EMIB-T Roadmap, Custom HBM, HBM4 Packaging Challenges, Microfluidic Cooling, Photonic Interconnects, and More

EMIB-T 技术路线图、定制 HBM、HBM4 封装挑战、微流控冷却、光子互连及其他

ECTC 2026 Roundup, Intel, TSMC, SK Hynix, Samsung, Micron, Marvell, Lightmatter, Microsoft

ECTC 2026 年度综述：英特尔、台积电、SK 海力士、三星、美光、美满电子、Lightmatter、微软

AFZAL AHMAD, DC, GERALD WONG, AND DYLAN PATEL

AFZAL AHMAD、DC、GERALD WONG 与 DYLAN PATEL

JUL 03, 2026 2026 年 7 月 3 日 · PAID · 付费

46 3 Share 分享

As transistor density scaling has slowed, advanced packaging has become the primary scaling vector. However, AI accelerators have grown so large and require such fast interconnects that the package itself is now hitting limits. Circular interposers constrain package size and wafer utilization, HBM4E doubles the I/O count while increasing speed, and multi-kilowatt packages are overwhelming conventional cooling architectures.

随着晶体管密度提升速度放缓，先进封装已成为主要的扩展方向。然而，AI 加速器的规模日益庞大，且对互连速度要求极高，以至于封装本身正面临瓶颈。环形中介层限制了封装尺寸与晶圆利用率，HBM4E 在提升速度的同时使 I/O 数量翻倍，而功耗达数千瓦级别的封装正在挑战传统散热架构的极限。

ECTC is the industry’s premier conference for all things packaging. This year’s disclosures were closely aligned with upcoming commercial products. Intel provided an overview of EMIB-T integration, package scaling and future roadmaps. Marvell showed how custom HBM can move interface logic off the accelerator while shortening package routing. TSMC and Microsoft pushed coolant directly into silicon,

while Marvell and Lightmatter brought optical interconnects onto the package.

ECTC 是行业最权威的封装技术会议。今年公布的技术与即将上市的商业产品高度契合。英特尔概述了 EMIB-T 集成技术、封装扩展及未来路线图；美满电子展示了定制化 HBM 如何将接口逻辑移出加速器并缩短封装走线；台积电与微软将冷却液直接注入硅基板，而美满电子与光互连企业 Lightmatter 则将光学互连引入封装领域。

This roundup covers ECTC 2026 technologies most likely to shape AI accelerator packages over the next several years.

本综述聚焦 ECTC 2026 中极有可能在未来数年内塑造 AI 加速器封装的技术成果。

SemiAnalysis is a reader-supported publication. To receive new posts and support our work, consider becoming a free or paid subscriber.

SemiAnalysis 是一份由读者支持的刊物。如需接收新文章并支持我们的工作，可考虑成为免费或付费订阅用户。

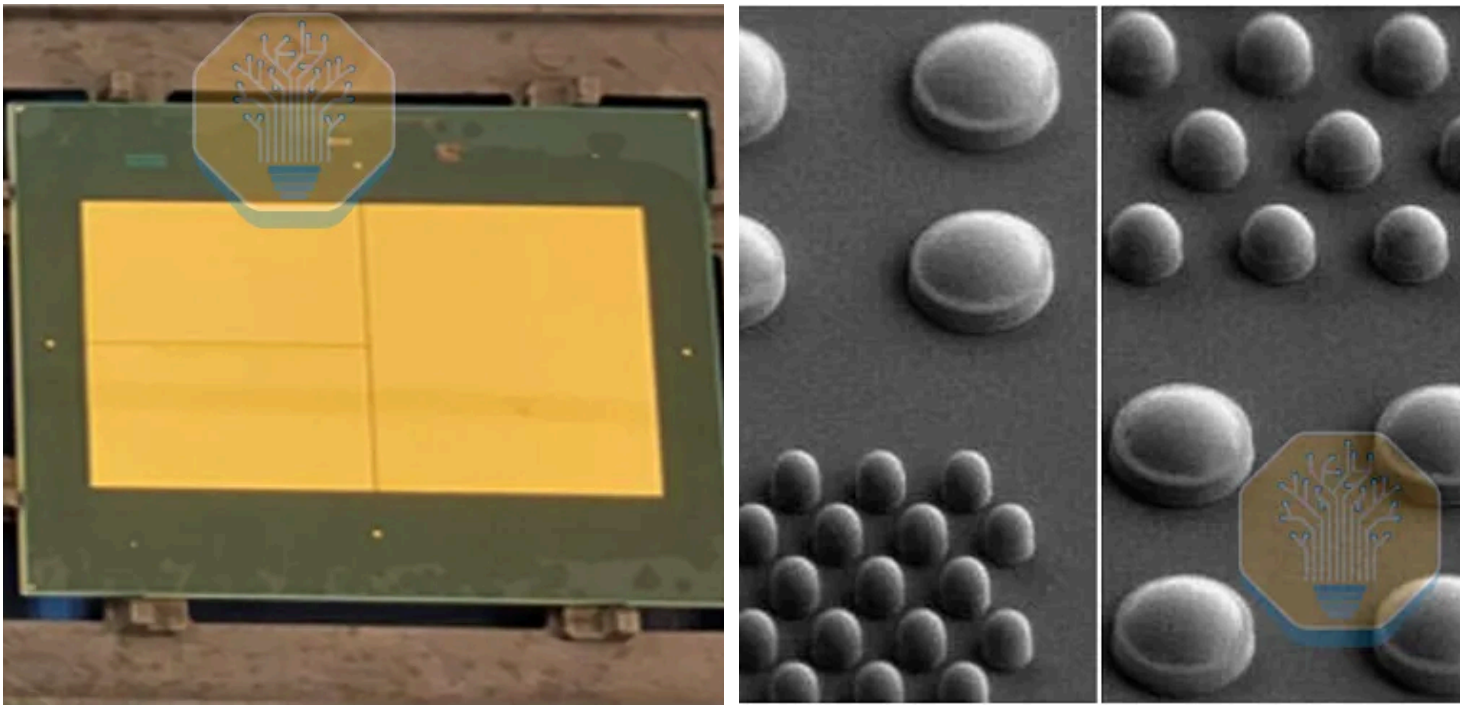
Intel EMIB-T 英特尔 EMIB-T

Intel was the largest corporate presenter at ECTC. The key disclosure was EMIB-T. This is the next generation of EMIB with through-silicon vias (TSVs). After the initial announcement, Intel has filled in more of the architecture and roadmap, including tighter bump pitch, larger packages and on-bridge capabilities. Their disclosures show why EMIB-T is expected to be used in Google's TPU v9, and why it is the most credible alternative to TSMC's CoWoS platform for large-package AI accelerators.

英特尔是 ECTC 大会上最大的企业参展方。其核心技术发布是 EMIB-T。这是下一代采用硅通孔（TSV）技术的 EMIB。在初步发布后，英特尔进一步充实了该技术的架构与路线图，包括更密集的凸点间距、更大尺寸的封装以及桥接功能。他们的披露表明，为何 EMIB-T 有望被用于谷歌 TPU v9，以及为何它最有可能成为台积电 CoWoS 平台在大型封装 AI 加速器领域的有力替代方案。

EMIB-T scaling test vehicle with 2× reticle silicon content. Top-down SEM images show 110, 55 and 36 µm mixed bump pitches. Source: Intel, “Scaling the EMIB-T Advanced Packaging Technology to Address the Future HPC/AI Demand,” ECTC 2026

EMIB-T 扩展测试载体（含 2 倍掩模版硅面积）。俯视扫描电镜图像显示了 110 微米、55 微米和 36 微米混合凸点间距。来源：英特尔，《扩展 EMIB-T 先进封装技术以满足未来 HPC/AI 需求》，ECTC 2026

