



# Chiplet互联架构分析 及其关键技术

奇异摩尔 祝俊东




## Agenda

Kiwimoore

# Chiplet 互联架构分析

Kiwi SoChiplet ， 高性能互联平台

## 持续提升 芯片性能



## 持续提升 系统性能



## 持续提升芯片性能

- 增加晶体管密度 ▶ **摩尔定律**
- 增加芯片面积 ▶ **Chiplet**
- 提升芯片内传输效率 ▶ **3DIC**
- 提升计算效率 ▶ **异构计算**
- 突破存储墙 ▶ **存算一体**

## 持续提升系统性能

- 增加计算规模 ▶ **超大规模集群**
- 提升 Cluster 内传输效率 ▶ **Chip2Chip Direct**
- 提升 Cluster 间传输效率 ▶ **无损数据传输RDMA**
- 降低传输负载 ▶ **Workload & Storage Acceleration**

## 持续提升芯片性能

增加晶体管密度

▶ 摩尔定律

增加芯片面积

▶ 新的瓶颈：互联

提升芯片内传输效率

▶ 新的途径：Chiplet, 网络加速

提升计算效率

▶ 异构计算

突破存储墙

▶ 存算一体

## 持续提升系统性能

增加计算规模

▶ 超大规模集群

提升 Cluster 内传输效率

▶ Chip2Chip Direct

降低传输负载

▶ 无损数据传输RDMA

▶ Workload & Storage Acceleration

# 高性能计算迈进Chiplet时代，三大主流形态

## Multi-Die

## Central IO Die

## Side Interface Die


intel.



Sapphire Rapids



LakeField




Ponte Vecchio




Falcon Shores  
Flexible Ratios and Configurations of Tiles




Meteor Lake




Zen



MI200




Zen2



Zen3+




MI300




RX7000




Grace Hopper



M1 Ultra




Dojo



BR100



昇腾910




AWS Graviton3






Multi-Die




Intel Sapphire Rapids

Central IO Die




AMD Bergama


Side Interface Die




AWS Graviton 4



Intel Emerald Rapids




AMD Genoa




Ampere Syrin

- Intel 7nm with 2.5D Advanced Package
- Core Counts 56
- Thread Counts 112
- Total LLC 105 MB
- PCIe 5.0 /CXL1.1 Lanes Up to 128
- Memory Bandwidth 307GB/s
- SPECrate@2017\_int\_base(2P) 990






- Intel 7nm with 2.5D Advanced Package
  - Core Counts 56
  - Thread Counts 112
  - Total LLC 105 MB
  - PCIe 5.0 /CXL1.1 Lanes Up to 128
  - Memory Bandwidth 307GB/s
  - SPECrate@2017\_int\_base(2P) 990
- 
- Intel Golden Cove Core
  - LLC 1.875MB
  - Mesh NoC Architecture
  - 15 Cores Per Tile, 14 be used



- Intel 7nm with 2.5D Advanced Package
- Core Counts 56
- Thread Counts 112
- Total LLC 105 MB
- PCIe 5.0 / CXL1.1 Lanes Up to 128
- Memory Bandwidth 307GB/s
- SPECrate@2017\_int\_base(2P) 990

- Intel Golden Cove Core
- LLC 1.875MB
- Mesh NoC Architecture
- 15 Cores Per Tile, 14 be used

- 400Gb/s Symmetric Crypto
- 160Gb/s Compression +
- 160Gb/s De-compression
- 400M Load Balancing Decisions per Second







- Intel 7nm with 2.5D Advanced Package
- Core Counts 56
- Thread Counts 112
- Total LLC 105 MB
- PCIe 5.0 / CXL1.1 Lanes Up to 128
- Memory Bandwidth 307GB/s
- SPECrate@2017\_int\_base(2P) 990

- Intel Golden Cove Core
- LLC 1.875MB
- Mesh NoC Architecture
- 15 Cores Per Tile, 14 be used

- 400Gb/s Symmetric Crypto
- 160Gb/s Compression +
- 160Gb/s De-compression
- 400M Load Balancing Decisions per Second





- HS D2D with 2.5D EMIIB & Si-Interposer
- D2D Bandwidth: 500GB/s
- Data Rate: 5GT/s
- Bump Pitch: 55um
- Energy Efficiency: 0.5pj/bit
- PHY Latency end-to-end TX+RX: 2.4ns

- Intel 7nm with 2.5D Advanced Package
- Core Counts 56
- Thread Counts 112
- Total LLC 105 MB
- PCIe 5.0 / CXL1.1 Lanes Up to 128
- Memory Bandwidth 307GB/s
- SPECrate@2017\_int\_base(2P) 990

- Intel Golden Cove Core
- LLC 1.875MB
- Mesh NoC Architecture
- 15 Cores Per Tile, 14 be used

