

Universal Chiplet Interconnect Express (UCIe): An Open Industry Standard for Innovations With Chiplets at Package Level

Debendra Das Sharma¹, *Senior Member, IEEE*, Gerald Pasdast, Zhiguo Qian, *Senior Member, IEEE*,
and Kemal Aygün², *Fellow, IEEE*
(Invited Paper)

Abstract—Universal Chiplet Interconnect Express (UCIe) is an open industry standard interconnect for developing an open chiplet ecosystem, where chiplets from any supplier can be packaged anywhere in an interoperable manner. This article delves into the architectural, circuit, channel, and packaging aspects that we developed that has been adopted in the UCIe 1.0 Specification. We present our results based on our channel and circuit implementation studies.

Index Terms—Accelerator, availability, chiplet, compute express link, co-packaged optics, memory expansion, packaging, Peripheral Component Interconnect (PCI) Express, pooling, power-efficiency, rack scale architecture, reliability.

I. INTRODUCTION

GORDON Moore predicted the “Day of Reckoning” in his seminal paper where he posited “Moore’s law” [1]: “It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected.” Today, we are past that inflection point. On-package integration of multiple dies has been widely deployed across the semiconductor industry, including mainstream volume central processing units (CPUs) and general purpose-graphics processor units (GP-GPUs) [2].

There are many drivers for on-package chiplets. Overcoming reticle limitations to deliver performance/functionality and yield challenges with larger dies is a primary reason most companies have their own proprietary solutions.

Lowering the overall portfolio cost with a time-to-market advantage would be a compelling driver for deploying chiplets. For example, the compute cores shown in Fig. 1 [3] can be implemented in an advanced process node to deliver leadership power-efficient performance whereas the fabric functionality comprehending memory and input/output (I/O) controller functions may be reused from a design already deployed in

Manuscript received 28 July 2022; revised 4 September 2022; accepted 8 September 2022. Date of publication 16 September 2022; date of current version 7 October 2022. Recommended for publication by Associate Editor W. T. Beyene upon evaluation of reviewers’ comments. (Corresponding author: Debendra Das Sharma.)

Debendra Das Sharma and Gerald Pasdast are with Intel Corporation, Santa Clara, CA 95052 USA (e-mail: debendra.das.sharma@intel.com).

Zhiguo Qian and Kemal Aygün are with Intel Corporation, Chandler, AZ 85226 USA.

This article has supplementary material provided by the authors and color versions of one or more figures available at <https://doi.org/10.1109/TCPMT.2022.3207195>.

Digital Object Identifier 10.1109/TCPMT.2022.3207195

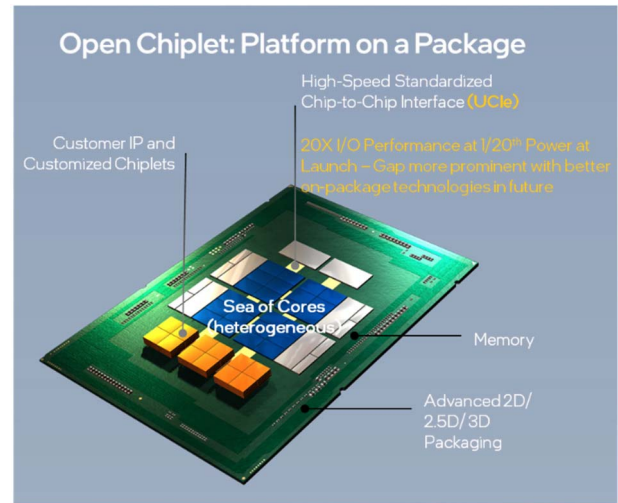


Fig. 1. UCIe-based open chiplet ecosystem: platform on a package.

an established process node. Such partitioning also results in smaller dies, which results in better yields. Furthermore, this approach helps mitigate IP porting costs, which are increasing significantly for the advanced process nodes [3].

Another value of chiplets is the ability to offer bespoke solutions. For example, one can choose different numbers of compute, memory and I/O, and accelerator chiplets depending on the need of a particular product segment. As a result, one does not need to do a different die design for different segments, lowering the design, validation, and product cost.

Universal Chiplet Interconnect Express (UCIe) [4] is an open industry standard interconnect, offering high-bandwidth, low-latency, power-efficient, and cost-effective on-package connectivity between heterogeneous chiplets to address the needs across the compute continuum. UCIe 1.0 Specification [4] comprehends all the layers of the stack [Fig. 2(a)], the only complete specification we are aware of with a well-defined compliance mechanism, targeting heterogeneous integration of components using Peripheral Component Interconnect (PCI)-Express¹ (PCIe¹) [5], [6] and Compute Express Link² (CXL²) [7] protocols and software infrastructure, for ensuring interoperability. This enables a designer to package

¹Registered trademark.

²Traditional trademark.

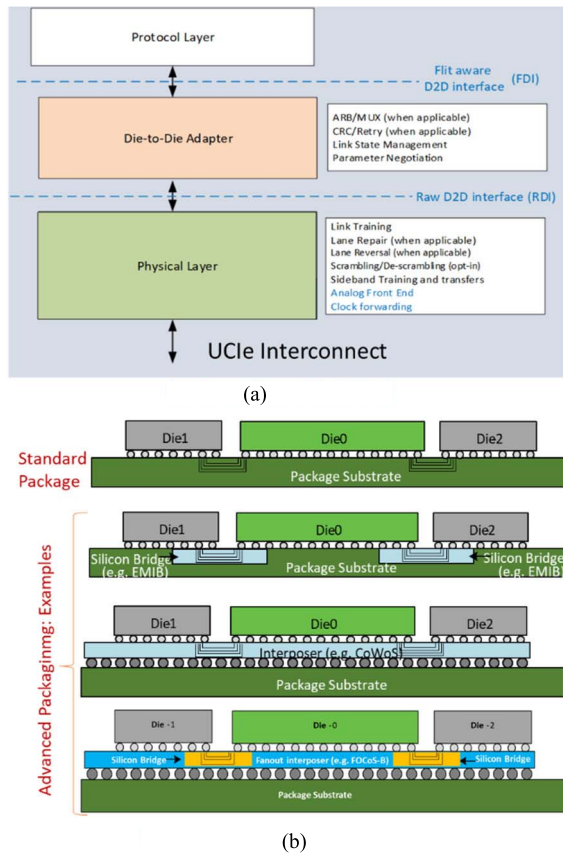


Fig. 2. UCIE: layering approach and different packaging choices. (a) Layering with UCIE. (b) Packaging options: 2-D and 2.5-D.