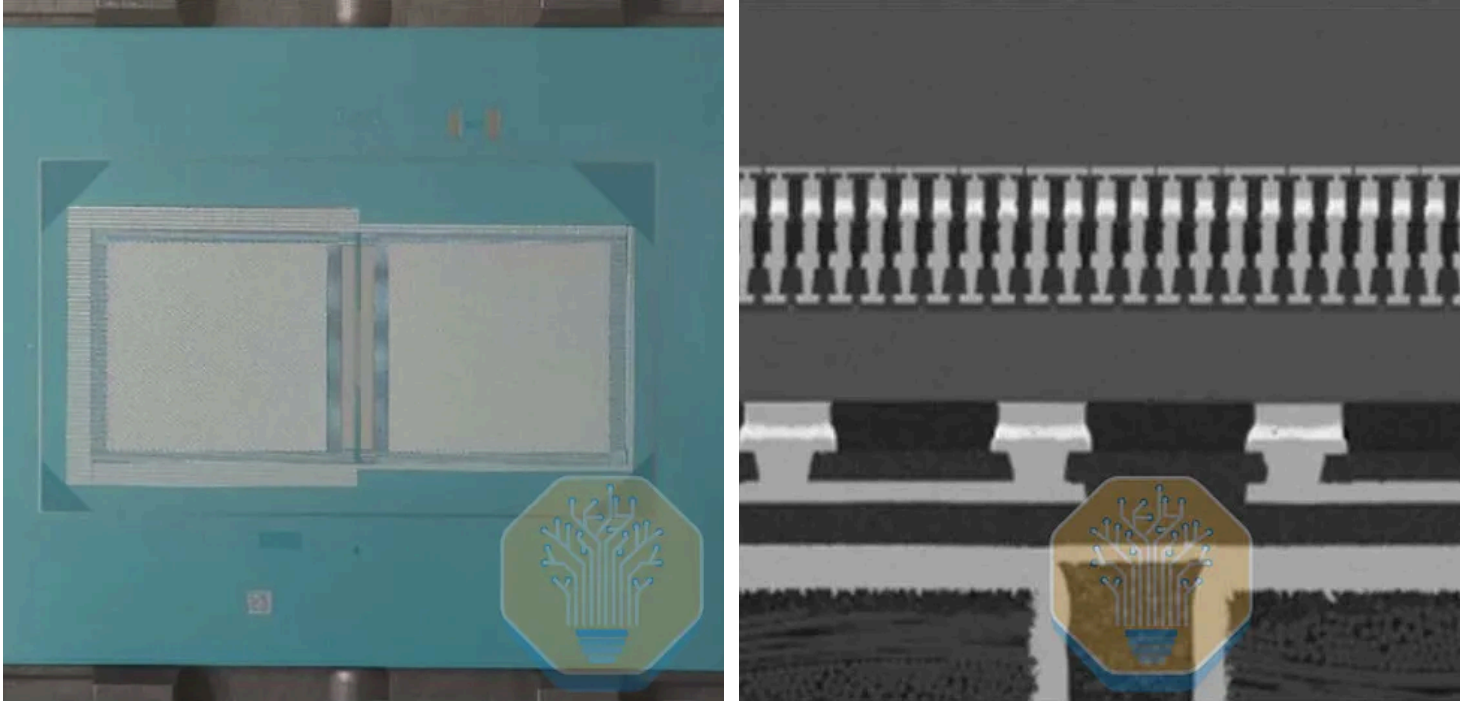


EMIB-T scaling test vehicle with 2× reticle silicon content. Top-down SEM images show 110, 55 and 36 µm mixed bump pitches. Source: Intel, “Scaling the EMIB-T Advanced Packaging Technology to Address the Future HPC/AI Demand,” ECTC 2026

EMIB-T 扩展测试载体（含 2 倍掩模版硅面积）。俯视扫描电镜图像显示了 110 微米、55 微米和 36 微米混合凸点间距。来源：英特尔，《扩展 EMIB-T 先进封装技术以满足未来 HPC/AI 需求》，ECTC 2026

Intel has validated EMIB-T at a 36/35 µm bump pitch on a package with 2× reticle-sized silicon. This is a reduction over the 45 µm pitch used in Granite Rapids, and a 65% increase in bump density. Granite Rapids-AP is a large package, measuring 70 mm × 105 mm, or slightly less than 9 reticles of area. Validation for 36/35 µm bump pitches is now expanding to 4.5× reticle silicon packages, with certification targeted by the end of 2026.

英特尔已在采用 2 倍掩模板尺寸硅片的封装上，成功验证了 36/35 微米凸点间距的 EMIB-T 技术。相比 Granite Rapids 使用的 45 微米间距，这实现了凸点密度 65%的提升。Granite Rapids-AP 是一款大型封装，尺寸为 70 毫米×105 毫米，面积略小于 9 个掩模板。目前 36/35 微米凸点间距的验证工作正扩展至 4.5 倍掩模板尺寸的硅片封装，预计于 2026 年底完成认证。



EMIB-T 25 µm bump pitch test vehicle with two 1-reticle silicon dies connected by a 3 mm × 18 mm bridge. SEM cross-section shows the post-assembly bump and via structure. Source: Intel, “Scaling the EMIB-T Advanced Packaging Technology to Address the Future HPC/AI Demand,” ECTC 2026

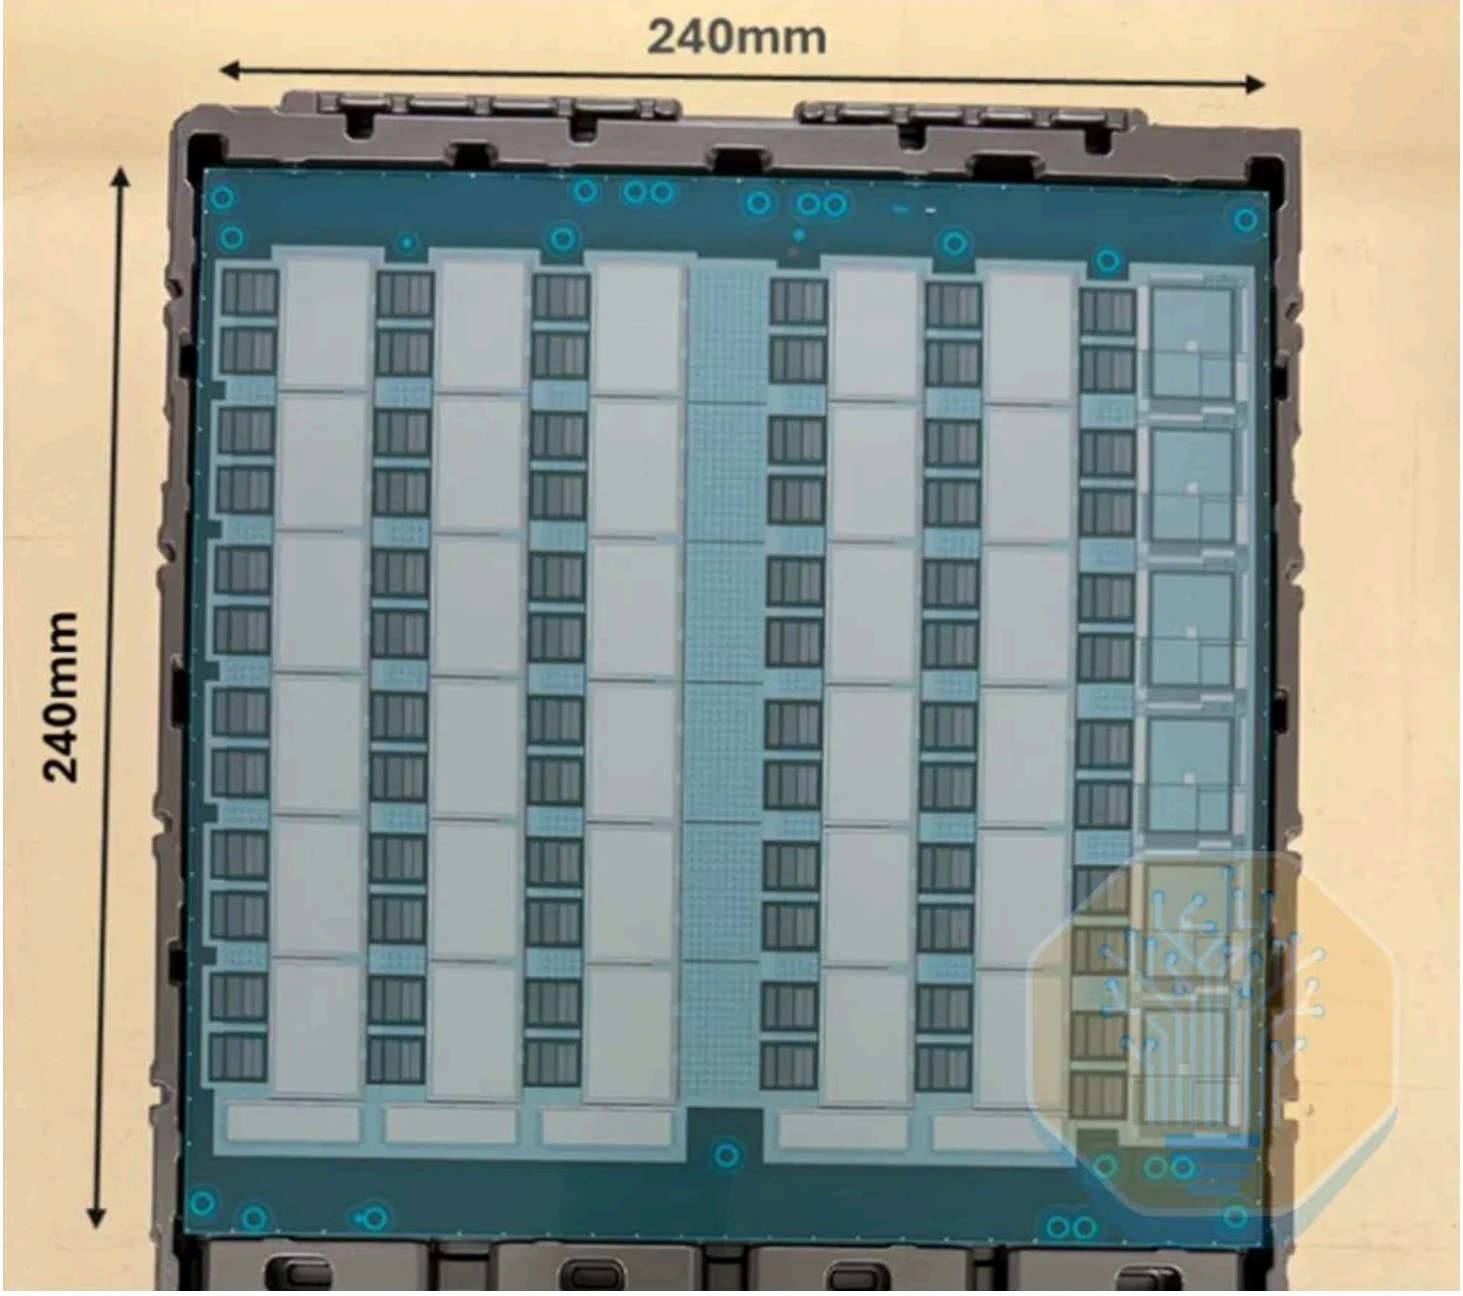
采用 25 微米凸点间距的 EMIB-T 测试载体，两个 1 倍掩模板硅芯片通过 3 毫米×18 毫米的桥接器连接。扫描电镜截面图展示了组装后的凸点与通孔结构。来源：英特尔，《面向未来 HPC/AI 需求扩展 EMIB-T 先进封装技术》，ECTC 2026