- 400Gb/s Symmetric Crypto
- 160Gb/s Compression +
- 160Gb/s De-compression
- 400M Load Balancing Decisions per Second

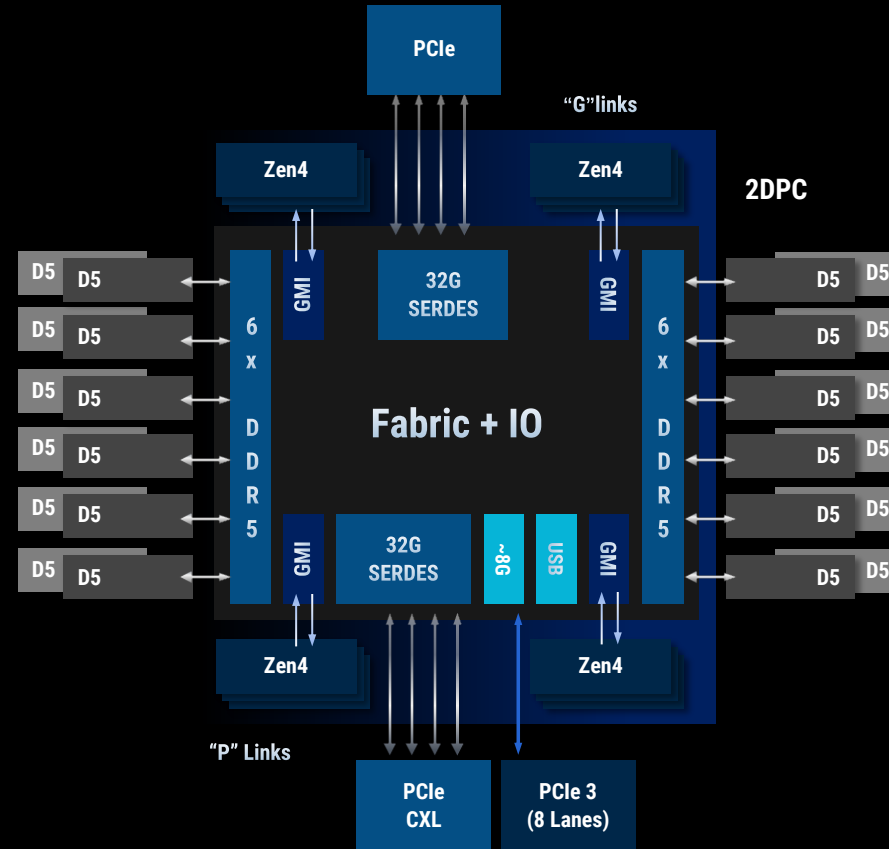


- HS D2D with 2.5D EMIIB & Si-Interposer
- D2D Bandwidth: 500GB/s
- Data Rate: 5GT/s
- Bump Pitch: 55um
- Energy Efficiency: 0.5pj/bit
- PHY Latency end-to-end TX+RX: 2.4ns

- 56 High Speed IO, 32 for PCIe5 (CXL 1.1), 24 for UPI
- Lanes Rate: 2.5-32Gbps
- Energy Efficiency: 6.48pj/bit




- Core Counts 96
- Threads Counts 192
- Total L3 Cache 384MB
- Memory Bandwidth 460.8GB/s
- IFOP(GMI3) Counts 12
- PCIe 5.0 Lanes 128
- DDR5 Memory Channels 12
- SPECrate@2017\_int\_base(2P) 1950






- Core Counts 96
- Threads Counts 192
- Total L3 Cache 384MB
- Memory Bandwidth 460.8GB/s
- IFOP(GMI3) Counts 12
- PCIe 5.0 Lanes 128
- DDR5 Memory Channels 12
- SPECrate@2017\_int\_base(2P) 1950




- 1CCD Core Counts 8
- TSMC N5 55 mm<sup>2</sup> 6.5B Trans
- 32 MB Shared L3 Cache