chips from different sources, including different fabs, using a wide range of packaging technologies. UCIE is an evolution of our prior work which has been implemented in Intel Sapphire Rapids CPU as our proprietary Multi-Die Fabric Interface (MDFI) [2]. The key metrics, characteristics, and simulation methodology delineated in this article have been demonstrated in Sapphire Rapids silicon [2].

This article delves into the requirements and usage models for UCIE in Section II. Our proposed approach is described in Section III, which has been mostly adopted in the UCIE Specification [4]. We present our results in Section IV and conclude in Section V.

II. USAGE MODELS, PACKAGING TECHNOLOGIES, AND PERFORMANCE METRICS TARGETED BY UCIE 1.0 SPECIFICATION

UCIE 1.0 supports two types of packaging, as shown in Fig. 2(b). The standard package (2-D), referred to as UCIE-S, is used for cost-effective performance. The advanced packaging (UCIE-A) is used for power-efficient performance. There are multiple commercially available options, that can deploy both UCIE-S and UCIE-A, some of which are shown in the diagram. UCIE 1.0 specification embraces all types of packaging choices in these categories. The industry-leading performance metrics of UCIE 1.0 specification is summarized in Table I [3].

TABLE I
UCIE 1.0 CHARACTERISTICS AND KEY METRICS

Characteristics / KPIs	Standard Package	Advanced Package	Comments
Characteristics			
Data Rate (GT/s)	4, 8, 12, 16, 24, 32		Lower speeds must be supported -interop (e.g., 4, 8, 12 for 12G device)
Width (each cluster)	16	64	Width degradation in Standard, spare lanes in Advanced
Bump Pitch (um)	100 – 130	25 - 55	Interoperate across bump pitches in each package type across nodes
Channel Reach (mm)	<= 25	<=2	
Target for Key Metrics			
B/W Shoreline (GB/s/mm)	28 – 224	165 – 1317	Conservatively estimated: AP: 45u for AP; Standard: 110u; Proportionate to data rate (4G – 32G)
B/W Density (GB/s/mm ²)	22-125	188-1350	
Power Efficiency target (pJ/b)	0.5	0.25	
Low-power entry/exit	0.5ns <=16G, 0.5-1ns >=24G		Power savings estimated at >= 85%
Latency (Tx + Rx)	< 2ns		Includes D2D Adapter and PHY (FDI to bump and back)
Reliability (FIT)	0 < FIT (Failure In Time) << 1		FIT: #failures in a billion hours (expecting ~1E-10) w/ CXI Flit Mode

III. OUR PROPOSED APPROACH FOR UCIE

Our approach is a well-specified layered standard, including protocol layer, adapter, and physical layer (PHY). We will first briefly explain these layers, and then focus on the unique circuit architecture and package channel design features to achieve the targeted performance, flexibility, and interoperability.

A. Layers

The PHY is responsible for the electrical signaling, clocking, link training, sideband, circuit architecture, and package interconnect channel.

The die-to-die adapter provides the link state management and parameter negotiation for the chiplets. It guarantees reliable delivery of data through its cyclic redundancy check (CRC) and link level retry mechanism, when enabled. Multiple protocols are supported with its underlying arbitration mechanism. A 256-byte (or 68 byte) flow control unit (FLIT) supports the underlying reliable transfer mechanism.

We map the PCIe and CXL protocols to UCIE natively as those are widely deployed at the board level across all segments of compute. This is done to ensure seamless interoperability by leveraging the existing ecosystem where board components can be brought on package. With PCIe and CXL, system-on-chip (SoC) construction, link management, and security solutions deployed in today’s platform can be seamlessly transported to UCIE.

The usage models addressed by our approach for a die-to-die interconnect such as UCIE are comprehensive: data transfer using direct memory access, software discovery, and error handling, are addressed through PCIe/CXL.io; the memory use cases are handled through CXL.Mem; and caching requirements for applications such as accelerators are addressed with CXL.cache. We also define a “streaming protocol,” which can be used to map any other protocol such as a proprietary symmetric cache coherency protocol (e.g., Ultra Path Interconnect). Our approach also enables the UCIE consortium to innovate new protocols to cover new usage models or enhance existing ones going forward.

We support different data rates, widths, bump-pitches, and channel reach to ensure the widest interoperability feasible, as detailed in Table I. The unit of construction of the interconnect is a cluster which composes of N single-ended, unidirectional, full-duplex data lanes ($N = 16$ for standard

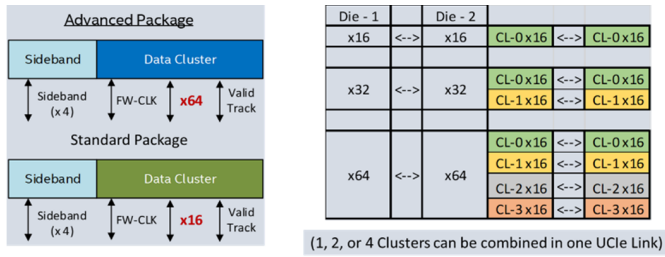


Fig. 3. Cluster(s) in standard and advanced package.

package and $N = 64$ for advanced package), one single-ended lane for valid, one lane for tracking, a differential forwarded clock per direction, and two single-ended lanes per direction for sideband (one for the 800-MHz clock and one for the data). The sideband interface is used for status exchange to facilitate link training in data cluster, register access mechanism even when the link is not trained, and is used for diagnostics. The advanced package supports spare lanes to handle faulty lanes (including clock, valid, and sideband) whereas the standard package supports width degradation to handle failures. Multiple clusters can be aggregated to deliver more performance per link, as shown in Fig. 3.

