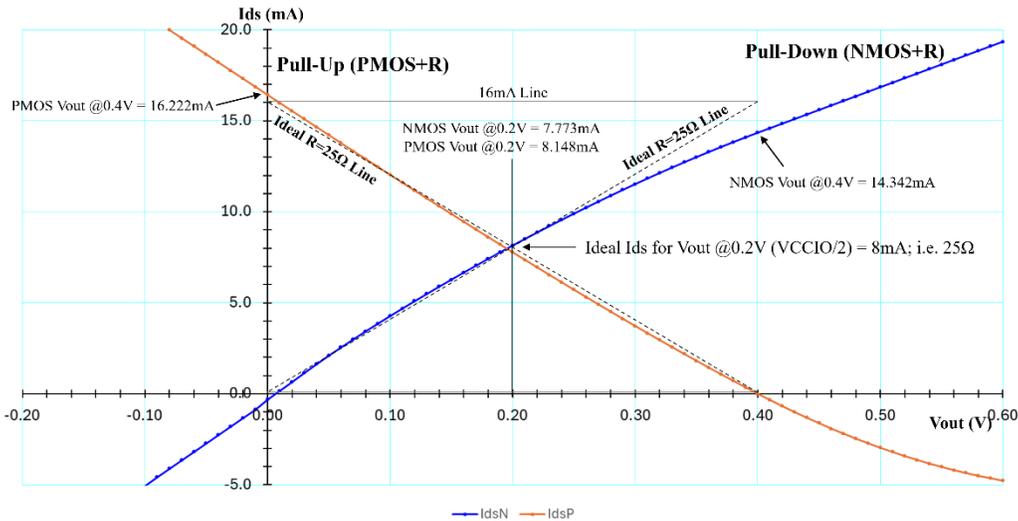


Figure 15: Transmitter pull-up and pull-down curves as designed as part of this effort, showing the non-linear and asymmetrical effects as compared to an ideal 25Ω load line.

UCIe Performance Metrics

UCIe provides varying performance metrics depending upon factors such as data rate and package type. This effort will focus primarily on 32 GT/s and faster within an Advanced Package configuration. Upon designing a channel, the first metric to check prior to running a channel simulation is the VTF as described in [10, Sec. 5.7.1]. VTF is a frequency-domain metric similar to scattering parameters, but the critical distinction is the requirement of termination resistance and capacitance on both the transmit and receive side. Instead of a power ratio as in S-parameters, VTF instead relies on a voltage ratio. VTF loss is analogous to insertion loss and is defined both at the DC point as well as at the Nyquist frequency. The VTF Loss ($L(f)$) over frequency as defined in [10] can be computed with:

$$L(f) = 20 \log_{10} \left| \frac{V_r(f)}{V_s(f)} \right|, \quad \text{Equation 1}$$

where $V_r(f)$ is the voltage at the receiver, and $V_s(f)$ is the voltage at the transmitter. After computing the VTF loss, it can be compared to the mask as shown in Figure 16.

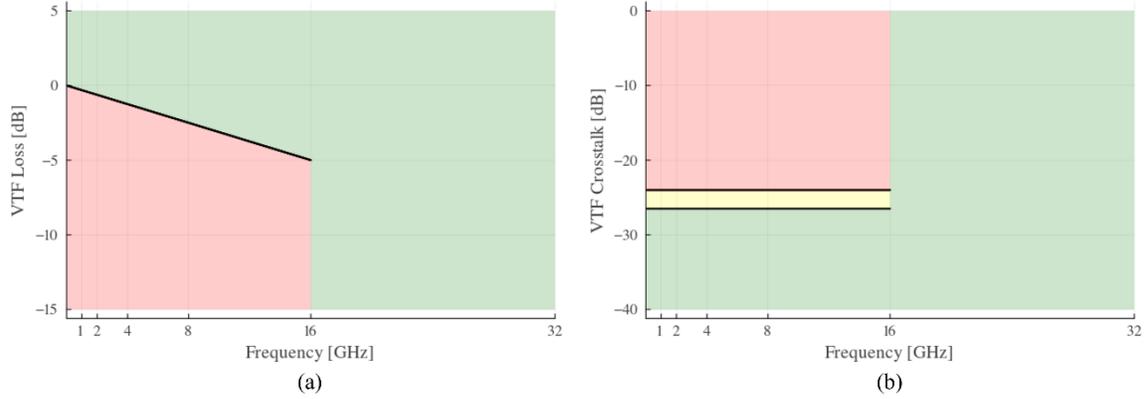


Figure 16: VTF loss (a) and crosstalk (b) masks for a 32 GT/s Advanced Package design. Note the yellow region in the XT plot indicating an area of potential violation depending upon the value of VTF loss at 16 GHz.