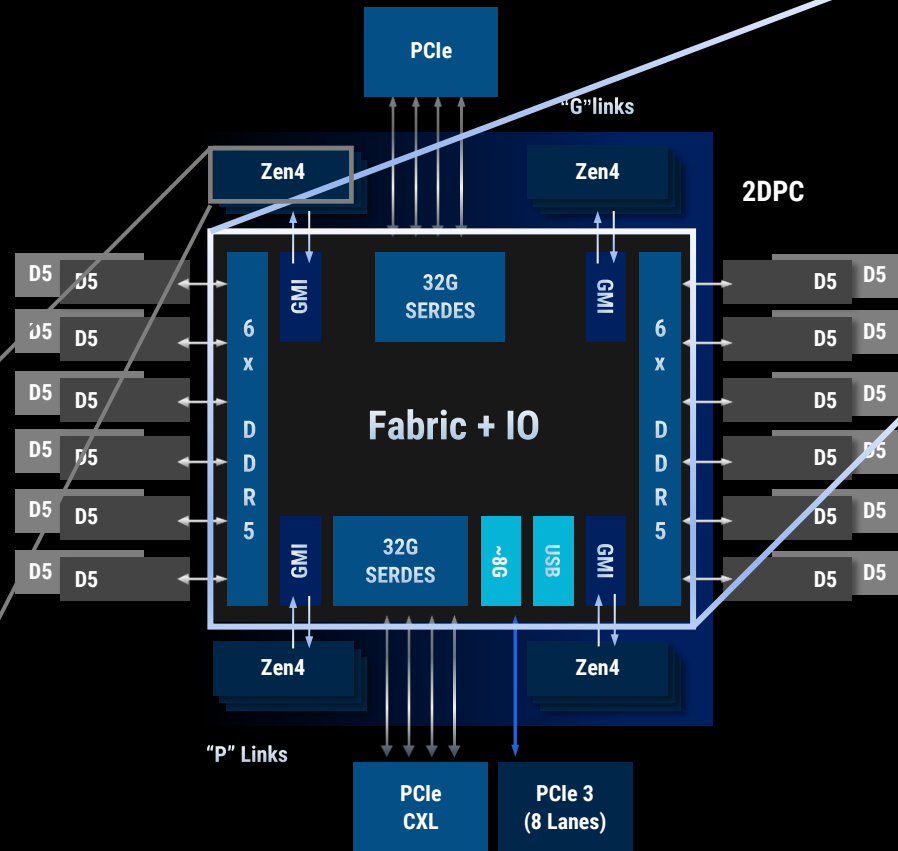


TSMC6 nm 386.88mm<sup>2</sup> 11B Trans


- Core Counts 96
- Threads Counts 192
- Total L3 Cache 384MB
- Memory Bandwidth 460.8GB/s
- IFOP(GMI3) Counts 12
- PCIe 5.0 Lanes 128
- DDR5 Memory Channels 12
- SPECrate@2017\_int\_base(2P) 1950



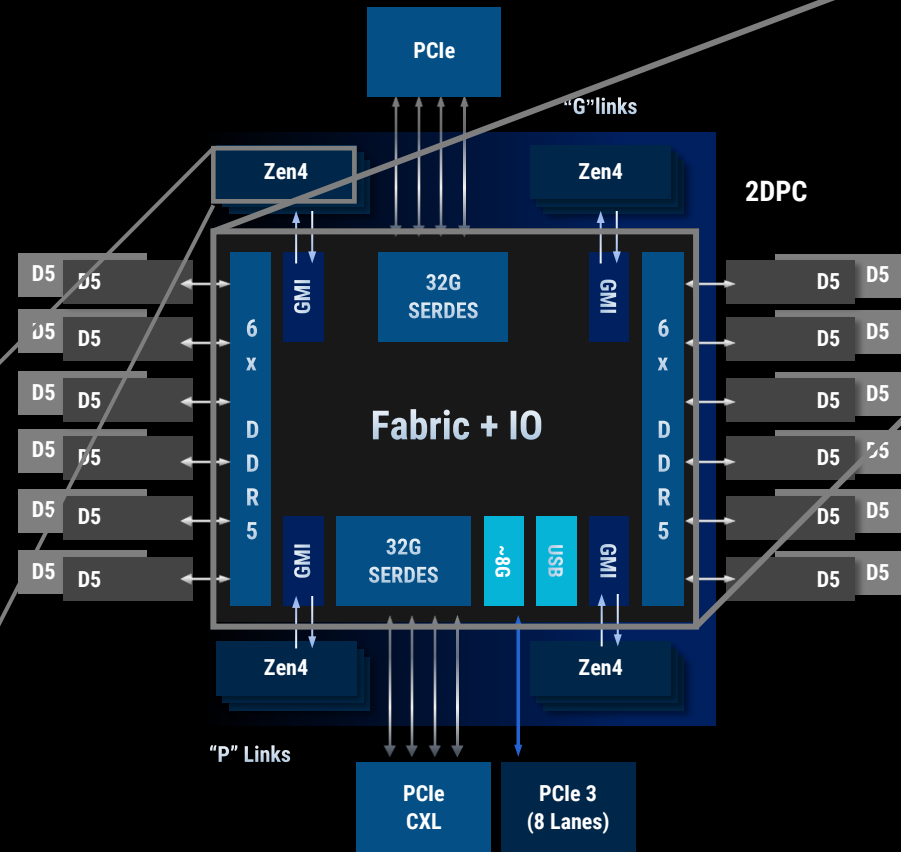
- 1CCD Core Counts 8
- TSMC N5 55 mm<sup>2</sup> 6.5B Trans
- 32 MB Shared L3 Cache