B. PHY Architecture

We have architected the UCle PHY layer with Integrated Device Manufacturer (IDM) and outsourced semiconductor assembly and test (OSAT) portability in mind. Most circuit components can be built with digital-type circuits such as push-pull transmitters (TXs), digital delay locked loops (DLL) and phase interpolators (PIs), inverter-based front-end receivers (RXs), strong-arm latches for samplers, and inverter-based clock distribution. Some components can be swapped with higher performance standard analog building blocks such as continuous-time amplifiers for the RX analog front end (AFE), on-die termination, inductors, and on-die regulators, which are portable to any modern IDM node.

We propose the same clocking and signaling schemes for both UCle-A and UCle-S. These consist of source-synchronous clocking and matched clock/data delay paths for robust performance in noisy supply environments along with non-return-to-zero (NRZ) signaling as optimal power/performance for the channel specifications to be discussed in the next section. The TX output swing has been specified with a wide operating range of 400–850 mV to allow implementation complexity versus channel power/performance optimization. The RX must be able to resolve an input eye height spec of $40 \text{ mV} \times 47 \text{ ps}$ width at 16 GT/s and $40 \text{ mV} \times 20 \text{ ps}$ at 32 GT/s. Parameter negotiation in early training phase will communicate the swing level to the receiving die and the RX trip point and other calibrations can happen at that time as well.

After training, the link will have approximately 0.5 unit interval (UI) separation between the clock and the data paths. This 0.5 UI target, which makes the link effectively a “matched architecture,” is critical for minimizing the impact of deterministic jitter (DJ) on the link timing performance. During power

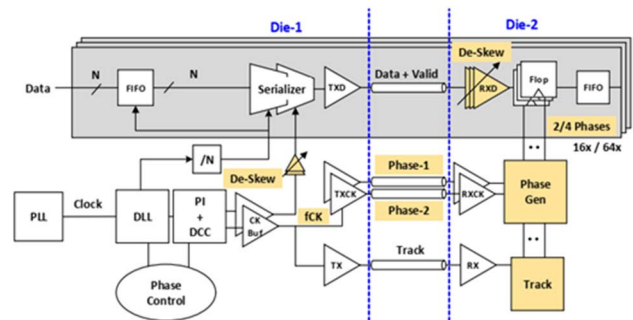


Fig. 4. High-level PHY architecture.

supply droop events, the 0.5 UI delay delta between the clock and data paths gets modulated as a factor of the magnitude of the supply droop multiplied by the alpha factor (i.e., the % of delay change with respect to % of VCC change) of the circuit paths. Typically, the higher the delay delta between the clock and data paths, the more skew develops between the two paths during a supply droop event. This additional skew directly translates to link performance degradation. The proposed 0.5 UI architecture allows for 40–50 mV of supply noise at 16 GT/s. In contrast, a 1.5 or 2.5 UI target would require significantly tighter supply noise specs or high bandwidth tracking mechanisms, which can be power-hungry. The matched architecture on the RX side requires delays through the data and clocking paths to be within 0.1 UI of each other to the sampling flop. De-skew buffers, typically consisting of two CMOS buffers with strength control, are added to each datapath lane for lane-to-lane de-skew calibration. The overall power and noise impact are negligible when taking the higher power supply noise tolerance into account. Fig. 4 shows a high-level overview of our proposed PHY architecture.

Wires from RDI interface [Fig. 2(a)] go through a clock-crossing first-in first-out (FIFO) to retime the signals between the protocol PLL and the PHY PLL domains. The FIFO output is serialized and transmitted with an impedance-compensated TX driver. The clocking path consists of a DLL to generate the necessary references (quadrature or equivalent) for the fine-skew adjuster (PI) and duty cycle corrector (DCC). On the RX die, the data and clock paths to the sampler flop are matched by adding some delay between the data RX AFE and the sampler flop (typically two inverter stages) to match the clock RX AFE + phase gen/clock distribution delays to the sampler flop.

Two phases of clocks are forwarded as even and odd clocks. For 4, 8, 12, and 16 GT/s, the two clocks are forwarded as 90° and 270° clocks running at half the data rate (e.g., 2 GHz for 4 GT/s and 4 GHz for 8 GT/s). This is with respect to 0° used for data on the transmit side, thus producing the 0.5 UI separation needed between the clock and data paths to the sampler. Both edges of the differential forwarded clock are used to sample at the RX, referred to as two-way interleaving. For 24 and 32 GT/s operation, an additional optional four-way interleaving, configured as 45°/135°, is supported for power optimization. Fig. 5 summarizes two-way or four-way clock interleaving options for implementation flexibility and power optimization. At higher data rates, it is often more power efficient to implement four-way interleaving versus two-way. The proposed overall clocking scheme provides the best