The loss at 0 Hz is dependent upon the termination resistances. According to [10], all Advanced Package implementations use high-impedance receiver termination resistance which effectively forces the $L(0)$ value to zero. The VTF Crosstalk (XT) is further defined in [10] which states that 19 aggressor signals must be included to compute the VTF XT for a single victim, specifically:

$$XT(f) = 10 \log_{10} \left(\sum_{i=1}^{19} \left| \frac{V_{ai}(f)}{V_s(f)} \right|^2 \right), \quad \text{Equation 2}$$

in which the conventional crosstalk calculation is inverted. A typical power-sum crosstalk computation would create a stimulus on all aggressors and then measure the effect on the receiver. Instead, this VTF XT computation uses a stimulus on the victim, then measures and sums the effects on all aggressor receivers. Further, VTF XT is solely analogous to far-end XT, ignoring any possible near-end effects present in the system. The limit for VTF loss in a 32 GT/s system is -5 dB at Nyquist. There are two limits present for VTF XT: one absolute limit of -24 dB, and another varying limit based on loss wherein VTF XT shall not exceed 1.5 times the VTF Loss at Nyquist, less an additional 19 dB. Effectively, this means that a system with higher loss will have a more stringent requirement for VTF XT. For example, if a system manages to have an $L(f_N)$ of -1 dB, it must meet the stricter of either -24 dB, or -20.5 dB for its VTF XT at Nyquist (thus defaulting to the -24 dB limit). If, however, a system manages to barely pass the VTF Loss requirement with $L(f_N)$ of -5 dB, it will then be required to meet a VTF XT limit of -26.5 dB at the Nyquist frequency. These limits do not provide any frequency dependence either—meaning the value calculated at the Nyquist frequency carries down to 0 Hz. This loss dependence is critical due to the tight margins for eye width and eye height.

A channel simulation is also required in order to validate performance against an eye mask as depicted in Figure 17. The channel simulation must account for the same nineteen aggressor signals and the resultant crosstalk. Further, the channel simulation is not to account for noise or jitter in the receiver and transmitter models [10].

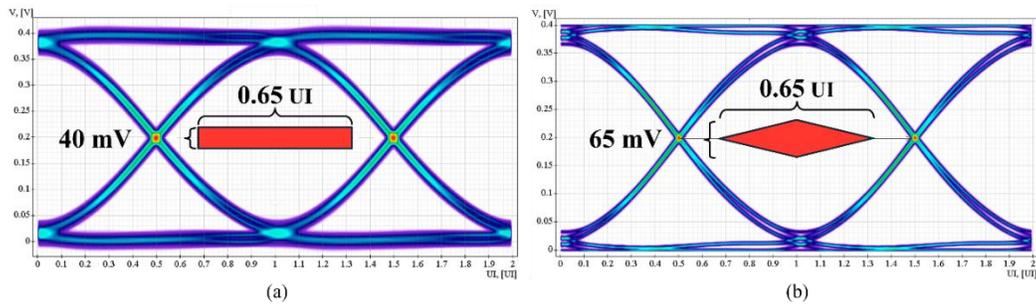


Figure 17: UCIE eye mask limit for all data rates at 24 and 32 GT/s (a) and above 32 GT/s (b).

All eye diagram plots presented within this effort will utilize one of these two masks depending upon the data rate. Every eye diagram will consist of an overlay of all twenty received data signals with the eye mask applied against all data signals together. Additionally, equalization—while allowed—will not be considered in any channel simulations as this effort is directly focused on channel and material comparisons rather than variations in the transmitter and receiver models.

20-Conductor Wallstrip Parallel Interconnect

The Advanced Package specification for UCIE considers a die beachfront width of 388.8 μm [10] for 64 data, 2 clock and 2 sideband signals in each direction, for a total of 136 conductors. We limit the number of routing layers to three and consider only twenty signals as required by the Voltage Transfer Function (VTF) equation in [10, Sec. 5.7.1], thus we must fit these signals within an overall width of 57.2 μm . This allows for the use of 2 μm trace width and 2 μm space between traces as shown in Figure 18. The dielectric, while unspecified in height within this publication, is larger than both the typical 3 μm limit of polyimide-based materials, and the thickness of the metal layers.

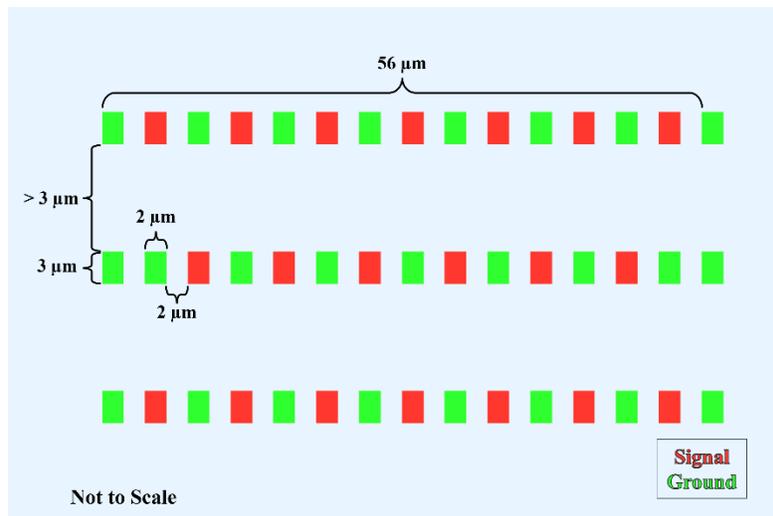


Figure 18: Wallstrip transmission line cross-section for a 20-conductor subset of the UCIE Advanced Package interface routed on three layers and fitting within the maximum channel width requirement.