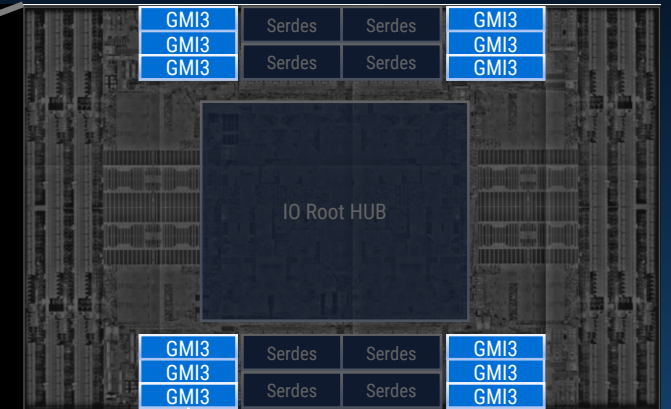
- Core Counts 96
- Threads Counts 192
- Total L3 Cache 384MB
- Memory Bandwidth 460.8GB/s
- IFOP(GMI3) Counts 12
- PCIe 5.0 Lanes 128
- DDR5 Memory Channels 12
- SPECrate@2017\_int\_base(2P) 1950



- 1CCD Core Counts 8
- TSMC N5 55 mm<sup>2</sup> 6.5B Trans
- 32 MB Shared L3 Cache



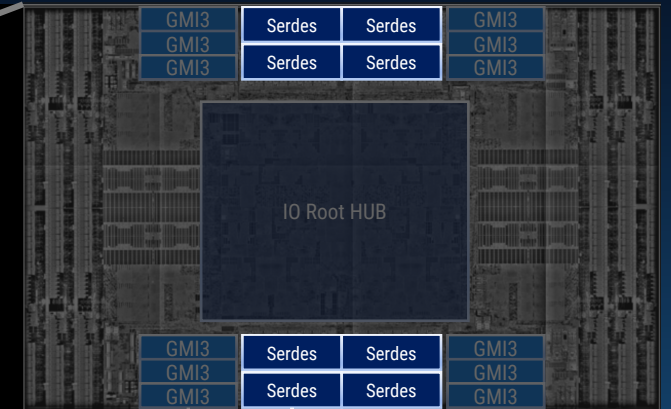


TSMC6 nm 386.88mm<sup>2</sup> 11B Trans



- 1CCD Core Counts 8
- TSMC N5 mm<sup>2</sup> 6.5B Trans
- 32 MB Shared L3 Cache

TSMC6 nm 386.88mm<sup>2</sup> 11B Trans

- Core Counts 96
- Threads Counts 192
- Total L3 Cache 384MB
- Memory Bandwidth 460.8GB/s
- IFOP(GMI3) Counts 12
- PCIe 5.0 Lanes 128
- DDR5 Memory Channels 12
- SPECrate@2017\_int\_base(2P) 1950



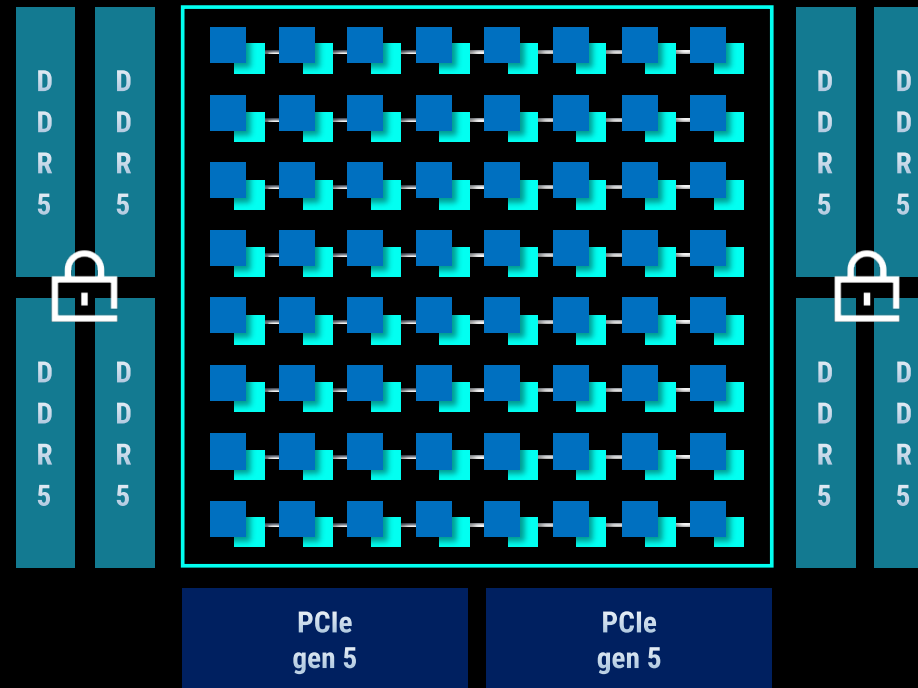
- 1CCD Core Counts 8
- TSMC N5 55 mm<sup>2</sup> 6.5B Trans
- 32 MB Shared L3 Cache