The next pitch step is also underway, with Intel testing 25 µm bump pitch on a vehicle with two 1-reticle silicon dies connected through a single 3 mm × 18 mm EMIB-T bridge.

下一个间距阶段也已启动：英特尔正在测试 25 微米凸点间距的测试载体，该载体通过单个 3 毫米×18 毫米的 EMIB-T 桥接器连接两个 1 倍掩模板硅芯片。

Further scaling gets harder. Below 25 µm, the solder volume in each bump becomes very small. Shorts, opens, and assembly-driven yield loss become much more likely. EMIB-T can keep scaling, but the limiter shifts from bridge routing density to bump formation, placement accuracy, and assembly yield.

进一步缩小间距将面临更大挑战。当间距低于 25 微米时，每个凸点中的焊料体积会变得非常小。短路、开路以及组装引发的良率损失将显著增加。EMIB-T 技术虽能持续缩小，但制约因素将从桥接布线密度转向凸点成型、贴装精度及组装工艺。



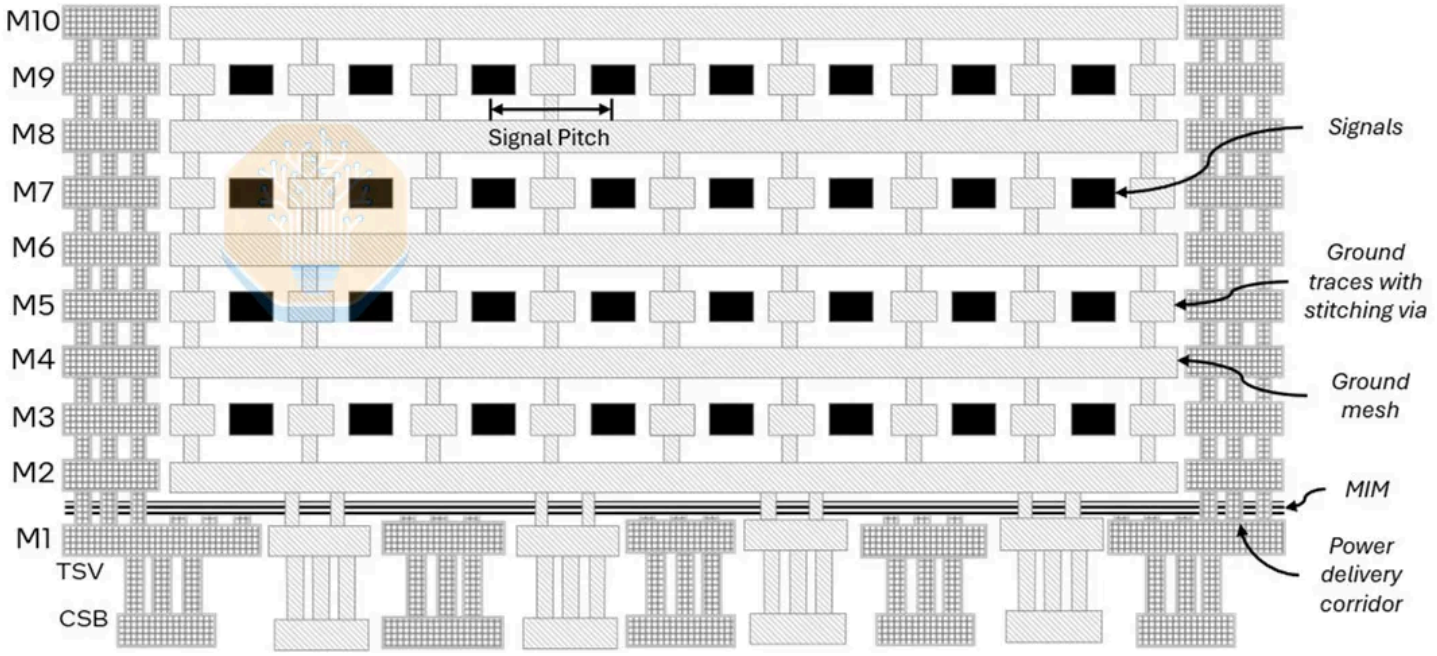
Intel's 240 mm × 240 mm quarter-panel EMIB-T test vehicle. Source: Intel, "Scaling the EMIB-T Advanced Packaging Technology to Address the Future HPC/AI Demand," ECTC 2026

英特尔 240 毫米×240 毫米四分之一面板 EMIB-T 测试载具。来源：英特尔，《扩展 EMIB-T 先进封装技术以满足未来高性能计算/人工智能需求》，ECTC 2026

Intel also showed how large EMIB-T packages can become. Full panel-sized packages are possible, but Intel framed quarter-panel packages as the practical target. It showed a 240 mm × 240 mm test vehicle, equivalent to roughly 67 reticles of area. However, we observed that the sample at the booth exhibited severe warpage. At that size, the bridge is only one part of the problem. Substrate handling, warpage, overlay, and panel-level patterning become first-order constraints. Intel is also evaluating advanced lithography approaches to keep overlay tight enough for these large substrates at

quarter-panel and even full-panel scale.