The Wallstrip layout is fully plane-less: the only reference metal used for the die-to-die transmission lines is in the form of strips rather than including a global plane above or below the structure as is shown in [23, Fig. 3]. We expect the slightly modified material set to have improved metal roughness as compared to the test vehicle which was measured in a previous section, so we configure the transmission line model with an R_q metal roughness of $0.0577 \mu\text{m}$, approximated from an R_z value of $0.2 \mu\text{m}$ [24]. Additionally, the metal conductivity is configured to be $5.1 \times 10^7 \text{ S/m}$. The dielectric is expected to have an ϵ_r of 3.3 and $\tan|\delta|$ of 0.0036 once fabricated.

Silicon Bridge Comparison Geometry

In order to build a representative comparison with typical silicon bridge technology, we configure the 2-D field solver with a layer and trace geometry generally accepted within industry to match a typical UCIE interface which successfully meets the standard requirements at 32 GT/s with a channel length of 1.65mm [25]. The geometry uses SiO_2 as its dielectric material, and copper as its conductor. The conductivity of the copper is set as along with an R_q metal roughness of $0.0577 \mu\text{m}$.

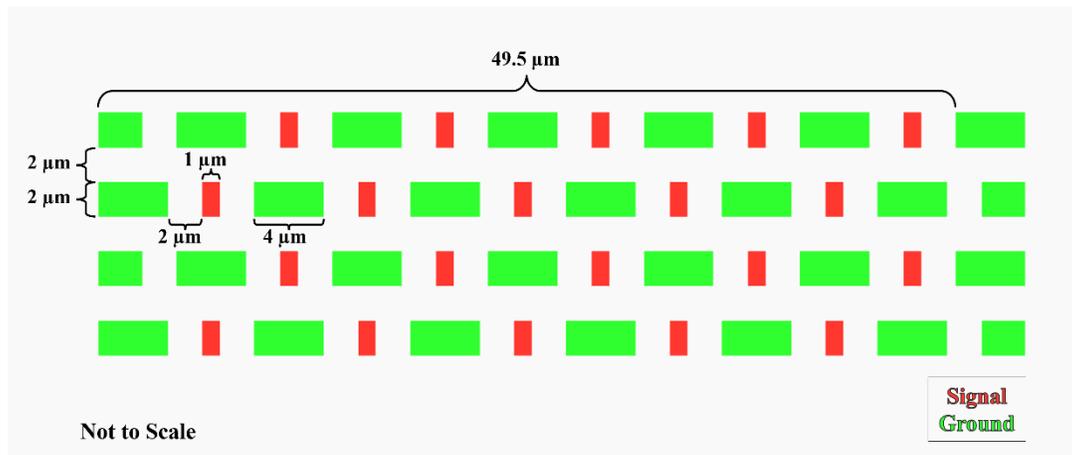


Figure 19: Silicon bridge geometry as-designed to compare with Wallstrip [25].

The VTF loss and crosstalk in this substrate is first matched to nearly replicate the results shown in [25], with our VTF loss matching within 6.5% at 16 GHz and our VTF XT matching within 0.4%. This simulated channel is a minimally passing version which meets the UCIE specification's VTF limits as plotted in Figure 20. This trace configuration does comply with the return path rule noted previously and creates a micro- and stripline structure.

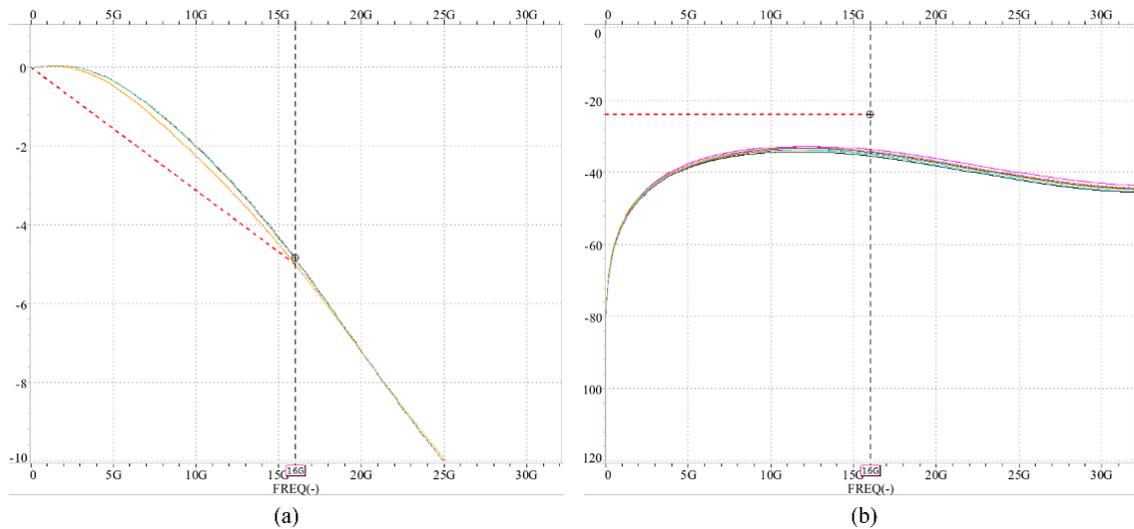


Figure 20: VTF Loss (a) and XT (b) for a 1.65mm length channel in a silicon bridge substrate matching the results from [25]. Vertical axis units in dB.

Results

Upon designing the Wallstrip substrate layout in accordance with manufacturing specifications and then validating that the performance of the silicon bridge substrate matches the behavior in the literature, many different channel simulation experiments were contrived and simulated. The goal of these channel simulations is to explore the design space with respect to various channel length and data rate tradeoffs.