- 1CCD Core Counts 8
- TSMC N5 55 mm<sup>2</sup> 6.5B Trans
- 32 MB Shared L3 Cache

- 16 Lanes, 32 Gbps
- Combination xGMI(C2C Link), CXL, PCIe, SATA Serdes


- 1 Computing Die(N5) + 2 PCIe5 Controller Die + 4 DD5 Controller Die
- Core Counts 64
- Threads Counts 64
- Memory Bandwidth: 300GB/s
- Total L3 Cache 32MB
- SPECrate@2017\_int\_base(3P) ~630






- 1 Computing Die(N5) + 2 PCIe5 Controller Die + 4 DD5 Controller Die
- Core Counts 64
- Threads Counts 64
- Memory Bandwidth: 300GB/s
- Total L3 Cache 32MB
- SPECrate@2017\_int\_base(3P) ~630

- ARM V1 Core
- L2 Cache 1MB
- L1 Instruction Cache 64KB
- L1 Data Cache 64 KB
- 282 mm<sup>2</sup>
- CMN 650 Fabric
- 8\*8 Mesh, 32 Node for Memory, 32 Node for CPU

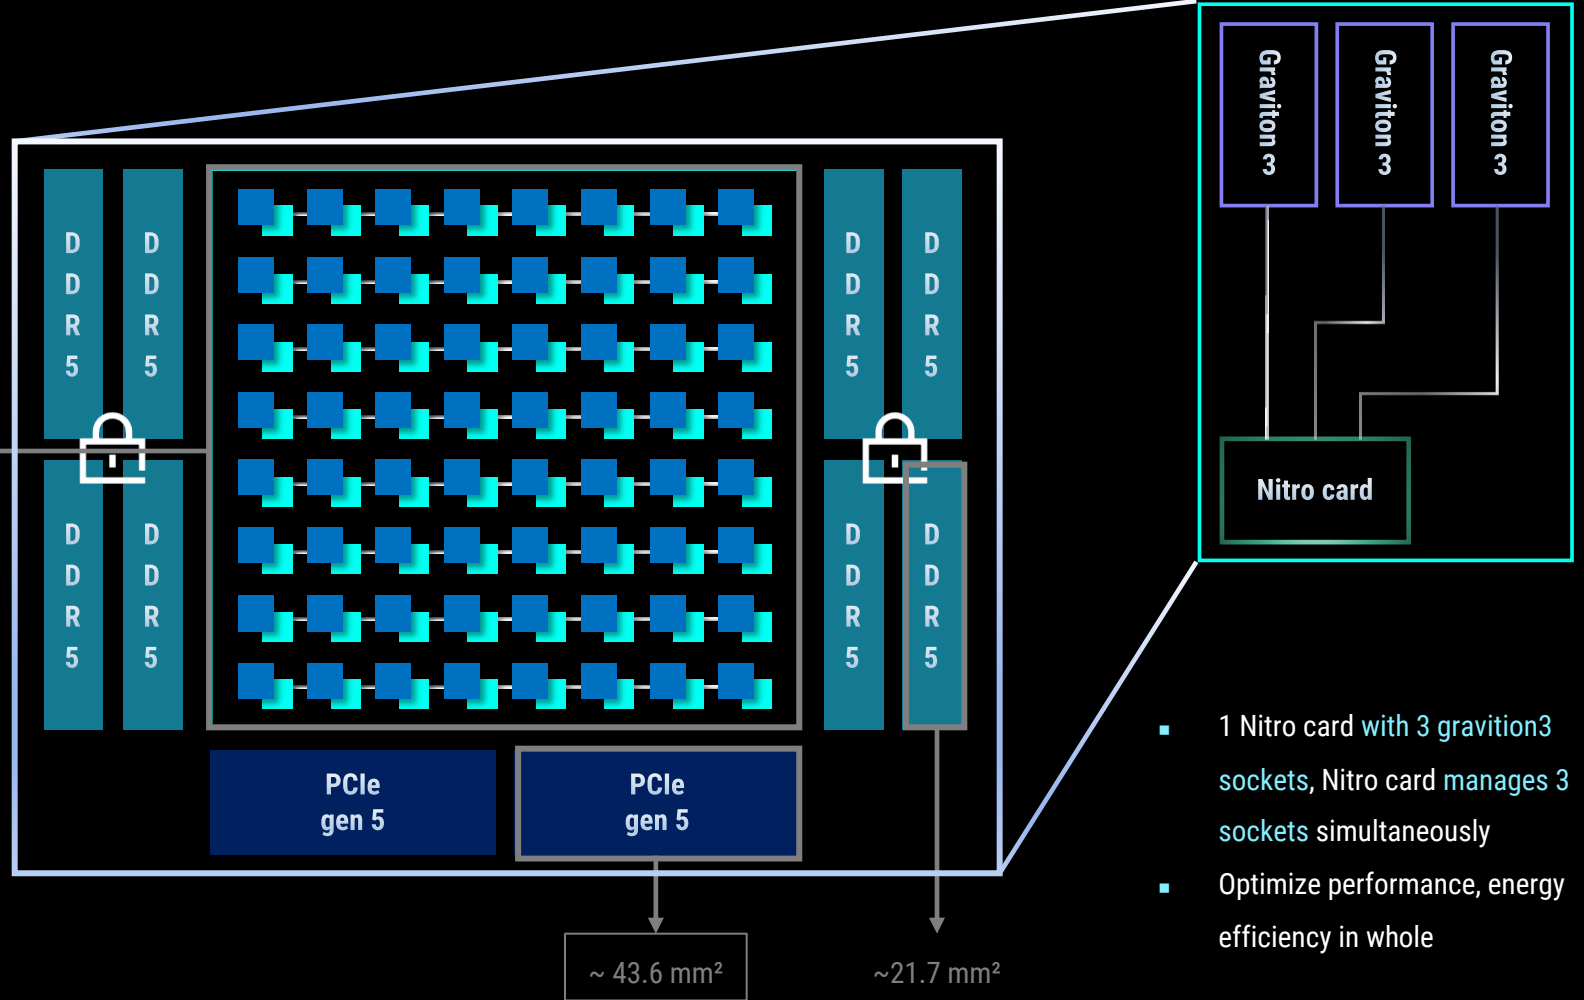