英特尔还展示了 EMIB-T 封装能够达到的规模。全面板级封装虽有可能，但英特尔将四分之一面板封装定位为实际目标。其展示的 240 毫米×240 毫米测试载体，面积约相当于 67 个掩模版。然而，我们注意到展台样品存在严重翘曲问题。在此尺寸下，桥接仅属问题之一。基板处理、翘曲控制、对准精度及面板级图案化已成为首要制约因素。为在四分之一面板乃至全面板尺寸的大型基板上保持足够的对准精度，英特尔正在评估先进光刻技术方案。



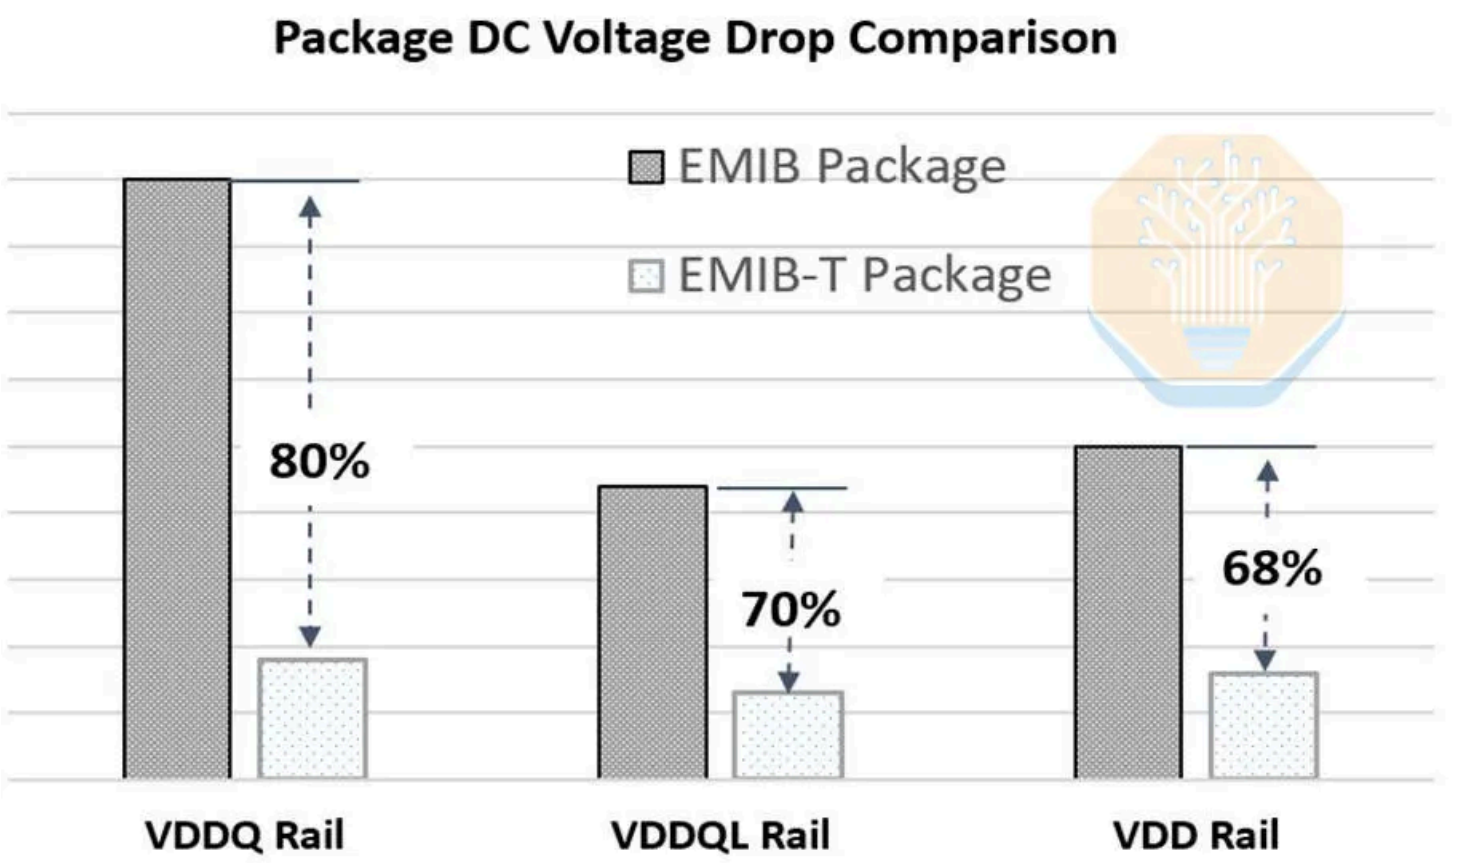
EMIB-T stack-up with routing, power mesh, TSVs, and MIM layers. Source: Intel, “Enabling 12+Gb/s HBM4E with EMIB-T Advanced Packaging Technology,” ECTC 2026

EMIB-T 堆叠结构包含走线、电源网格、TSV 和 MIM 层。来源：英特尔，《利用 EMIB-T 先进封装技术实现 12+Gb/s HBM4E》，ECTC 2026

While bump pitch and package size are important, the bridges matter just as much. EMIB-T is much more complex than the EMIB used in current products. It adds TSVs, more metal layers, a power mesh, and MIM capacitor layers so the bridge can carry both high-density signals and vertical power delivery. Intel showed a cross-section with 10 metal layers, including 4 routing layers, and MIM capacitors between M1 and

M2. Intel focused on its improvements for HBM4E.

尽管凸点间距和封装尺寸至关重要，但桥接器的同样举足轻重。EMIB-T 比当前产品中使用的 EMIB 复杂得多。它增加了硅通孔（TSV）、更多金属层、电源网格和金属-绝缘体-金属（MIM）电容层，使桥接器既能传输高密度信号，又能实现垂直供电。英特尔展示的横截面包含 10 个金属层（包括 4 个布线层）以及位于 M1 与 M2 层之间的 MIM 电容。英特尔重点介绍了针对 HBM4E 的改进方案。



Package DC voltage drop using conventional EMIB vs EMIB-T. Source: Intel, “Enabling 12+Gb/s HBM4E with EMIB-T Advanced Packaging Technology,” ECTC 2026

采用传统 EMIB 与 EMIB-T 时的封装直流电压降。来源：Intel，《以 EMIB-T 先进封装技术实现 12+Gb/s HBM4E》，ECTC 2026

The “T” in EMIB-T is for TSV. Their job is power delivery. In conventional EMIB, power is delivered vertically through the substrate in non-bridge regions, but power delivery near the bridges must spread laterally through the package and die-side routing. With TSVs in the bridge, power can be delivered directly through the bridge, reducing the current path distance significantly. Intel claims that DC voltage drop can

be reduced by 68-80% with these TSVs.