Channel Performance: 2mm Length at 32 GT/s

The first comparison will explore the channel performance with a die-to-die interface length of 2mm. The results of this study will inform the system design team regarding the benefits and downsides of pursuing a silicon bridge style package design over a Wallstrip design with an organic redistribution layer substrate.

While the Wallstrip design nearly violates the VTF Crosstalk limit of -24 dB, its channel eye opening exceeds that of the traditional silicon bridge design owing specifically to the significantly better insertion loss. The silicon bridge design has a non-trivial improvement in crosstalk performance compared to Wallstrip, and only slightly fails the loss requirement. The VTF loss and crosstalk for both the silicon bridge and Wallstrip designs are plotted in Figure 21.

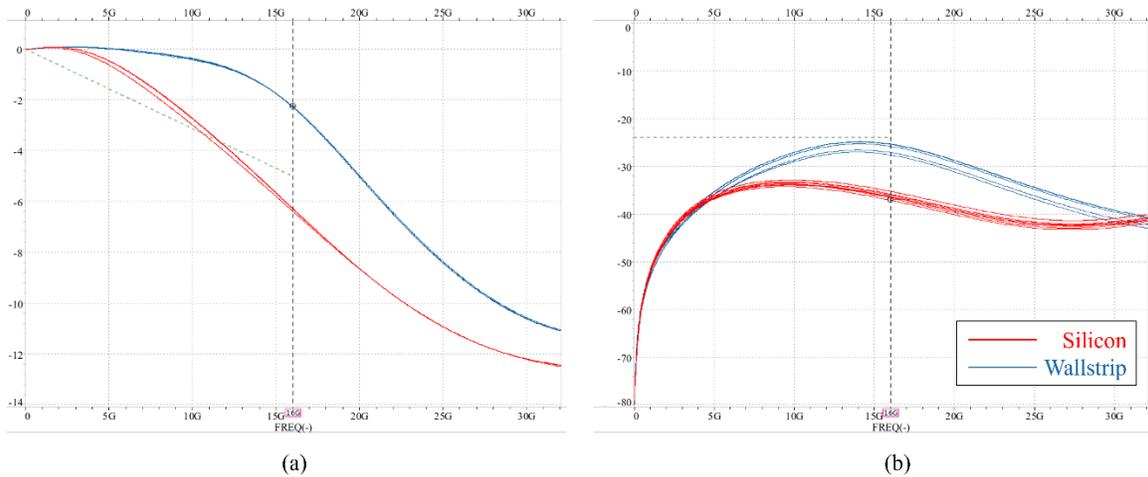


Figure 21: VTF Loss (a) and XT (b) for the Wallstrip and silicon bridge configurations with 2mm channel length. Vertical axis units in dB; UCIe specification limit shown with dotted line.

Both design options pass the UCIe specification’s minimum channel eye mask, with the Wallstrip design affording 3.72 ps of margin above the requirement and the silicon bridge design providing 2.21 ps of margin, depicted in Figure 22. As will be explored in a later section, despite both designs passing the requirements, the Wallstrip gives more flexibility to increase data rates beyond 32 GT/s as future product requirements also increase.

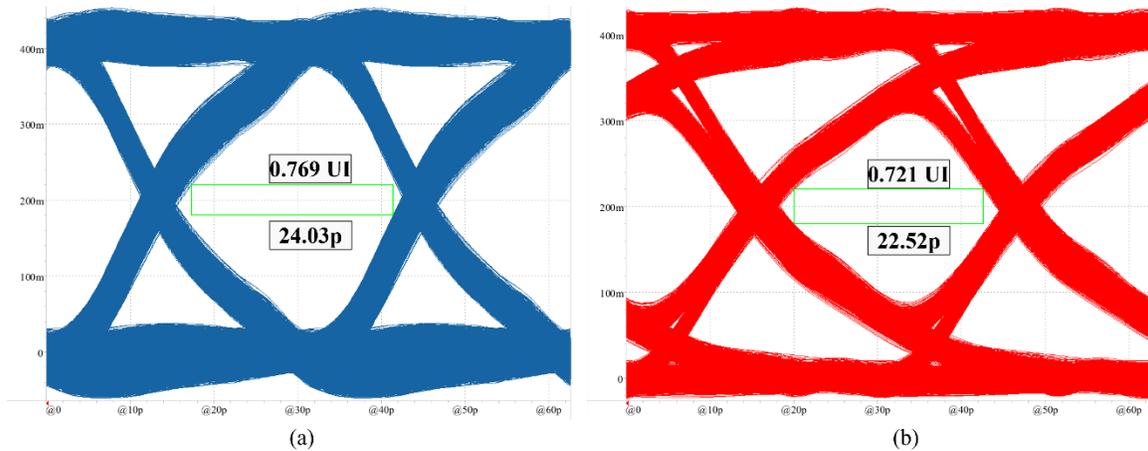


Figure 22: Channel eye diagrams for the Wallstrip (a) and silicon bridge (b) configurations with 2mm channel length at 32 GT/s. A rectangular aperture with 40mV height and maximum sustainable width is shown; vertical axis units in V. Both Wallstrip and silicon bridge pass the minimum eye aperture width requirement of 0.65 UI.

