



Universal Chiplet Interconnect Express (UCle)[®]: Building an open chiplet ecosystem

Dr. Debendra Das Sharma
Intel Senior Fellow and Chief Architect, I/O Technologies and Standards
Promoter Member of UCle

Universal Chiplet Interconnect Express (UCle)[®] is an open industry standard interconnect offering high-bandwidth, low-latency, power-efficient, and cost-effective on-package connectivity between chiplets. It addresses the projected growing demands of compute, memory, storage, and connectivity across the entire compute continuum spanning cloud, edge, enterprise, 5G, automotive, high-performance computing, and hand-held segments. UCle provides the ability to package dies from different sources, including different fabs, different designs, and different packaging technologies.

Motivation for on-package integration of Chiplets

In his seminal paper "[Cramming more components onto integrated circuits](#)" (published in Electronics, Volume 38, Number 8, on April 19, 1965), Gordon Moore posited that the number of transistors in an integrated circuit will double every two years. Popularly known as "Moore's law", it has held up for more than 50 years thus far. In the same paper, Gordon Moore also predicted the "Day of Reckoning": "*It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected.*" Today we see on-package integration of multiple dies in mainstream commercial offerings such as client CPUs, server CPUs, GP-GPUs, etc.

There are many drivers for on-package chiplets. As the die size increases to meet the growing performance demands, designs are running up against the die reticle limit. Examples include multi-core CPUs with core count in the hundreds or very large fanout switches. Even when a die can fit within the reticle limit, multiple smaller dies connected in a package may be preferable for yield optimization as well as die reuse across multiple market segments. On-package connectivity of identical dies enables these scale-up applications.

Another motivation for on-package integration is to lower the overall portfolio cost both from a product as well as project point of view and derive a time to market advantage. For example, the compute cores shown in Figure 1 can be implemented in an advanced process node to deliver leadership power-efficient performance at higher cost whereas the memory and I/O controller functionality may be reused from a design already deployed in an established (n-1 or n-2) process node. Such partitioning also results in smaller dies which results in better yield. IP porting costs across process nodes are high and increasing very rapidly for the advanced process nodes, as shown in Figure 2. Since we don't have to port all the IPs of dies whose functionality does not change, we save on costs in addition to getting the time to market advantage. Chiplet integration on package also enables a customer to make different trade-offs for different market segments by choosing different numbers and

types of dies. For example, one can choose different numbers of compute, memory, and I/O dies depending on the need of the segment. One does not need to do a different die design for different segments, resulting in lower product SKU cost.

On-package integration of chiplets enables a fast and cost-effective way to provide bespoke solutions. For example, different usages may need different acceleration capability but with the same cores, memory, and I/O, as shown in Figure 1. It also allows the co-packaging of dies where the optimal process node choice is made based on the functionality. For example, memory, logic, analog, and co-packaged optics each needs a different process technology which can be packaged with chiplets. Since package traces are short and offers dense routing, applications requiring high bandwidth such as memory access (e.g., High Bandwidth Memory), are implemented as on-package integration.

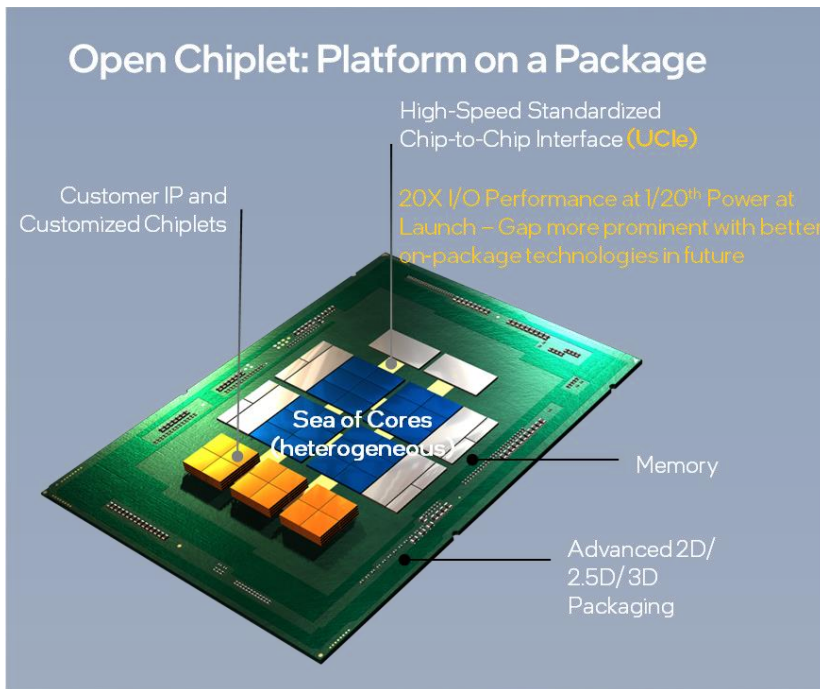


Figure 1: UCIe to enable an Open Chiplet Ecosystem delivering Platform on a Package

UCIe is a strategic on-package interconnect that comprehends these usage models in a forward-looking manner and is poised to transform the industry.