- 1 Computing Die(N5) + 2 PCIe5 Controller Die + 4 DD5 Controller Die
- Core Counts 64
- Threads Counts 64
- Memory Bandwidth: 300GB/s
- Total L3 Cache 32MB
- SPECrate@2017\_int\_base(3P) ~630

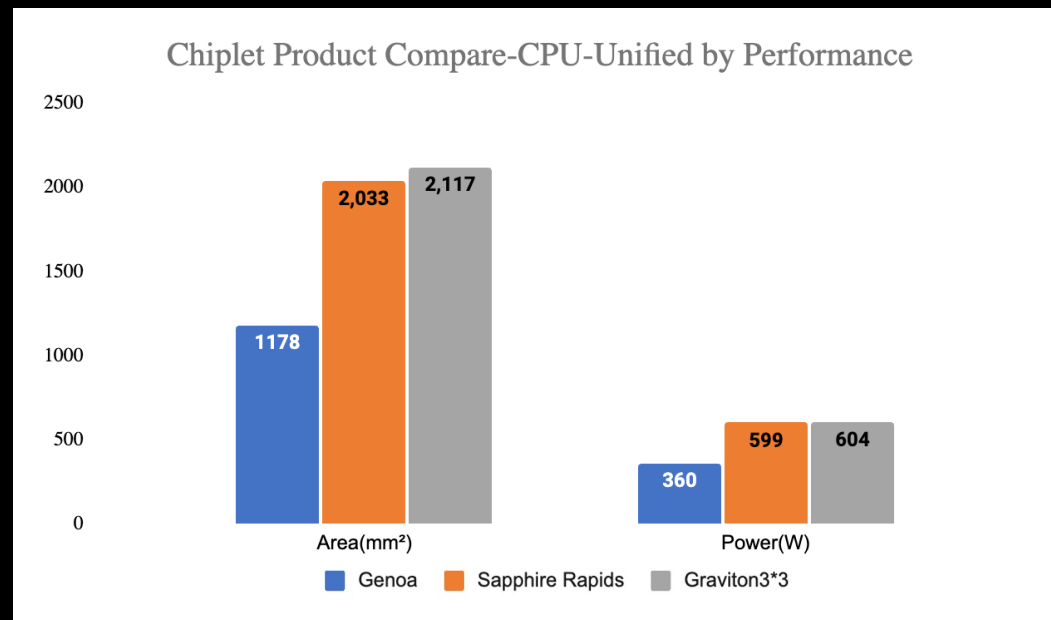

- ARM V1 Core
- L2 Cache 1MB
- L1 Instruction Cache 64KB
- L1 Data Cache 64 KB
- 282 mm<sup>2</sup>
- CMN 650 Fabric
- 8\*8 Mesh, 32 Node for Memory, 32 Node for CPU



- 1 Computing Die(N5) + 2 PCIe5 Controller Die + 4 DD5 Controller Die
  - Core Counts 64
  - Threads Counts 64
  - Memory Bandwidth: 300GB/s
  - Total L3 Cache 32MB
  - SPECrate@2017\_int\_base(3P) ~630
- 
- ARM V1 Core
  - L2 Cache 1MB
  - L1 Instruction Cache 64KB
  - L1 Data Cache 64 KB
  - 282 mm<sup>2</sup>
  - CMN 650 Fabric
  - 8\*8 Mesh, 32 Node for Memory, 32 Node for CPU