Channel Performance: Maximum Length at 32 GT/s

While both the silicon bridge and the Wallstrip pass the minimum eye mask requirement at 32 GT/s with a 2mm channel length, we next will determine the maximum length channel which will sustain 32 GT/s data rates in both substrate configurations. After exploring the results of the previous section, the system designer may also be interested in the maximum allowable length for a die-to-die interconnect. Understanding this limitation will help in

determining the quantity and placement of individual chiplet dies on the package substrate, as well as driving potential substrate extents.

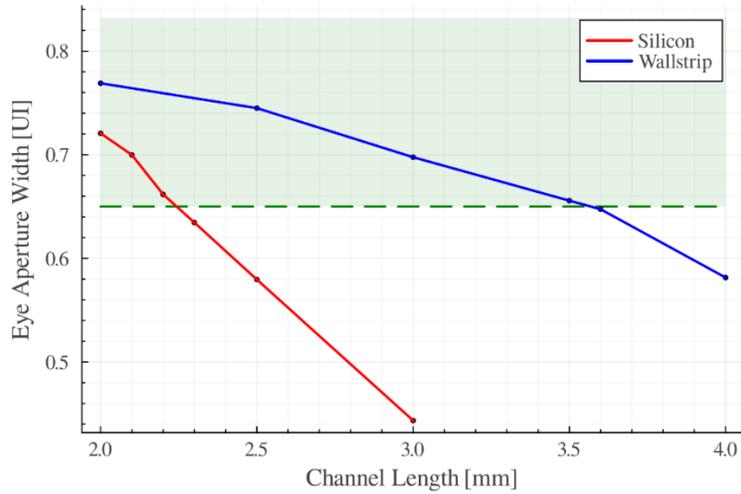


Figure 23: Plot comparing the Wallstrip and silicon bridge configurations describing the width of a 40 mV eye aperture across variation in channel length.

The width of a rectangular eye aperture with a 40 mV height is plotted in Figure 23 and this aperture width tends to decrease linearly as the channel length is increased. Notice that the silicon bridge shows a least-squares relationship of -0.279 UI/mm of channel length, whereas the rate of decline in the Wallstrip substrate is three times slower at -0.091 UI/mm. The end result of this improvement is that the silicon bridge channel cannot exceed 2.2mm of length while the Wallstrip channel can route up to 3.5mm before failing the UCIE eye mask requirement as depicted in Figure 24.

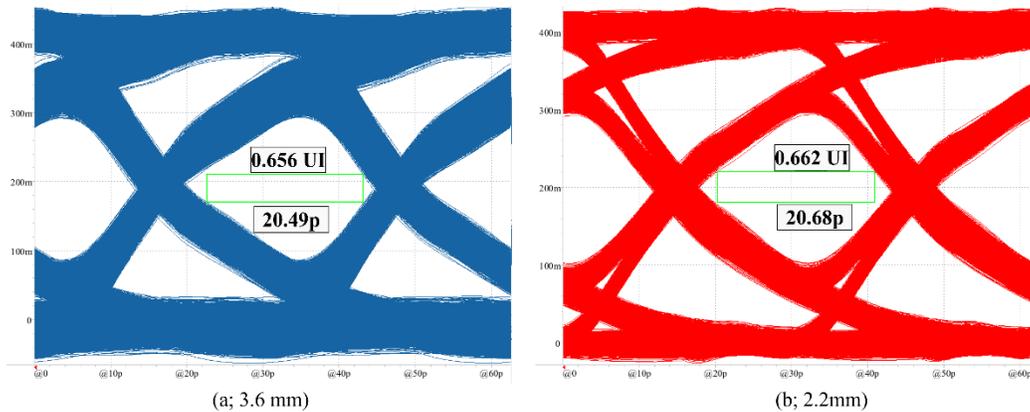


Figure 24: Minimally passing channel eye diagrams for the Wallstrip (a) and silicon bridge (b) configurations at 32 GT/s with respective maximum channel lengths.

Channel Performance: Maximum Data Rate

Finally, we investigate the maximum data rate possible with both the Wallstrip and the silicon bridge substrates while keeping the channel length fixed at 2mm. While all other cases have kept the transmit and receive termination capacitance fixed at 250fF and 200fF

respectively and in accordance with the specification for 32 GT/s systems, this study will reduce both the transmitter and receiver termination capacitance to 180 fF as is specified for 64 GT/s systems [10, Sec. 5.3.2, 5.4.2]. As we are now exploring data rates above 32 GT/s, the channel eye metric will use a diamond-shaped mask rather than rectangle as was described in a previous section and depicted in Figure 17.

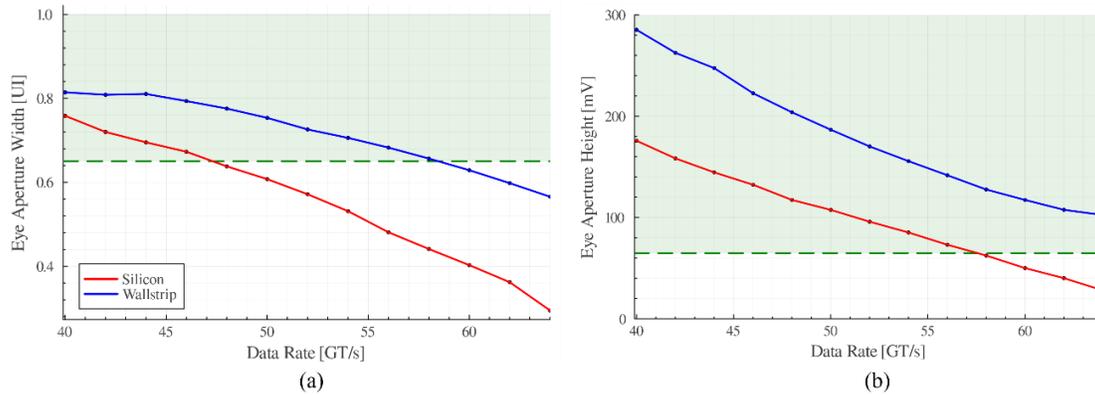


Figure 25: Comparison of Silicon Bridge and Wallstrip configurations with 2mm channel length exploring eye width (a) and eye height (b) across a range of data rates; the acceptable limits are highlighted in green.