Factors influencing wide industry adoption of a standard

The recipe for developing a successful ecosystem is shown in Figure 3. UCIe has been formed with the institutional learning of decades in creating and successfully driving thriving open ecosystems such as Peripheral Component Interconnect® with PCI Express®, Universal Serial Bus®, Computer Express Link (CXL)®, etc.




Figure 2: Design cost across different process nodes (Source: IBS, as cited in IEEE Heterogeneous Integration Roadmap)

An open industry standards body defining a specification with compelling key performance indicators (KPIs) catering to a wide range of usages, including a comprehensive compliance and interoperability mechanism is essential to develop a healthy ecosystem. The UCle Specification Rev 1.0 is complete with the industry leading KPIs, debug support, and compliance considerations. On-package integration of dies has matured as a technology across the industry covering manufacturing, assembly, and test companies. We see compelling products being offered from multiple foundries as well as through the Outsourced Semiconductor Assembly and Test (OSAT) companies in the market place using proprietary interconnects. UCle is the result of industry leaders working together to develop a common standard so that multiple chiplets from different sources can interoperate seamlessly. While the UCle promoters cover a wide intersection of cloud, semiconductor manufacturing, OSAT, IP suppliers, and chip designers, the UCle consortium is open to all. UCle is poised to be the ubiquitous on-package interconnect for chiplets driving a thriving open chiplet ecosystem.

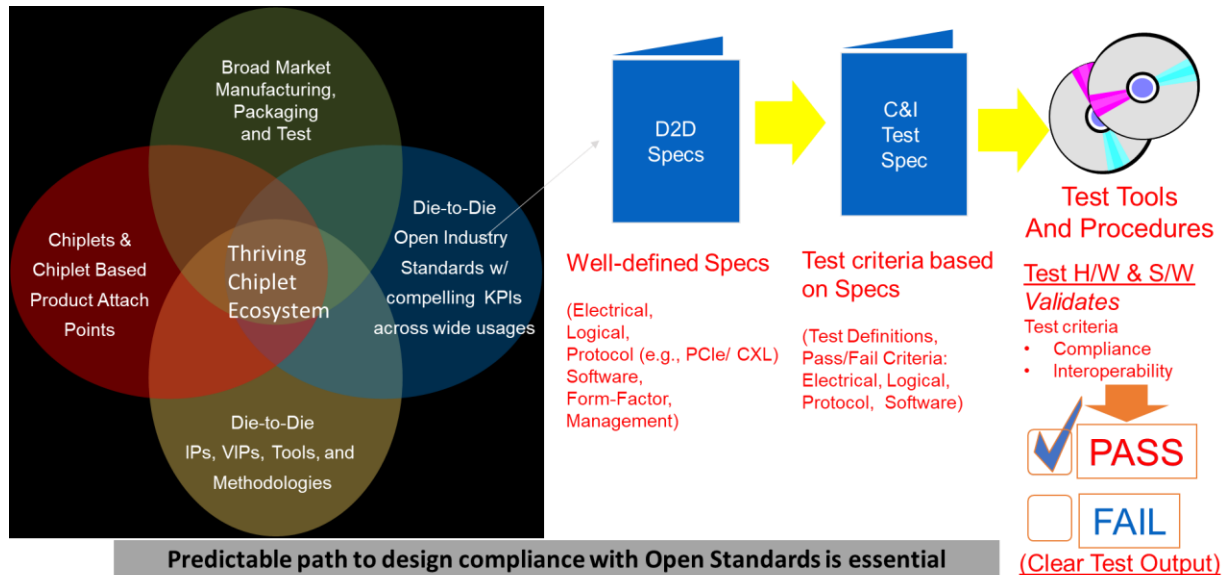


Figure 3: Ingredients of a successful and broad interoperable chiplet ecosystem

Usage Models and KPIs driven by UCle 1.0 Specification

UCle is a layered protocol, as shown in Figure 4a. The physical layer is responsible for the electrical signaling, clocking, link training, sideband, etc. The Die-to-Die adapter provides the link state management and parameter negotiation for the chiplets. It optionally guarantees reliable delivery of data through its cyclic redundancy check (CRC) and link level retry mechanism. When multiple protocols are supported, it defines the underlying arbitration mechanism. A 256-byte FLIT (flow control unit) defines the underlying transfer mechanism when the adapter is responsible for reliable transfer.