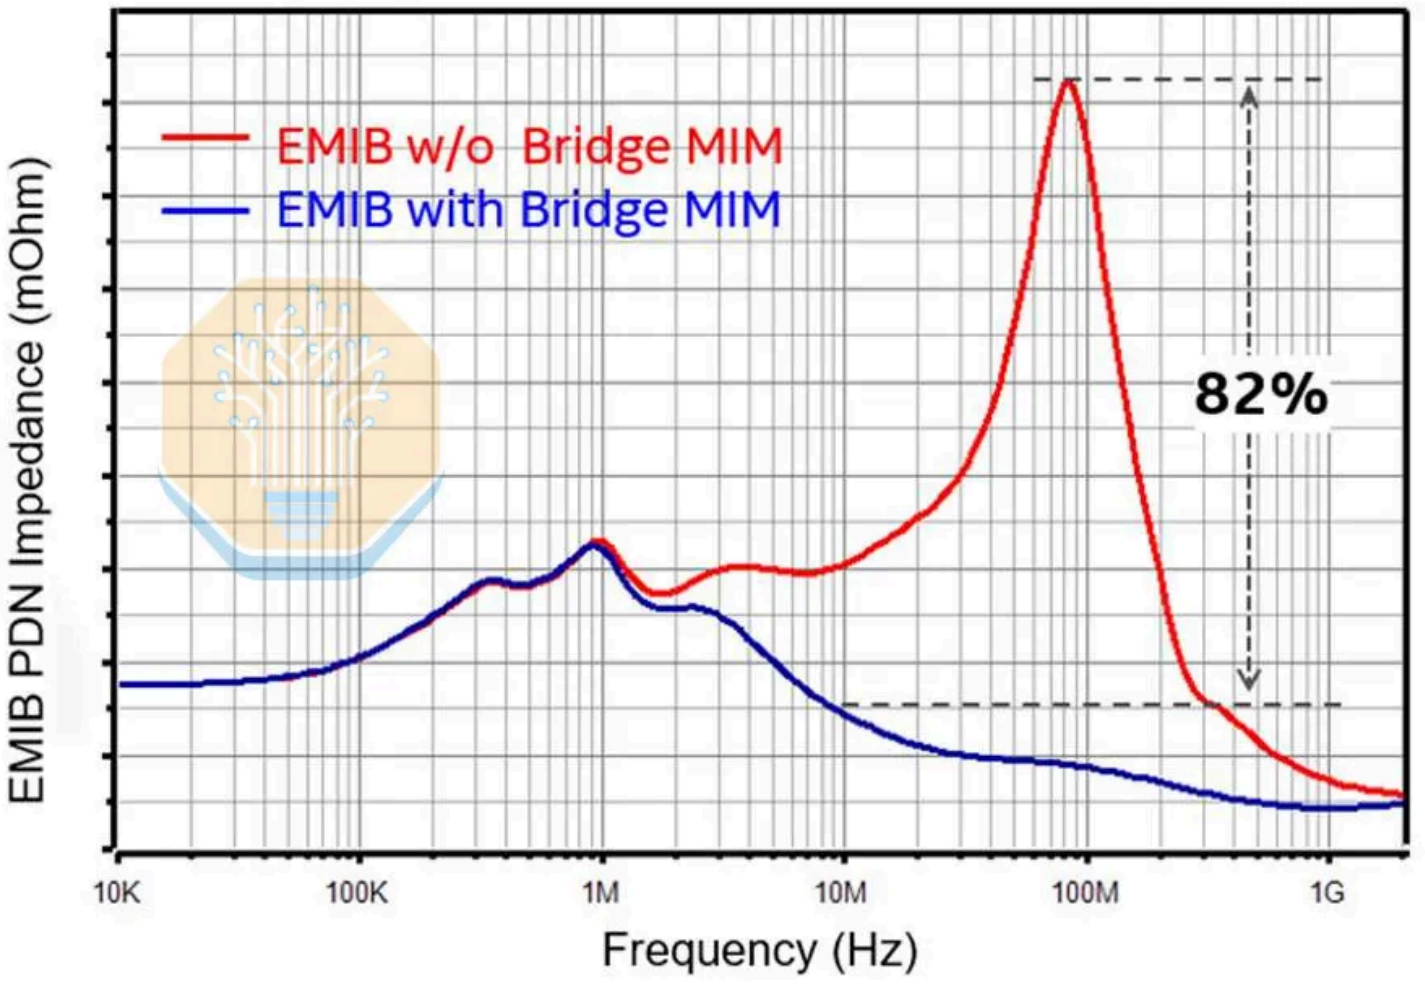
EMIB-T 中的“T”代表 TSV（硅通孔），其职责是供电传输。传统 EMIB 架构中，电源通过基板在非桥接区域垂直传输，但桥接区域附近的电源必须通过封装和芯片侧布线横向扩散。通过在桥接结构中嵌入 TSV，电源可直接通过桥接件垂直传输，大幅缩短电流路径。英特尔宣称，采用该技术后直流压降可降低 68-80%。

HBM4E is difficult because the interconnect must scale both signal density and power delivery at the same time. HBM4 doubles the pin count versus HBM3, and the PHY needs additional power rails such as VDDQ and VDDQL. These rails take routing area away from signals, raising signal density in the remaining space.

HBM4E 存在难度，因为互连必须同时扩展信号密度和电力输送。与 HBM3 相比，HBM4 的引脚数量翻倍，且物理层需要额外的电源轨，例如 VDDQ 和 VDDQL。这些电源轨占用了信号的路由面积，导致剩余空间内的信号密度增加。

To combat this, Intel did not route every HBM channel the same way. It placed the longest signal paths on cleaner routing layers. On M9, only ~28% of the longest channel length runs through the tightest routing region, while on lower layers such as M3, that number rises to ~84%, but those channels are shorter. This keeps crosstalk and insertion loss from being dominated by the worst routes.

为解决这一问题，英特尔并未对所有 HBM 通道采用相同布线方式。它将最长的信号路径布置在干扰较少的布线层上。在 M9 层，仅~28%的最长通道长度经过最密集布线区域，而在 M3 等较低层，该比例升至~84%，但这些通道长度更短。此举可避免串扰和插入损耗被最差路径主导。

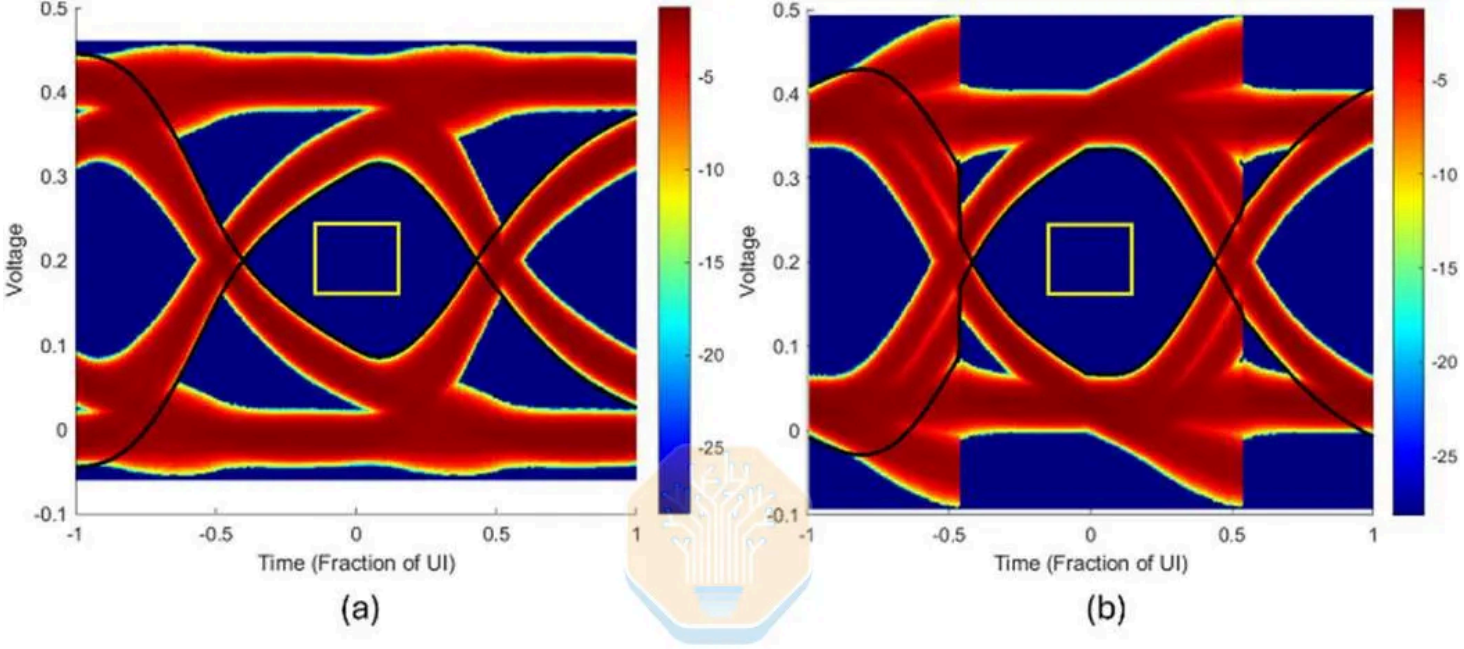


EMIB-T package PDN AC impedance for HBM4E with and without bridge MIM capacitors. Source: Intel, "Enabling 12+Gb/s HBM4E with EMIB-T Advanced Packaging Technology," ECTC 2026

EMIB-T 封装中 HBM4E 的 PDN 交流阻抗（包含与不包含桥接 MIM 电容两种情形）。来源：英特尔，《采用 EMIB-T 先进封装技术实现 12+Gb/s HBM4E》，ECTC 2026

Power delivery is also moving into the bridge. EMIB-M introduced metal-insulator-metal (MIM) capacitors between M1 and M2, and EMIB-T builds on that approach. Intel disclosed a capacitance density of 500 nF/mm^2, roughly comparable to Intel 18A MIMs. Intel claims these on-bridge capacitors improve power delivery network (PDN) AC impedance by more than 82% versus an EMIB-T package without bridge MIM capacitors.