The plots in Figure 25 show that both Wallstrip and silicon bridge have an easier time meeting the eye height requirement, with a least-squares fit of -8.2 mV per GT/s and -6.5 mV per GT/s respectively; in-fact, the silicon bridge has a somewhat slower rate of decline in its eye aperture height as data rates increase, although the Wallstrip eye height is consistently between 65 and 100 mV larger than the silicon bridge across the range of data. The same relationship does not hold for the eye aperture width, however. In this case, the silicon bridge cannot exceed a data rate of 46 GT/s without violating the eye mask requirement. The Wallstrip substrate, on the other hand, does not fail the eye mask limit until speeds above 58 GT/s. It must be noted that these results compare the interconnect only, using identical transceiver and receiver models and do not account for any equalization in either of said models. Figure 26 shows the overlaid eye diagrams for both Silicon and Wallstrip at the highest respective data rates which still pass the minimum eye mask limit. Figure 25

Recall that the silicon bridge compared within this work requires one additional layer than the Wallstrip for routing, but the overall width of the route channel is 12% narrower as shown in Figure 18 and Figure 19. While both channel configurations carry the same 20 data bits, but at different maximum data rates, we may then contrive a minor modification to the standard measure of bandwidth density [26] by exploring the total throughput in a given multi-layer routing area. The Bandwidth Density per Routing Area (BDRA) measure is defined as,

$$\text{BDRA} = \frac{\text{Data Rate [bits/s]} \times \text{Signal Count}}{\text{Routing Width [mm]} \times \text{Layer Count}} \quad \text{Equation 3}$$

Instead of focusing on die shoreline width, the BRDA instead considers the overall impact to substrate routing resources where any additional layer area used to route die-to-die data interfaces could have instead been used to route additional power distribution area or the entire layer even removed to reduce the overall cost of the chiplet package.

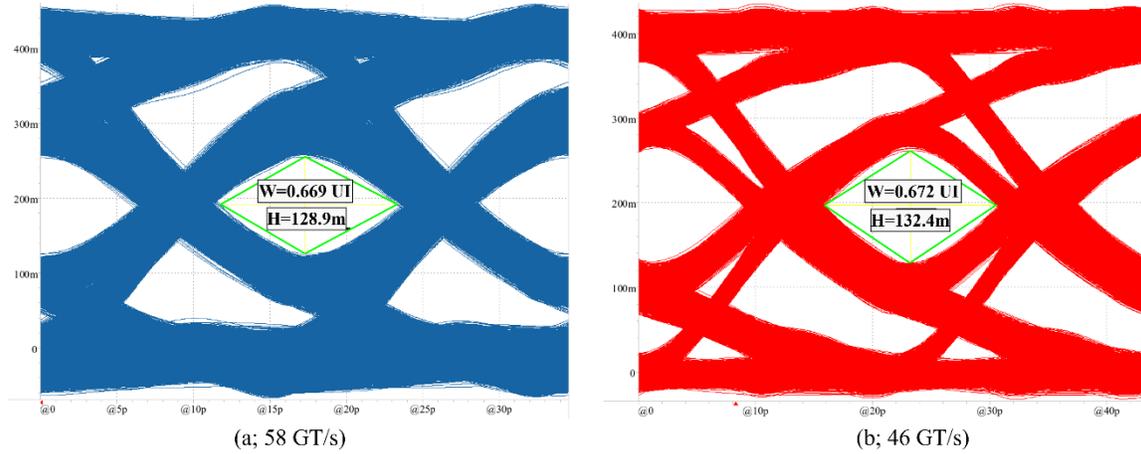


Figure 26: Minimally passing channel eye diagrams for the Wallstrip (a) and Silicon Bridge (b) configurations with 2mm channel length at respective maximum data rates.

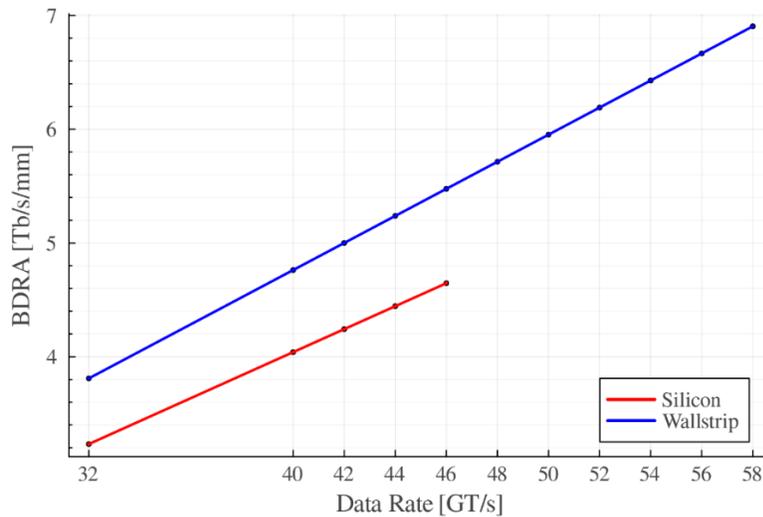


Figure 27: The Bandwidth Density per Routing Area (BDRA) of a silicon bridge as compared to the Wallstrip substrate for a 2mm channel length, provided for all data rates which pass the UCIE minimum eye mask specification.