UCle maps PCIe and CXL protocols 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. With PCIe and CXL, SoC construction, link management, and security solutions that are already deployed can be leveraged to UCle. The usage models addressed are also comprehensive: data transfer using direct memory access, software discovery, error handling, etc., are addressed with 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. UCle also defines a “streaming protocol” which can be used to map any other protocol. Further, the UCle consortium can innovate on protocols in the future optimized for chiplets as usage models evolve in the future.

UCle 1.0 defines two types of packaging, as shown in Figure 4b. The standard package (2D) is used for cost-effective performance. The advanced packaging is used for power-efficient performance. There are multiple commercially available options, some of which are shown in the diagram. UCle specification embraces all types of packaging choices in these categories.

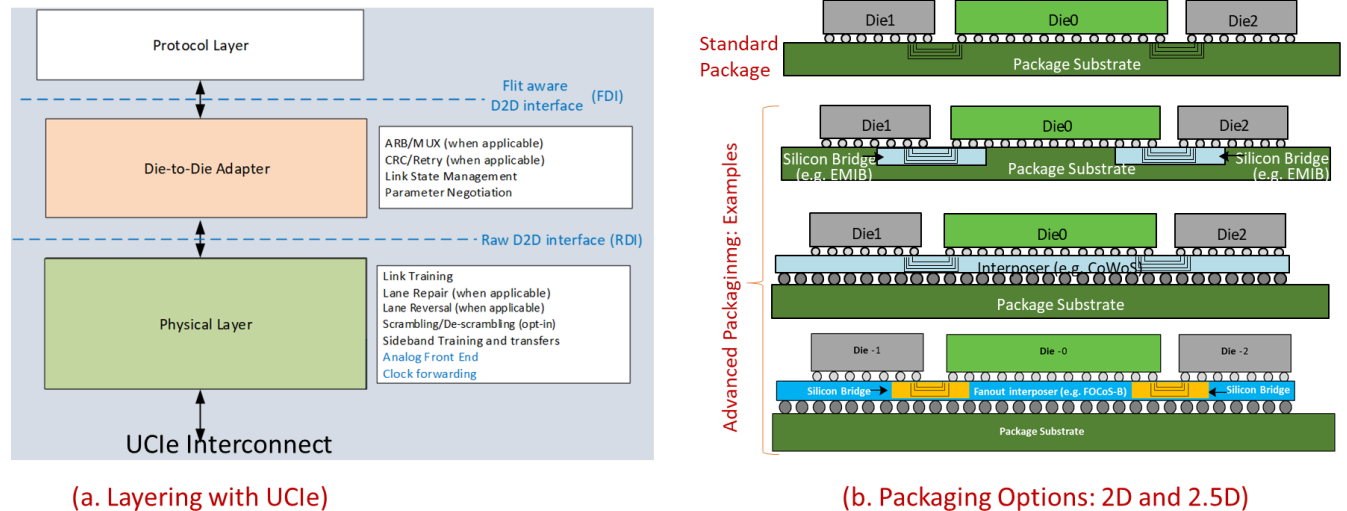


Figure 4: UCle : Layering Approach and different packaging choices

UCle supports two broad usage models. The first is package level integration to deliver power-efficient and cost-effective performance, as shown in Figure 5a. Components attached at the board level such as memory, accelerators, networking devices, modem, etc. can be integrated at the package level with applicability from hand-held to high-end servers with dies from multiple sources connected through different packaging options even on the same package. The second usage is to provide off-package connectivity using different type of media (e.g., optical, electrical cable, mmWave) using UCle

Retimers to transport the underlying protocols (e.g., PCIe, CXL) at the rack or even the pod level for enabling resource pooling, resource sharing, and even message passing using load-store semantics beyond the node level to the rack/ pod level to derive better power-efficient and cost-effective performance at the edge and data centers.