电源传输也正在向桥接技术演进。EMIB-M 在 M1 与 M2 金属层之间引入了金属-绝缘体-金属（MIM）电容器，而 EMIB-T 则在此基础上进一步升级。英特尔披露其电容密度达到 500 nF/mm^2，大致与 Intel 18A 工艺的 MIM 电容相当。英特尔宣称，与未采用桥接 MIM 电容器的 EMIB-T 封装相比，这些桥接电容器可将供电网络（PDN）的交流阻抗降低超过 82%。



Signaling performance of HBM4E EMIB-T channel at 12 Gb/s without and with DFE. Source: Intel, "Enabling 12+Gb/s HBM4E with EMIB-T Advanced Packaging Technology," ECTC 2026

HBM4E EMIB-T 通道在 12 Gb/s 速率下未启用和启用 DFE 时的信号性能。来源：英特尔，《利用 EMIB-T 先进封装技术实现 12+Gb/s HBM4E》，ECTC 2026

Intel also simulated EMIB-T with HBM4E. At 12 Gb/s, Intel showed ~67% UI eye width without receiver equalization. With a one-tap decision feedback equalizer (DFE), that improves to ~72.5%. A DFE is a receiver-side circuit that reduces interference from prior bits after the signal passes through the package channel.

英特尔还使用 HBM4E 模拟了 EMIB-T。在 12 Gb/s 速率下，英特尔展示了无需接收器均衡即可达到~67%的 UI 眼宽。采用单抽头判决反馈均衡器（DFE）后，该指标提升至~72.5%。DFE 是一种接收端电路，可在信号经过封装通道后，降低前序比特位造成的干扰。

Intel also simulated higher speeds of 12.8 Gb/s, 14 Gb/s and 16 Gb/s. The UI eye width remained above 60% at all tested speeds, with modest pad capacitance reduction.

英特尔还模拟了 12.8 Gb/s、14 Gb/s 和 16 Gb/s 的更高速度。在所有测试速度下，UI 眼图宽度均保持在 60%以上，且焊盘电容略有降低。

Intel's EMIB roadmap goes beyond passive bridges with only routing and capacitors. Future versions include higher-density on-bridge MIM capacitors, larger high-aspect-ratio bridge dies, sub-25 µm bump pitch, active bridges, and embedded voltage regulators inside EMIB dies. Intel has also disclosed substrate-core embedded deep-trench capacitor (DTC) concepts and >2500 nF/mm^2 eMIM-T capacitors embedded

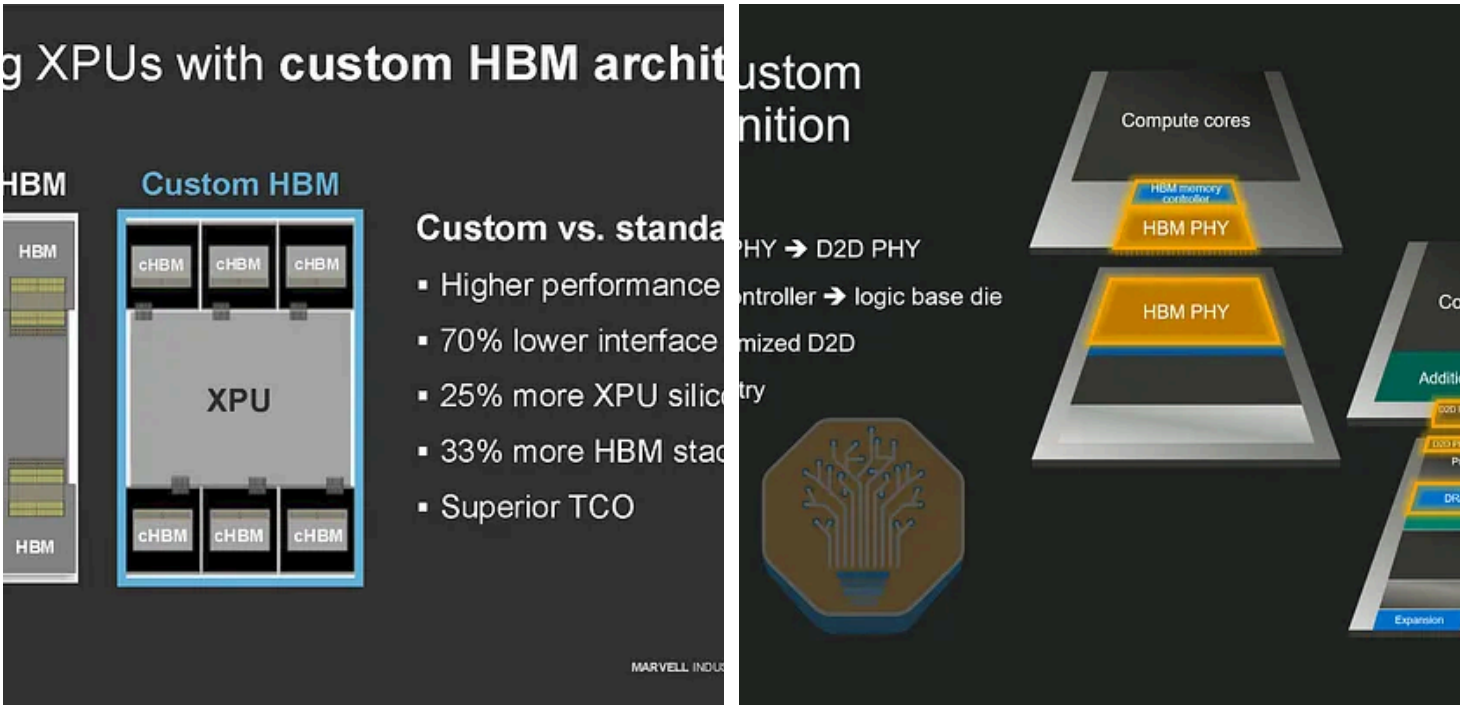
below the base die, although we have not seen either in shipping EMIB products.

英特尔 EMIB 路线图已超越仅包含布线和电容器的无源桥接方案。未来版本将包括更高密度的板载 MIM 电容器、更大高深宽比的桥接芯片、低于 25 微米的凸点间距、有源桥接，以及嵌入 EMIB 芯片内部的电压调节器。英特尔还公开了基板核心嵌入式深沟槽电容器（DTC）概念和基底层下嵌入的>2500 nF/mm^2 eMIM-T 电容器，不过这两项技术均未在量产 EMIB 产品中出现。

EMIB-T is still behind TSMC's CoWoS platform on several vectors. TSMC has already deployed DTC/eDTC integration and is further along on integrated voltage regulators and active local silicon interconnect (LSI). EMIB-T narrows the gap, but Intel is still catching up to an ecosystem that has been executing in volume for years.

EMIB-T 在多个维度上仍落后于台积电的 CoWoS 平台。台积电已部署 DTC/eDTC 集成技术，并在集成电压调节器和有源本地硅互连（LSI）方面走得更远。EMIB-T 缩小了差距，但英特尔仍在追赶一个已大规模量产多年的生态系统。

Marvell Custom HBM　Marvell 定制 HBM



Marvell Custom HBM Summary and Base Die Floorplan. Source: Marvell, Industry Analyst Day 2024 and Hot Chips 2025