The maximum BDRA of the 2 mm Wallstrip channel is 6.9 Tb/s/mm which noticeably exceeds the typical die shoreline bandwidth density of modern devices [27] thus ensuring that a Wallstrip substrate will provide adequate overhead to support increased die bandwidth densities as chiplet die technology improves. This also exceeds the 4.6 Tb/s/mm capable with a silicon bridge, thus providing a 50% improvement in overall bandwidth density of the package substrate. Figure 27 depicts a comparison of BDRA between both

silicon bridge and Wallstrip, showing that Wallstrip has a higher rate of BDRA increase over frequency as well as a higher overall magnitude.

Conclusion

In this work, we have explored some of the performance metrics and implementation details used by the most common die-to-die chiplet interface standards. We have then introduced the Wallstrip transmission line structure, which allows for denser routing than a traditional microstrip or stripline leading to adequate crosstalk and improved insertion loss. The effort has shown the viability of today's 2-D field solver simulation tools to match the measured performance of small feature sizes such as those used in today's chiplet package technology. Next, we reviewed some of the critical requirements of a UCIE interface and presented the design of a vendor-agnostic SPICE model capable of representing a UCIE transmitter, including its de-emphasis equalization feature.

Upon introducing the specific organic substrate featuring a Wallstrip transmission line, we have explored said substrate compared to a more conventional silicon bridge substrate used in UCIE. We have shown that, despite silicon bridge featuring better crosstalk, the lower insertion loss in the Wallstrip substrate leads to larger eye margin when implementing a 2mm length UCIE channel at 32 GT/s. Next, we showed that the Wallstrip substrate can achieve up to a 3.5mm channel length while still passing the UCIE eye mask for 32 GT/s, a 59% longer channel than the silicon bridge can support. We have also demonstrated the ability of this Wallstrip channel to operate at much higher data rates than the silicon bridge, supporting 58 GT/s with a 2mm channel length. Finally, we have devised a modification to the standard bandwidth density metric, presenting the Bandwidth Density per Routing Area (BDRA) and showing that a Wallstrip topology transmission line in an organic substrate is capable of a 50% improvement in BDRA. This will improve flexibility and allow the heterogeneous system architect to take advantage of various improvements such as reducing drive strength (thereby improving energy efficiency), lengthening die-to-die channels for optimal assembly, or utilizing the same package substrate design for next generation silicon improvements.

Future Work

Despite the knowledge gained from this effort, there is still much to be done in designing the optimal Wallstrip die-to-die interface in an organic chiplet substrate. The next iteration of this work will focus on the critical area of die bump-out patterns and via geometries. These 3-D structures serve as large impedance discontinuities and since the electromagnetic field patterns are significantly perturbed, they contribute the most significant effects to channel crosstalk. Next steps will also explore the effects of equalization, both on the transmitter and the receiver side. While we have already prepared a transmitter model capable of implementing de-emphasis, we also plan to implement the feed-forward equalization (FFE) added in Revision 3 of the UCIE specification. Additionally, the receiver is expected to implement continuous time linear equalization (CTLE) in future work. Finally, additional work in this area will develop further physical test and measurement structures which will allow us to refine our material properties data set for more accurate simulation models. Ultimately, this will allow a Wallstrip-based chiplet package substrate to meet the current 64 GT/s data rate specified in UCIE as well

as position this organic substrate to exceed future iterations of this and other die-to-die interface standards.