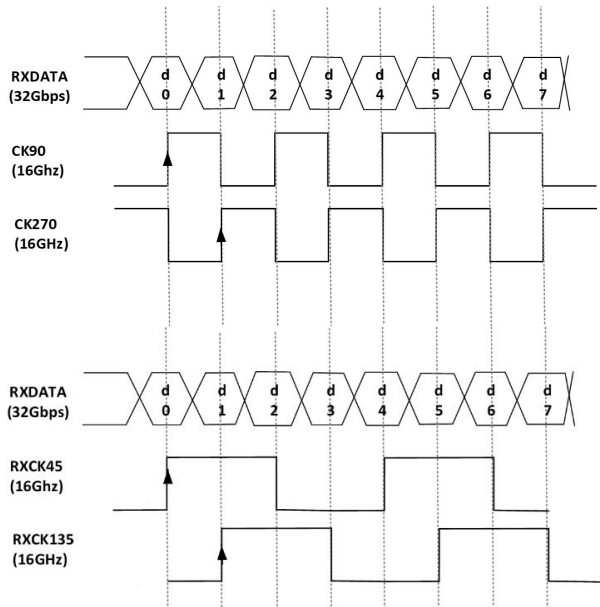


Fig. 5. Two-way and four-way clock interleaving.

TABLE II
RAW WIRE BER SPECIFICATION

Data (GT/s)	4	8	12	16	24	32
UCIe-A	1e-27	1e-27	1e-27	1e-15	1e-15	1e-15
UCIe-S	1e-27	1e-27	1e-15	1e-15	1e-15	1e-15

optimization point when taking entry/exit latency into account and the corresponding high di/dt 's and the higher power supply noise. This is especially important at lower data rates, which will also be very relevant for future potential 3-D die-to-die standards.

Some additional details of the PHY architecture include a valid lane to enable clock gating (<1 ns) when traffic is idle. We estimate $\geq 85\%$ of the full-activity power can be saved in this idle state by gating most of the PHY clocks just beyond the main trunk of the PLL output distribution to each PHY module. This is particularly helpful for power reduction method for workloads running at lower than 100% utilization. We have also allocated a Track Lane, which adjusts clock-to-data skew in the background due to temperature drift.

Source synchronous clocking together with maintaining a 0.5 UI clock-to-data skew separation allow for very robust link performance in power-supply noisy environments. This enables lower VCC operation for an optimal balance of best power/latency performance while avoiding super tight supply regulation to ease SoC integration. Table II summarizes the raw wire bit error rate (BER) needed to hit the failure in time (FIT) rate of $\ll 1.0$, as shown in Table I. At lower operating data rates, the PHY raw wire BER is $1e - 27$. At higher data rates, the raw wire BER is $1e - 15$; with a 16-bit CRC achieves the target FIT.

C. Standard Package Channel Design

We define the UCIe standard module based on the state-of-the-art organic flip-chip packaging capabilities to achieve




Fig. 6. Illustration of an organic flip-chip package.




Fig. 7. Bump map of one module for standard package. Signal bumps are in blue, ground bumps in green, and power bumps in red.

the performance targets in Table I. Our proposal incorporates flexibilities to envelope a wide range of technology offerings in the packaging industry. We propose a fixed shoreline length for the module to facilitate the interoperability between various chiplets.

The organic flip-chip package, illustrated in Fig. 6, is the mainstream packaging solution today [8]. The technology envelope has increased tremendously in the past 30 years. Currently, the maximum layer count is greater than 20 (e.g., two core layers with nine build-up layers on both front and back sides), and the largest form factor is more than 3000 mm^2 . To keep pace with Moore's Law scaling, the minimal pitch of controlled chip collapse connection (C4) bumps has reduced to about $100 \mu\text{m}$, and the minimal pitch for routing traces has shrunk to about $20 \mu\text{m}$. These lead to about 20 IO/mm escape density at the die edge with each routing layer. To maintain affordability, these pitches and densities are expected to scale slowly. As a result, higher IO bandwidth density needs to rely more on faster data rates and larger layer counts.

The basic UCIe-S block, either for TX or for RX, comprises 20 signals in unidirectional single-ended mode. A recommended bump map for them is illustrated in Fig. 7. The first 10 signals closer to the die edge escape the bump field in one routing layer, while the other ten signals in the back escape using the same trace design strategy in the next routing layer. The width of the block is chosen to be $571.5 \mu\text{m}$, so the pitch along the die edge, P_y , is $190.5 \mu\text{m}$. The other dimensions are flexible, based on the technology option selected. Table III lists two design cases based on 110- and $130\text{-}\mu\text{m}$ minimal bump pitches. The pitch in the diagonal direction P and the pitch in the depth direction P_x are adjusted accordingly. The other dimensions need to comply to the following two conditions:

$$P = D + L + 2S \quad (1)$$

$$P_y = D + 3L + 4S \quad (2)$$

TABLE III
UCIE STANDARD MODULE AT DIFFERENT BUMP PITCHES

P (μm)	P_x (μm)	P_y (μm)	Escape Density (IO/mm/layer)	Depth (μm)	Areal Density (IO/mm ²)
110	110	190.5	17.5	715	49
130	177	190.5	17.5	1151	30.4




Fig. 8. Four UCIE standard modules with stacking and checkerboard arrangement. The red arrows indicate the routing connections.

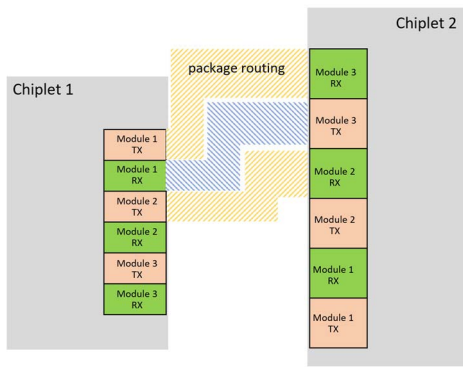


Fig. 9. Illustration of the routing channel for three standard package modules with mismatched PHY block between two chiplets.

where D is the diameter of via pad, L is the width of trace, and S is the spacing around the trace. With $571.5 \mu\text{m}$ block width, the escape density is 17.5 IO/mm for each routing layer, and the aggregate density is 35 IO/mm for two routing layers.