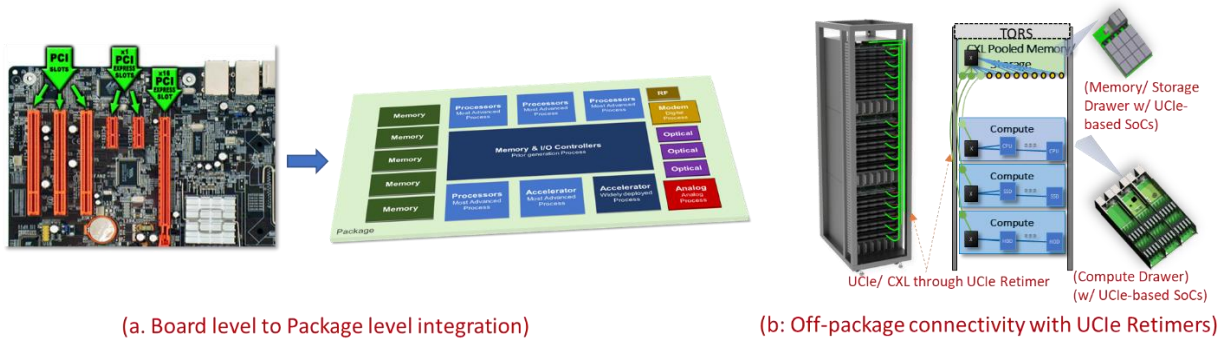


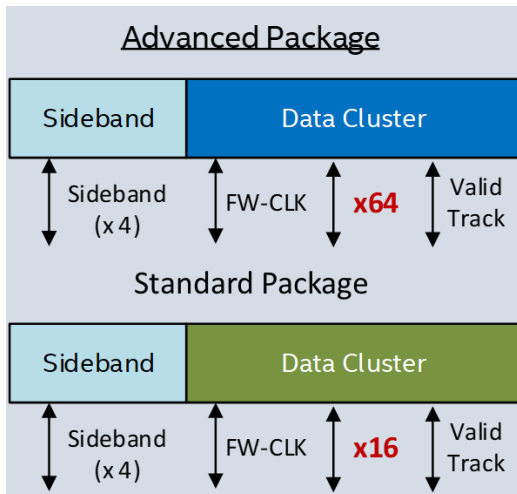
Figure 5: Usage Models supported by UCle: on-package integration as well as off-package connectivity with different media (e.g., optics, mmWave, electrical cable)

UCle supports different data rates, widths, bump-pitches, and channel reach to ensure the widest interoperability feasible, as detailed in Table 1. It defines a sideband interface for ease of design and validation. The unit of construction of the interconnect is a cluster which comprises of N single-ended, unidirectional, full-duplex Data Lanes ($N = 16$ for standard package and 64 for advanced package), one single-ended Lane for Valid, one lane for tracking, a differential forwarded clock per direction, and 2 lanes per direction for sideband (single-ended, one 800 MHz clock and one data). The advanced package supports spare lanes to handle faulty lanes (including clock, valid, sideband, etc) where as the standard package supports width degradation to handle failures. Multiple clusters can be aggregated to deliver more performance per Link, as shown in Figure 6.

Table 1 summarizes the key metrics for both the packaging options. A die with the standard package design is expected to interop with any other design on the standard package. Similarly, a die with the advanced package design will interoperate with any other die designed for the advanced package, even within the wide range of bump pitch from 25u to 55u. It should be noted that the KPI table conservatively estimates performance for the most widely deployed bump pitch today. For example, 45u is used for advanced packaging. The bandwidth density will go up by up to 3.24X if we go with a denser bump pitch of 25u. Even at 45u, the bandwidth density of 1300+ (both for linear as well as area) is about 20X what we can achieve with the most efficient PCIe SERDES. Similarly, PCIe PHY have a power efficiency of $\sim 10\text{pJ/b}$ today which can be lowered by up to 20X with the UCle based designs due to their shorter channel reach. UCle also enables for a linear power-bandwidth consumption curve with very fast entry and exit times (sub-ns vs multiple micro-seconds for SERDES based designs) while saving 90+% power. Thus, in addition to being really low power, it also is very effective in power savings, offering compelling power-efficient ultra-high performance. What is important is as the technology advances, these savings would be even more significant. UCle 1.0 has been defined to meet the projected needs of a wide range of challenging applications through almost the end of this decade.

Table 1: UCle 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



Die - 1		Die - 2			
x16	<-->	x16	CL-0 x16	<-->	CL-0 x16
x32	<-->	x32	CL-0 x16	<-->	CL-0 x16
			CL-1 x16	<-->	CL-1 x16
x64	<-->	x64	CL-0 x16	<-->	CL-0 x16
			CL-1 x16	<-->	CL-1 x16
			CL-2 x16	<-->	CL-2 x16
			CL-3 x16	<-->	CL-3 x16

(1, 2, or 4 Clusters can be combined in one UCle Link)

Figure 6: Cluster Width; 1, 2, or 4 Clusters can be combined in each packaging option to deliver higher bandwidth

Conclusions

There is a huge demand for an open chiplet ecosystem that will unleash innovations across the compute continuum. UCle 1.0 offers compelling power-efficient and cost-effective performance. The fact that it is an open standard with a plug-and-play model, modeled after several successful standards, and launched by the right set of industry leaders will ensure its wide-spread adoption. We foresee the next generation of innovations will happen at the chiplet level allowing an ensemble of chiplets offering different capabilities for the customer to choose from that best addresses their application requirements.

In the future, we expect the consortium to drive even more power-efficient and cost-effective solutions as bump pitches continue to shrink and 3D integration becomes mainstream. Those may require wider links running slower 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. UCle is well poised to enable innovations in the ecosystem to take full advantage of these technological advances as they unfold.