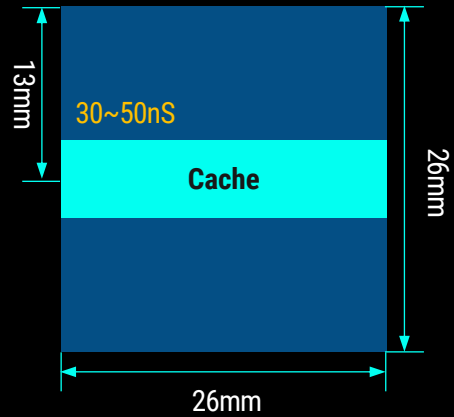
- 1 Nitro card with 3 gravition3 sockets, Nitro card manages 3 sockets simultaneously
- Optimize performance, energy efficiency in whole




- IO Die Central架构提供更好的性能，更小的面积，更低的功耗
- IO Die Central架构提供更好的多芯粒互联性能，更均衡的延时，更高的灵活性
- IO Die Central架构中，计算芯粒和互联芯粒采用不同的制程，量产成本更低
- Multi-Die架构设计更为简单，且在2-Die产品中具有性价比优势

提升芯片内传输效率 ▶ 3DIC


SoC



3DIC




MI300



Source: AMD

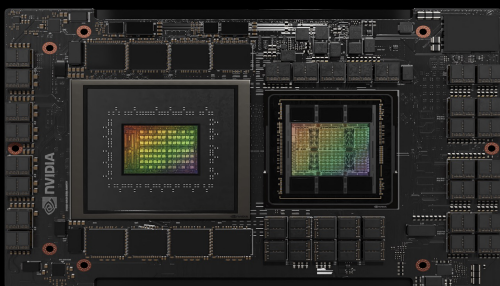
Ponte Vecchio




Source: Intel



## 提升计算效率 ▶ 异构计算



4 PetaFLOPS TE | 72 Arm CPUs  
96GB HBM3 | 576GB GPU Memory



### Chiplet: 异构计算的黄金搭档

- Tbps 级带宽, nS 级延时, 满足异构核间通信需求
- 模块化组合, 通用处理器+专用 DSA, 满足不同应用需求
- 模块化升级, 满足算法不断迭代需求

异构计算单元 Grace+Hopper, 实现提供 5 倍性能提升

Source: Nvidia

## Agenda

Kiwimoore

# 互联时代的挑战和关键技术

## Kiwi SoChiplet, 高性能互联平台

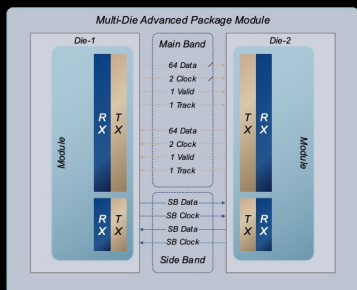


# 奇异摩尔, 全球领先的 Chiplet 高性能互联产品及解决方案商

- 以互联为中心, 基于 Chiplet 架构和高性能 RDMA 技术
- 提供全链路高性能互联产品及解决方案
- 助力客户持续提升性能、减少设计周期, 降低量产成本