Our proposed UCIE-S module includes one TX block and one RX block. Therefore, the full module width is $1143 \mu\text{m}$. The escape trace ordering is symmetric between TX and RX, so that a single PHY design can be used to interconnect all chiplets. The standard module also supports stacking to further increase the escape density to 70 IO/mm with four routing layers. These modules are arranged in a checkerboard pattern (Fig. 8). The modules at the die edge are connected using two top routing routing layers, while the ones in the back are connected using two deeper routing layers. We recommend adhering to the same block width. If the block width is significantly different between the two chiplets as shown in Fig. 9, it requires room for fan-in and fan-out routing. This increases the channel length and requires significant die-to-die distance, thus is not feasible when the PHY blocks of the two chiplets are facing each other across a tiny chip gap.

The areal density correlates to the bump pitch. Coarser pitch causes larger bump area depth and lower areal density, as shown in Table III. Advances in packaging technology has been pushing the bump pitch smaller to increase the areal density. Reducing the ground bumps also increases areal density. The bump map in Fig. 7 has good ground isolations to ensure that the channel through the deep package layers can meet 32 GT/s requirements. However, the ground bumps can be reduced to save silicon area if the target data rate is lower or if there is no module stacking and the via-stack height is kept short. This gives flexibility to accommodate different trade-off points between bandwidth density, silicon area, and package layer count.

D. Advanced Package Channel Design

In the past decade, new advanced packaging architectures have emerged, which have achieved significant reduction of package feature sizes [8]. To better utilize the capabilities of these advanced technologies, we define a separate UCIE-A module to support the performance targets in Table I. Similar to the standard module, the advanced module offers flexibility to envelope a wide range of packaging technology offerings available in the industry. The proposed shoreline width of the module is critical for the interoperability between independently developed chiplets. We have built-in redundancy for repair, which is important to achieve good packaging yield.

Advanced packaging technologies in the industry have enabled the bump pitch to be less than $55 \mu\text{m}$ and improved the routing trace pitch to be only a couple of μm . Many of these technologies leverage the silicon manufacturing capabilities. The small via size and good via alignment enable the vias to be enveloped by the trace. This creates high flexibility for signals to switch layers and swizzle routing orders. These are dramatically different than the standard packaging solutions.

The bump map of the standard module in Fig. 7 is not suitable for advanced packaging technologies. It forces a 16-bit cluster design and will require a stacking of at least ten modules to fully utilize the routing density capability of the advanced packaging. The corresponding on-die data path in and out of these modules is very complex and prevents the modularized PHY design. It does not incorporate the redundant bits for repair, which is required by the advanced packaging. In addition, the checkerboard module arrangement in Fig. 8 causes part of the channel to be significantly longer, which is going to limit the bandwidth and power efficiency.

Therefore, the advanced module is defined with different size and form factor. An example bump map [4] based on $45\text{-}\mu\text{m}$ pitch is shown in Fig. 10. Similar to the standard module, it comprises one TX block and one RX block. The TX block is close to the die edge, while the RX block is at the back. Both comprise 74 signals, among which there are 64 data lanes and ten overhead signals. A special type of overhead signal is the redundancy signal for repair. Advanced packaging solutions usually involve tens of thousands of micro-bump connections at fine pitch. The advanced module allocates two redundant bumps for every 32 data signals to repair

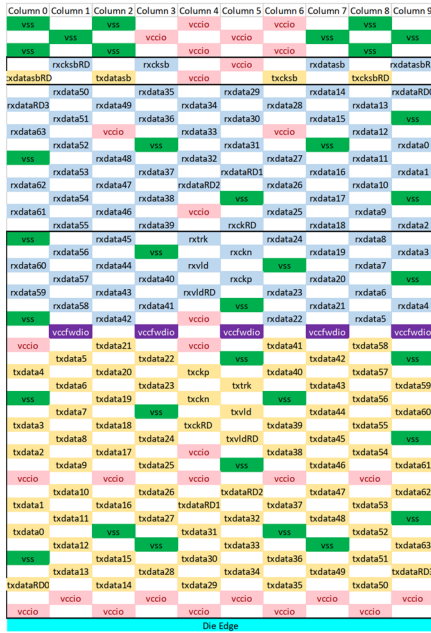


Fig. 10. Example bump map of one UCIE advanced module, with TX signals in yellow, RX signals in blue, ground in green, and power in red and purple.

the potential assembly failures. This is necessary for good manufacturing yield.

The module width is fixed at $388.8 \mu\text{m}$. When using advanced packaging, the two chiplets are usually placed right next to each other to reduce the channel length, which is a critical factor for power efficiency and transceiver design. However, if the module width is significantly different between the two chiplets, there is little room for the fan-in and fan-out to make the connections. This is similar as the problem of the standard package module in Fig. 9. Since the advanced package channel is very sensitive to the channel length due to the strong RC behavior, the mismatched module width can significantly degrade the channel bandwidth and power efficiency. Therefore, a fixed module width is foundational for chiplet interoperability.

With $388.8\text{-}\mu\text{m}$ module width, the $45\text{-}\mu\text{m}$ pitch bump map has ten columns, as shown in Fig. 10. The bump pitch along the die edge is $77.76 \mu\text{m}$, and the bump pitch is about $45 \mu\text{m}$ in both the depth direction and the diagonal direction. This follows the hexagonal pattern, which maximizes the bump density. For tighter bump pitches, the number of columns and rows can be adjusted to achieve the maximal bump density. For example, if the packaging technology supports $25\text{-}\mu\text{m}$ minimal bump pitch, the number of columns can increase to 18, and the pitch along the die edge reduces to $43.2 \mu\text{m}$, so that the module width remains at $388.8 \mu\text{m}$. The pitch along the depth and diagonal direction is about $25 \mu\text{m}$. This also follows a hexagonal pattern.