Marvell 定制 HBM 总结与基底芯片布局。来源：Marvell，2024 年行业分析师日及 2025 年 Hot Chips 大会

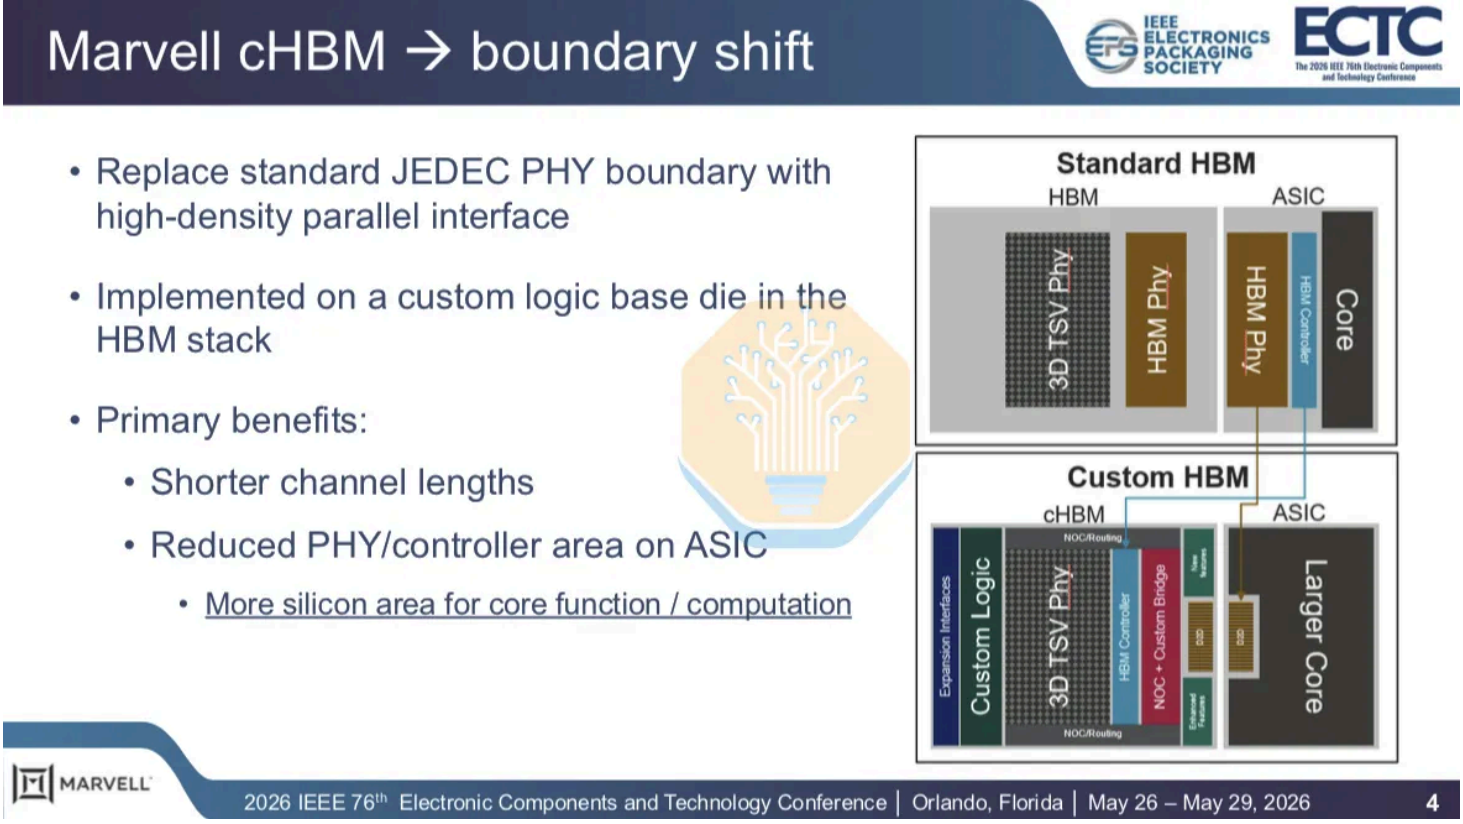
At Marvell's 2024 Industry Analyst Day, Marvell announced custom HBM. At the time, it was a vague claim with few technical details. HBM had been built around JEDEC compatibility: standard DRAM stacks from the memory vendors, standard HBM PHYs

on the accelerator, and a fixed wide interface between them. At Hot Chips 2025, Marvell showed the floorplan of the custom base die. At ECTC, it finally provided the package-level details for custom HBM4E.

在 Marvell 2024 年行业分析师日上，Marvell 宣布推出定制化 HBM。当时这只是一个模糊的声明，几乎没有透露技术细节。HBM 一直围绕 JEDEC 兼容性构建：来自内存供应商的标准 DRAM 堆叠、加速器上的标准 HBM PHY，以及两者之间固定的宽接口。在 Hot Chips 2025 会议上，Marvell 展示了定制化基础芯片的版图。而在 ECTC 上，该公司终于提供了定制化 HBM4E 的封装级细节。

The JEDEC specification fixes the interface between the HBM stack and host. That is good for interoperability: any memory vendor's HBM can be paired with any compliant host. However, it is bad for power, performance and area. The host ASIC must implement the standard HBM PHY and route a very wide, parallel interface with standardized pad placement and breakout rules. As packages get larger and HBM speeds rise, that fixed boundary makes it harder to optimize shoreline, routing density, power delivery and signal integrity.

JEDEC 规范固定了 HBM 堆栈与主控之间的接口。这对于互操作性有利：任何内存供应商的 HBM 都能与任何兼容的主控配对。然而，这在功耗、性能和面积方面存在弊端。主控 ASIC 必须实现标准 HBM 物理层，并按照标准化焊盘布局和扇出规则布设极宽并行接口。随着封装尺寸增大和 HBM 速率提升，这种固定边界使得优化海岸线布局、布线密度、供电网络及信号完整性变得更加困难。

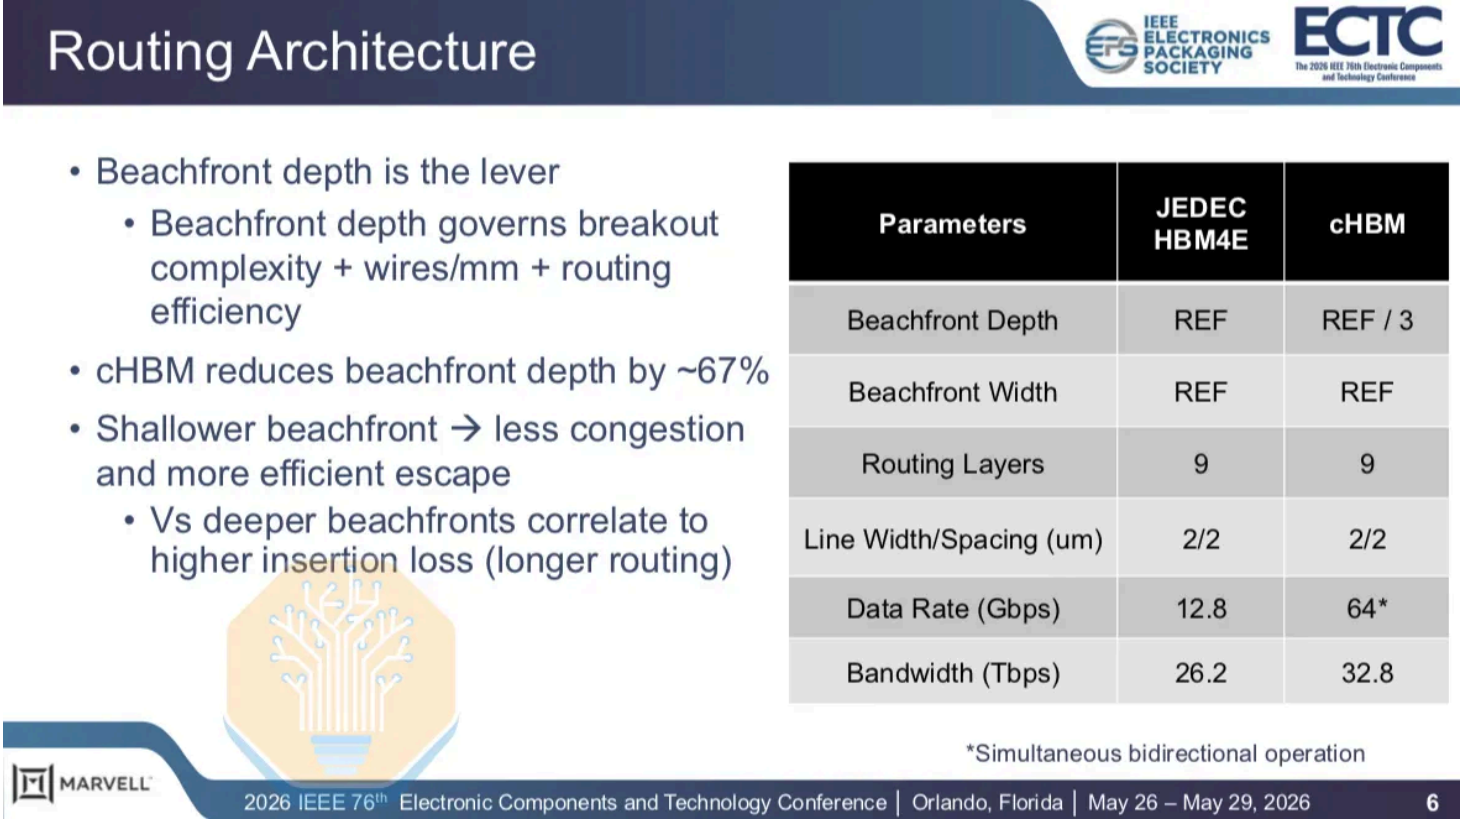


Standard HBM vs Custom HBM base die and host floorplan. Source: Marvell, "Marvell custom HBM routing and signal integrity analysis," ECTC 2026

标准 HBM 与定制 HBM 基础芯片及主芯片布局对比。来源：Marvell，《Marvell 定制 HBM 布线及信号完整性分析》，ECTC 2026

Custom HBM keeps the DRAM core dies unchanged. Instead, a custom base die, with an optimized die-to-die interface, is fabricated on an advanced logic process. The custom base die can integrate the HBM controller, management and monitoring capabilities, custom logic, and expansion interfaces.