## 高速 Die2Die IP 全面覆盖 2.x/2.5/3DChiplet 形态



UCIe Compatible

高带宽低延时低功耗

32Gbps, 0.5pJ, 5ns


全面覆盖

2.x/2.5/3D 等不同 Chiplet 封装形态

国际标准

支持Chiplet国际联盟, UCIe 标准

## 2.5D IO Die 高速互联芯粒



超高速 D2D 接口

Die2Die 2.5D, Multi-Channel, UCIe


全集成高速接口

DDR&PCIe MC/PHY, Multi-Channel

超高速互连网络

Coherent OCI

## 3D Base Die 3D 近存超高速互联芯粒



高密度 3D 堆叠

3D 堆叠实现 200% 高密度集成

全集成高速接口

高速 Memory & PCIe MC/PHY

超高速片间互连

Coherent OCI


分布式 3D 近存

数百 MB Cache  
3D 高带宽低延时

大电流 IVR

数百瓦高速大电流  
分布式供电网络

## nDSA 数据调度及网络加速芯粒



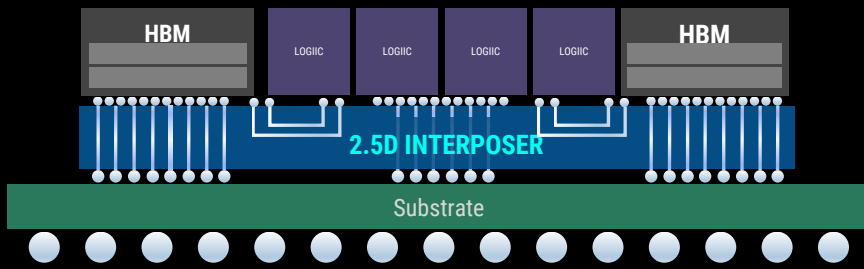
高性能 RDMA

集成多种硬件加速单元

基于 Chiplet 架构

可编程 DP Engine

## 2.5D Interposer 及 Chiplet 设计



丰富 Chiplet 项目设计经验

10+, Global

全流程 Turnkey 服务

设计-封装-测试-量产

Interposer 战略合作

解决量产问题

## 奇异摩尔与智原科技联合发布2.5D/3DIC 整体解决方案

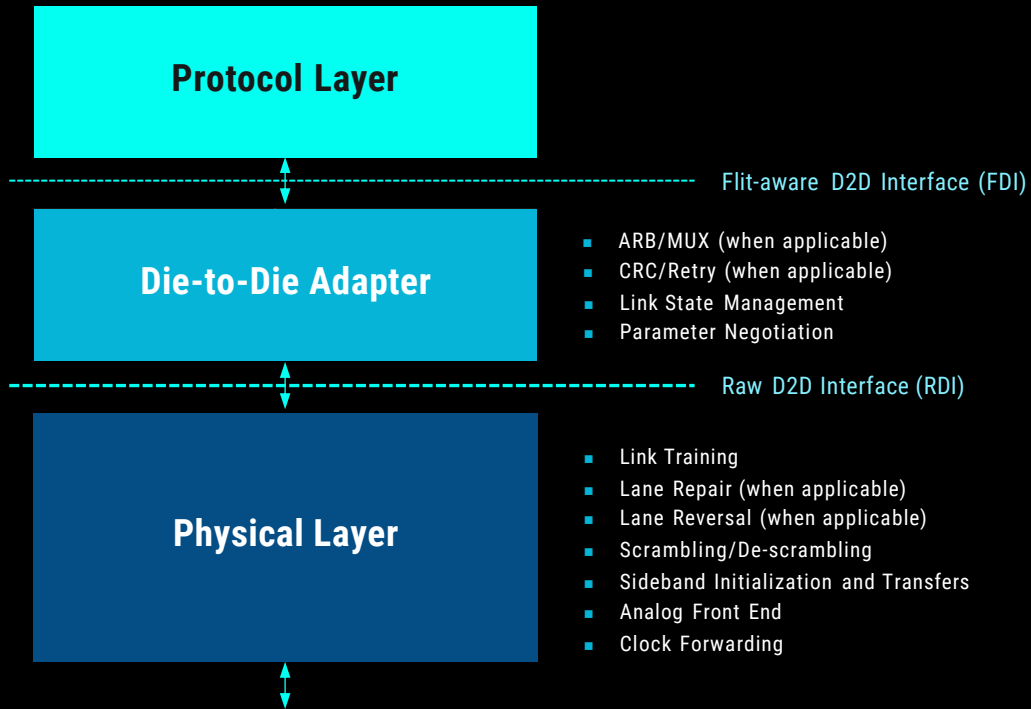
2023年11月11日，智原科技与奇异摩尔宣布共同推出 2.5D interposer 及 3DIC 整体解决方案，双方将基于晶圆对晶圆 3DIC 堆叠封装平台，为行业提供 2.5D interposer 及 3DIC “从设计、封装、测试至量产的全链路服务”。

**奇异摩尔 CEO 田陌晨表示：**很荣幸与智原科技达成战略合作。我们致力于与智原科技进一步扩大及深化合作伙伴关系。基于双方的合作，我们能更好的为客户提供从芯粒产品，设计、封装、测试到量产的全链路解决方案。希望双方能更深入的挖掘 Chiplet 与互联创新技术应用，促进 Chiplet 生态成熟和商业化落地。

**智原科技营运长林世钦表示：**非常高兴与奇异摩尔携手，共同发布 2.5D/3DIC 整体解决方案。凭借我们在 SoC 设计方面的专业知识，以及与晶圆和封测领域的顶尖企业的合作，能为 3DIC 先进封装服务提供全方位支持。这一合作标志着我们在小芯片整合领域能够更充分的发挥各种尖端应用潜力，以满足客户需求。



## 全球首批高性能 UCle V1.1 IP Controller & PHY



- Compatible UCle 1.1 Standard
- Support multi-package: standard & advantage
- Support for 4, 8, 12, 16, 24 and 32 GT/s data rates
- Supports PCIe, CXL, and streaming protocols
- Supports single and multiple PHY modules
- support link training, repair, redundance
- Link State Management

# 奇异摩尔

远见，超越芯所未见