The escape density at the die edge is about 400 IO/mm for the advanced module. The areal density scales with the bump pitch. At $45\text{-}\mu\text{m}$ pitch, the bump field depth is about 1 mm, so the areal density is about 400 IO/mm². It can improve




Fig. 11. Four UCIE advanced modules in uniform arrangement. The red arrows indicate the routing connections.

quadratically as the bump pitch shrinks. Advanced packages have fine design rules of vias and traces, so that the bump escape is much less restrictive than the organic package. The TX and RX modules can be arranged uniformly along the die edge instead of the checkerboard pattern. As shown in Fig. 11, all the TX modules can be placed at the die edge, while all the RX modules are behind them. This has two advantages: first, it only needs a single flavor of the TX and RX block design, hence simplifies the circuit design. The second is that it matches the routing length in both directions. In other words, it reduces the worst case routing length. This significantly improves the bandwidth of these lossy channels.

The bump map in Fig. 10 of the advanced module is not suitable for the standard package with $110 \mu\text{m}$ bump pitch. The module becomes at least 2.5-mm deep, while the ground shielding is far from sufficient for the long vias in the standard package. It will need at least eight routing layers to break out all the signals.

IV. PACKAGE CHANNEL PERFORMANCE RESULTS

We simulated the reference channels for both the UCIE-S and UCIE-A modules to the validate the electrical performance.

A. Standard Package Channel Performance

The standard package channel is based on the stacked module configuration, shown in Fig. 8. Each module uses the bump map of Fig. 7 with $110\text{-}\mu\text{m}$ bump pitch. The package substrate is assumed to be 8–2–8, meaning eight build-up layers on both front and back side of two core layers. The routing connection of the stacked UCIE-S modules requires four routing layers, which are the 2nd, 4th, 6th, and 8th metal layer counting from the package surface. The worst case channel is in the 8th metal layer, as it has the longest vertical via stack height with the highest crosstalk.

The channel length depends on the placement of the two chiplets. Longer channel causes higher loss, hence worse signaling margin. The characteristics of a 25-mm long channel are plotted in Fig. 12. The loss and cumulative crosstalk are based on the voltage transfer function (VTF) [4, 9] instead of the S-parameters. It combines the termination and capacitive loading of TX and RX with the channel for a comprehensive evaluation. The VTF loss is -8.77 dB and




Fig. 12. VTF-based loss and crosstalk of a 25-mm-long reference UCle standard channel, assuming 30 Ω and 125 fF at TX and 50 Ω and 125 fF at RX.

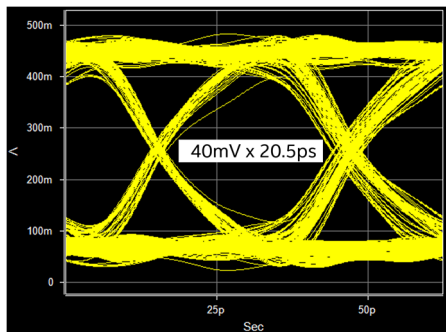


Fig. 13. RX eye diagram of the reference UCle standard channel at 32 GT/s with a 2-dB TX de-emphasis.

the cumulative VTF crosstalk is -31.3 dB at 16 GHz. They are based on the TX and RX requirements of the 32 GT/s standard package channel in the UCle specification [4]: 30-Ω TX termination, 50-Ω RX termination, and 125-fF equivalent capacitance for both TX and RX. The low die capacitance usually requires a low-voltage electrostatic discharge (ESD) protection, on-die inductor coils, and the TX and RX circuit loading optimizations. As the resistance termination and capacitive loading of TX and RX are incorporated into the VTF loss and crosstalk in Fig. 12, there are small reflections in the channel characteristics. These are fully captured in the time domain simulations. With 2-dB TX de-emphasis, the RX eye diagram at 32 GT/s is shown in Fig. 13. The worst case eye width opening at 40-mV eye height is more than 65% UI based on the peak distortion analysis. The TX de-emphasis has $\sim 10\%$ UI contribution. This enables about 224 GB/s/mm data bandwidth density across the die edge, excluding the overhead of clock and control signals.

The on-package channel can be divided into three segments: the bump break-out region of the first chiplet, the routing between the first and the second chiplet, and the bump break-in region of the second chiplet. The routing is typically a 50-Ω transmission line, which can be 2 mm–10 s of mm in length. The bump break-out and break-in segments are very short. The whole channel is relatively simple. Fig. 14 shows the relationship between the margin at 16 GT/s and the termination settings for the reference channel in Fig. 12. The optimal RX setting is about 50 Ω. This confirms that it is always preferred

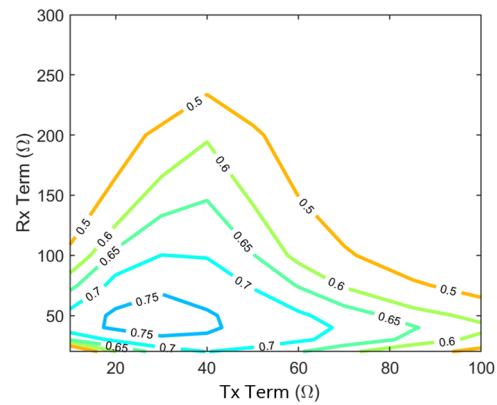


Fig. 14. TX and RX termination sensitivity on the signaling margin of a 25-mm-long UCle standard channel at 16 GT/s. The contours are the margin in UI.

to match the RX to the channel impedance to minimize the RX reflections. However, the TX termination can be lower than the channel impedance. The lower TX termination boosts the voltage level into the channel and increases the RX voltage swing and the signaling margins. However, larger mismatch at the TX side will cause unwanted reflections. As a result, the optimal TX setting is about 30 Ω. The sensitivity to resistance terminations is not significantly affected by the TX de-emphasis, the capacitive loading, or the data rate. The termination can be adjusted for lower data rate and shorter reach applications to trade off signaling margin for better energy efficiency.