References

- [1] “Heterogeneous Integration Roadmap - IEEE Electronics Packaging Society.” Accessed: Dec. 03, 2025. [Online]. Available: <https://eps.ieee.org/technology/heterogeneous-integration-roadmap.html>
- [2] J. H. Lau, *Chiplet Design and Heterogeneous Integration Packaging*. Singapore: Springer Nature Singapore, 2023. doi: 10.1007/978-981-19-9917-8.
- [3] Z. Yang *et al.*, “Challenges and Opportunities to Enable Large-Scale Computing via Heterogeneous Chiplets,” Mar. 04, 2024, *arXiv*: arXiv:2311.16417. doi: 10.48550/arXiv.2311.16417.
- [4] H. Peng, S. Davidson, R. Shi, S. L. Song, and M. Taylor, “Chiplet Cloud: Building AI Supercomputers for Serving Large Generative Language Models,” May 20, 2024, *arXiv*: arXiv:2307.02666. doi: 10.48550/arXiv.2307.02666.
- [5] C. A. Harper, Ed., *Electronic packaging and interconnection handbook*. New York: McGraw-Hill, Inc, 1991.
- [6] R. Kruger, “Chiplet Design Considerations,” *Semiconductor Engineering*. Accessed: Oct. 08, 2025. [Online]. Available: <https://semiengineering.com/chiplet-design-considerations/>
- [7] “High Bandwidth Memory (HBM4) DRAM | JEDEC.” Accessed: Oct. 06, 2025. [Online]. Available: <https://www.jedec.org/standards-documents/docs/jesd270-4>
- [8] *Bunch of Wires (BoW) PHY Specification*, Mar. 01, 2023.
- [9] chipsalliance, “Advanced Interface Bus (AIB) Specification Revision 2.0.3,” GitHub. Accessed: Oct. 13, 2025. [Online]. Available: https://github.com/chipsalliance/AIB-specification/blob/096a97006a1e46dbab51b4e16ca99e6193dd16cb/AIB_Specification%202_0.pdf
- [10] *Universal Chiplet Interconnect Express™ (UCIe™) Specification*, Aug. 05, 2025.
- [11] *Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps, 11G+ bps, 25G+ bps, 56G+ bps and 112G+ bps I/O*, May 05, 2022.
- [12] I. I. S. Prakash T, and Y. Jagadeeswari, “Innovative Layout Optimization Methodology and Via Routing Pattern to Enable UCIe-36Gbps in Organic Interposer,” presented at the DesignCon 2025, 2025.
- [13] B. C. Wadell, *Transmission line design handbook*. in The Artech House microwave library. Boston, Mass.: Artech House, 1991.
- [14] S. T. Todd, X. T. Huang, J. E. Bowers, and N. C. MacDonald, “Fabrication, Modeling, and Characterization of High-Aspect-Ratio Coplanar Waveguide,” *IEEE Trans. Microw. Theory Tech.*, p. 5639066, Dec. 2010, doi: 10.1109/TMTT.2010.2086531.
- [15] E. Bogatin, *Signal Integrity - Simplified*. Upper Saddle River, NJ: Prentice Hall, 2008.
- [16] D. M. Pozar, *Microwave Engineering*, 4th ed. Hoboken, NJ: John Wiley & Sons, Inc, 2012.
- [17] S. H. Hall and H. L. Heck, *Advanced signal integrity for high-speed digital designs*. Hoboken, N.J: Wiley IEEE, 2009.

- [18] “IEEE Draft Standard for Electrical Characterization of Printed Circuit Board and Related Interconnects at Frequencies up to 50GHz,” *P370D8 July 2020*, pp. 1–150, 2020.
- [19] “Resonant Test Structures: Primer and Signal Integrity Applications | 2016-09-17 | Signal Integrity Journal.” Accessed: Nov. 07, 2025. [Online]. Available: <https://www.signalintegrityjournal.com/articles/194-resonant-test-structures-primer-and-signal-integrity-applications>
- [20] Y. Shlepnev, A. Neves, T. Dagostino, and S. McMorrow, “Practical identification of dispersive dielectric models with generalized modal S-parameters for analysis of interconnects in 6-100 Gb/s applications,” presented at the DesignCon 2010, 2010.
- [21] E. Bracken, “A Causal Huray Model for Surface Roughness,” presented at the DesignCon 2012,
- [22] N. H. E. Weste and D. M. Harris, *CMOS VLSI design: a circuits and systems perspective*, 4. ed. Boston, Mass.: Addison-Wesley, 2011.
- [23] Y. Uematsu, N. Ushifusa, and H. Onozeki, “Electrical Transmission Properties of HBM Interface on 2.1-D System in Package Using Organic Interposer,” in *2017 IEEE 67th Electronic Components and Technology Conference (ECTC)*, Orlando, FL, USA: IEEE, May 2017, pp. 1943–1949. doi: 10.1109/ECTC.2017.34.
- [24] “Simonovich-Cannonball Conductor Roughness Model Demystified,” Bert Simonovich’s Design Notes. Accessed: Nov. 19, 2025. [Online]. Available: <https://blog.lamsimenterprises.com/2019/03/29/cannonball-huray-model-demystified/>
- [25] S. Li, M.-S. Lin, W.-C. Chen, and C.-C. Tsai, “High-Bandwidth Chiplet Interconnects for Advanced Packaging Technologies in AI/ML Applications: Challenges and Solutions,” *IEEE Open J. Solid-State Circuits Soc.*, vol. 4, pp. 351–364, 2024, doi: 10.1109/OJSSCS.2024.3506694.
- [26] J. M. Wilson, W. J. Turner, and J. W. Poulton, “Ground-Referenced Single-Ended Signaling: Applications for High-Density, Short-Haul Communication Systems,” *IEEE Solid-State Circuits Mag.*, vol. 11, no. 2, pp. 54–68, 2019, doi: 10.1109/MSSC.2019.2910618.
- [27] M. Swaminathan, M. Kathaperumal, K. Moon, H. Sharma, P. Murali, and S. Ravichandran, “Materials for heterogeneous integration,” *MRS Bull.*, vol. 46, no. 10, pp. 967–977, Oct. 2021, doi: 10.1557/s43577-021-00212-2.

Acknowledgements

The authors are indebted to all of our colleagues at Chipletz who have provided assistance, information, and feedback both during the process of designing the Wallstrip interface as well as creating this publication. In particular, we thank Siddharth Ravichandran for sharing his knowledge in the package substrate technology space and for enduring a near-endless stream of inquiries. We also thank our manufacturing partners for providing high-quality substrates for use in developing this technology. Finally, we thank the executive staff at Chipletz for providing the opportunity to explore this exciting new paradigm in advanced package signal integrity.