B. Advanced Package Channel Performance

A reference advanced package channel is put together to validate the UCle-A module performance. The bump escape and the routing traces are the two key components of the physical channel that require optimization. The bump-via crosstalk is highly sensitive to the location of ground shielding. Therefore, the optimal placement of shielding bumps needs to strike the balance between the silicon area and the crosstalk level. The routing trace performance is greatly affected by the metal stack-up. This is a key area of the interconnect technology development for optimization between the channel reach, the routing density, and the bandwidth. The reference channel is based on the 45-μm pitch bump map in Fig. 10. The routing trace design is based on the 1 μm minimal width and spacing design rules. The signals of opposite directions are separated into two routing layers with a ground reference layer in between. The channel length is assumed to be 1.5 mm. The VTF loss and cumulative crosstalk of 20 signals are overlaid in Fig. 15. The worst case VTF loss is -2.73 dB at 8 GHz. The worst case cumulative VTF crosstalk is -24.3 dB at 8 GHz. The VTF metrics are based on the TX and RX requirements of the 16 GT/s advanced package channel in the UCle specification [4]: 250-fF capacitive loading at 25-Ω TX and 200-fF capacitive loading at unterminated RX. Since it is difficult to fit on-die inductors in the fine pitch bump field, the capacitive loading is higher for the advanced package case.

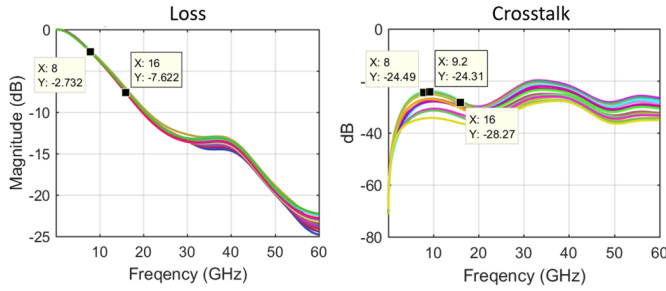


Fig. 15. VTF-based loss and crosstalk of a 1.5-mm long 16 GT/s reference UCIE advanced channel, assuming 250 fF loading at 25- Ω TX and 200 fF loading at unterminated RX.

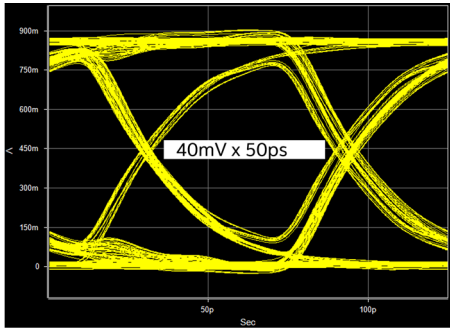


Fig. 16. RX eye diagram of the 1.5-mm-long reference UCIE advanced channel at 16 GT/s.

The corresponding RX eye diagram at 16 GT/s is shown in Fig. 16. The unterminated RX increases the voltage swing. The eye is widely open due to the low loss and low crosstalk up to the Nyquist frequency. Based on the peak distortion analysis, the worst case eye width opening at 40-mV eye height is more than 80% UI without using any equalization circuits. This enables about 658 GB/s/mm bandwidth density at 16 GT/s across the die edge, excluding the overhead signals. This is already three times of what the 32 GT/s standard module can achieve. At the same data rate, the advanced module delivers sixfold bandwidth density of the standard module. Advanced packaging technologies are rapidly evolving. Design features keep shrinking, and the layer count keeps increasing. These technology advancements will continue to reduce the channel loss and crosstalk to support higher data rates such as 32 GT/s.

Since advanced packaging channel is extremely short, it has a different sensitivity to the TX and RX termination than the standard channel. Fig. 17 shows the relationship between the margin at 16 GT/s and the TX and RX termination settings. It prefers a stronger TX and does not show significant sensitivity to RX termination. Hence, we set the UCIE-A TX termination to 25 Ω with an unterminated RX. This maximizes the signaling margin, simplifies RX design, and reduces power consumption.

V. CONCLUSION

The industry needs an open chiptlet ecosystem that will unleash innovations across the compute continuum.




Fig. 17. TX and RX termination sensitivity on the signaling margin of UCIE advanced channel at 16 GT/s. The contours are the margin in UI.

Our approach with UCIE 1.0 specification offers compelling power-efficient and cost-effective performance with plug-and-play and compliance aspects addressed upfront. We foresee the next generation of innovations will happen at the chiptlet level allowing an ensemble of chiptlets offering different capabilities for the customer to choose from that best addresses their application requirements.

In the future, we will have more sensitivity studies of the clocking architecture and the corresponding power noise impact on signaling margin. We expect innovations for even more power-efficient and cost-effective solutions as bump pitches continue to shrink and 3-D package integration becomes mainstream. Those may require wider links running at slower speeds and get closer to on-die connectivity from a latency, bandwidth, and power-efficiency point of view. Advances in packaging and semiconductor manufacturing technologies will revolutionize the compute landscape in the coming decades. UCIE is well poised to enable innovations in the ecosystem to take full advantage of these technological advances as they unfold.

REFERENCES

- [1] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, pp. 1–6, Apr. 1965.
- [2] N. Nassif *et al.*, "Sapphire rapids: The next-generation Intel Xeon scalable processor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2022, pp. 44–46, doi: 10.1109/ISSCC42614.2022.9731107.
- [3] D. D. Sharma, "Universal chiptlet interconnect express (UCIE)[®]: Building an open chiptlet ecosystem," UCIE Consortium, Beaverton, OR, USA, White Paper, Mar. 2022.
- [4] (Feb. 17, 2022). *Universal Chiptlet Interconnect Express (UCIE) Specification Rev 1.0*. [Online]. Available: <https://www.uciexpress.org>
- [5] *PCI Express[®] Base Specification Revision 5.0, Version 1.0*, PCI-SIG, Beaverton, OR, USA, May 28, 2019.
- [6] *PCI Express[®] Base Specification Revision 6.0, Version 1.0*, PCI-SIG, Beaverton, OR, USA, Jan. 11, 2022.
- [7] CXL Consortium. (Sep. 9, 2020). *Compute Express Link 2.0 Specification*. [Online]. Available: <https://www.computeexpresslink.org>
- [8] (2021). *IEEE Electronics Packaging Society Heterogeneous Integration Roadmap*. [Online]. Available: <https://eps.ieee.org/hir>
- [9] R. Mahajan *et al.*, "Embedded multi-die interconnect bridge (EMIB)—A localized, high density multi-chip packaging (MCP) interconnect," *IEEE Trans. Compon., Packag., Manuf. Technol.*, vol. 9, no. 10, pp. 1952–1962, Oct. 2019.



Debendra Das Sharma (Senior Member, IEEE) was born in Odisha, India, in 1967. He received the B.Tech. degree (Hons.) in computer science and engineering from IIT Kharagpur, Kharagpur, India, in 1989, and the Ph.D. degree in computer systems engineering from the University of Massachusetts, Amherst, MA, USA, in 1995.

He joined Hewlett-Packard, Roseville, CA, USA, in 1994, and Intel, Santa Clara, CA, USA, in 2001. He is currently a Senior Fellow with Intel. He is responsible for delivering Intel-wide critical inter-

connect technologies in Peripheral Component Interconnect Express (PCI Express), Compute Express Link (CXL), Universal Chiplet Interconnect Express (UCIe), Coherency Interconnect, Multi-Chip Package Interconnect, and Rack Scale Architecture. He has been leading the development of PCI-Express, CXL, and UCIe inside Intel as well as across the industry since their inception. He holds 160+ U.S. patents and more than 400 patents worldwide.

Dr. Das Sharma has been awarded the Distinguished Alumnus Award by IIT, in 2019, the 2021 IEEE Region 6 Engineer of the Year Award, the PCI-SIG Lifetime Contribution Award in 2022, and the 2022 IEEE CAS Industrial Pioneer Award. He is currently the Chair of UCIe Board, a Director of PCI-SIG Board, and the Chair of the CXL Board Technical Task Force.



Gerald Pasdast received the B.S. degree in electrical engineering from San Jose State University, San Jose, CA, USA, in 1996.

From 1996 to 1997, he was a Test and Product Engineer with Cypress Semiconductor, San Jose, CA, working on FIFO and dual port RAMs. Since 1997, he has been with Intel Corporation, Santa Clara, CA, USA, worked on various IO technologies. He is currently a Senior Principal Engineer with the Central IP Team and his area of focus has been on die-to-die (D2D) IO architecture and

technology since 2010 on several proprietary PHY running on both standard package (2-D) traces as well as advanced packages including EMIB (2.5-D) and Foveros (3-D chip stacking) that have been productized on server/high performance computing (HPC), graphics, and client CPUs. Most recently, he has coauthored the Universal Chiplet Interconnect express (UCIe) spec1.0. He has over 30 granted patents with an additional 20 pending mostly in the area of D2D architecture and packaging.

Mr. Pasdast is the Interim Chair of the UCIe Consortium Form Factor and Compliance WG.



Zhiguo Qian (Senior Member, IEEE) received the bachelor's and master's degrees in electrical engineering from Southeast University, Nanjing, China, in 2001 and 2004, respectively, and the Ph.D. degree in electrical and computer engineering from the University of Illinois, Urbana-Champaign, Urbana, IL, USA, in 2009.

From 2009 to 2010, he was a Principal Engineer with ANSYS, San Jose, CA, USA. In 2010, he joined Intel Corporation, Chandler, AZ, USA, where he is currently a Principal Engineer. He has

authored over 70 peer-reviewed technical articles and holds 32 U.S. patents. His current research interests include signal and power integrity of 2.5-D and 3-D semiconductor packaging, high speed I/O channel, and computational electromagnetics.

Dr. Qian was a two-time recipient of the Intel Achievement Award. He received the IEEE Phoenix Section Young Engineer of the Year Award in 2015. He has also served as a Semiconductor Research Corporation (SRC) Industrial Liaison for academic research projects.



Kemal Aygün (Fellow, IEEE) received the Ph.D. degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign, Urbana, IL, USA, in 2002.

In 2003, he joined the Intel Corporation, Chandler, AZ, USA, where he is currently an Intel Fellow and manages the High Speed I/O (HSIO) Team, Electrical Core Competency Group. He has coauthored five book chapters, more than 80 journal and conference publications, and holds 78 U.S. patents. His current research interests include novel technologies along

with electrical modeling and characterization techniques for microelectronic packaging.

Dr. Aygün was a recipient of the Semiconductor Research Corporation (SRC) Global Research Collaboration (GRC) Mahboob Khan Outstanding Mentor Award in 2008 and 2015 for his contributions in mentoring SRC GRC academic research projects. He was the General Chair of the 2020 IEEE Electrical Performance of Electronic Packaging and Systems Conference. He has been acting as a Distinguished Lecturer of the IEEE Electronics Packaging Society (EPS) and is a Co-Chair of the EPS Technical Committee on Electrical Design, Modeling, and Simulation.