定制 HBM 保持 DRAM 核心芯片不变，而是在先进逻辑工艺上制造定制基底芯片，配备优化芯片间接口。定制基底芯片可集成 HBM 控制器、管理与监控功能、定制逻辑及扩展接口。



JEDEC HBM4E vs Marvell Custom HBM. Source: Marvell, "Marvell custom HBM routing and signal integrity analysis," ECTC 2026

JEDEC HBM4E 与 Marvell 定制 HBM 对比。来源：Marvell，"Marvell 定制 HBM 布线与信号完整性分析，"ECTC 2026

Marvell claims this reduces the host ASIC footprint dedicated to HBM PHYs and associated logic by ~60%, directly freeing up area for more compute, cache or I/O. The custom interface moves much of the memory-side interface into the HBM base die.

Marvell 声称这可将专用于 HBM PHY 及相关逻辑的主机 ASIC 面积减少~60%，直接释放出更多空间用于计算、缓存或 I/O。定制接口将内存侧接口的大部分功能移入 HBM 基底芯片。

Marvell's example uses 1024 channels at 32 Gb/s, reaching 4.1 TB/s, equivalent to a 2048-bit JEDEC HBM4(E) interface at 16 Gb/s.

Marvell 的实例采用 1024 通道、32 Gb/s 速率，达成 4.1 TB/s 带宽，相当于 2048 位 JEDEC HBM4(E)接口在 16 Gb/s 速率下的性能。

The package also gets easier to route, with the custom interface shortening the interposer channel length from 6.5 mm to 1.5 mm, allowing Marvell to keep the same 9

routing layers and 2/2 μm line/space (L/S) while increasing bandwidth.

封装也更易于布线，定制接口将中介层通道长度从 6.5 毫米缩短至 1.5 毫米，使 Marvell 能够保持相同的 9 层布线层和 2/2 微米线宽/线距（L/S），同时提升带宽。

Marvell's example uses an organic redistribution layer (RDL) interposer instead of silicon, reducing packaging costs. Organic RDL is limited to much coarser L/S than the silicon interposers in CoWoS-S or the silicon bridges in CoWoS-L and EMIB-T, which makes layout harder. Marvell relies on customized shielding and routing patterns in different sections to maximize bandwidth density while controlling crosstalk.

Marvell 的例子采用有机再分布层（RDL）中介层替代硅基方案，从而降低封装成本。与 CoWoS-S 中的硅中介层或 CoWoS-L 及 EMIB-T 中的硅桥相比，有机 RDL 的线宽/线距（L/S）要大得多，这增加了布线难度。Marvell 通过不同区域定制化屏蔽与布线图案，在控制串扰的同时最大化带宽密度。



Nvidia Rubin GPU wafer. Source: TSMC, Computex 2026. Photo: SemiAnalysis

英伟达 Rubin GPU 晶圆。来源：台积电，2026 年台北国际电脑展。图片：SemiAnalysis

At GTC, Nvidia announced that Feynman would use custom HBM. Nvidia's reasons are likely similar to Marvell: higher bandwidth, lower power, and less accelerator die

area dedicated to HBM. We estimate that ~16% of Rubin GPU die area is dedicated to HBM-related logic and PHYs. Custom HBM allows Nvidia to offload much of that burden onto the HBM base die.

在 GTC 大会上，英伟达宣布 Feynman 将采用定制 HBM 内存。其考量与 Marvell 相仿：旨在实现更高带宽、更低功耗，同时减少加速器芯片中 HBM 相关逻辑及物理层的占用面积。据我们估算，~16%的 Rubin GPU 芯片面积需分配给 HBM 相关逻辑与物理层。定制 HBM 技术使英伟达能够将这部分负担转移至 HBM 基底芯片。

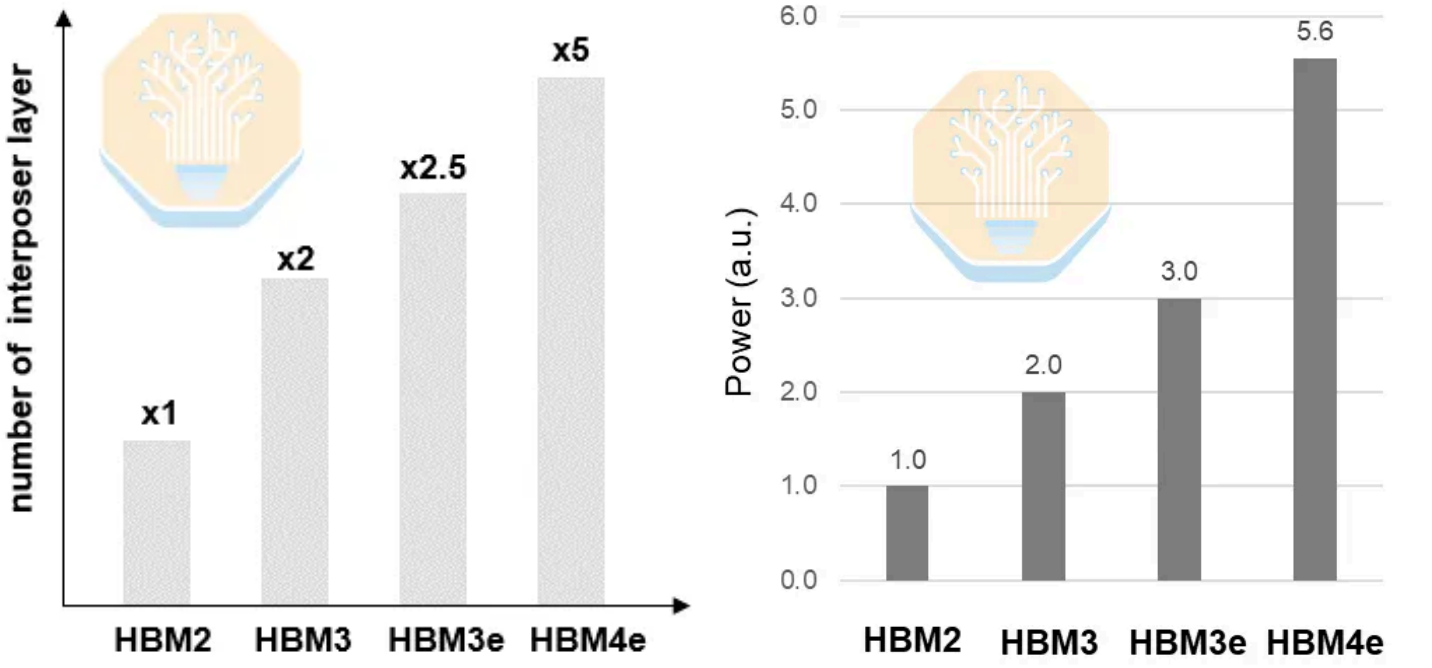
Custom HBM also enables expansion interfaces beyond the standard HBM link. Instead of forcing all memory traffic through the limited accelerator die shoreline, the base die can act as a secondary memory controller and fan out to additional memory. That memory could be higher-capacity, lower-bandwidth LPDDR, or even a second layer of HBM. This gives the accelerator a path to add capacity without consuming the valuable die shoreline needed for external I/O. This is relevant to AMD's upcoming MI450 and future MI500 GPUs, which will support LPDDR for increased memory capacity.

定制 HBM 还支持超越标准 HBM 链路的扩展接口。它不是强制所有内存流量通过受限的加速器芯片边缘（die shoreline），而是让基础芯片充当二级内存控制器，并扇出到额外内存。这些内存可以是容量更大、带宽更低的 LPDDR，甚至可以是第二层 HBM。这让加速器无需消耗用于外部 I/O 的宝贵芯片边缘空间，即可增加容量。这与 AMD 即将推出的 MI450 及未来 MI500 GPU 相关，这些产品将支持 LPDDR 以提升内存容量。

SemiAnalysis is a reader-supported publication. To receive new posts and support our work, consider becoming a free or paid subscriber.

SemiAnalysis 是一份由读者支持的刊物。如需接收新文章并支持我们的工作，可考虑成为免费或付费订阅用户。

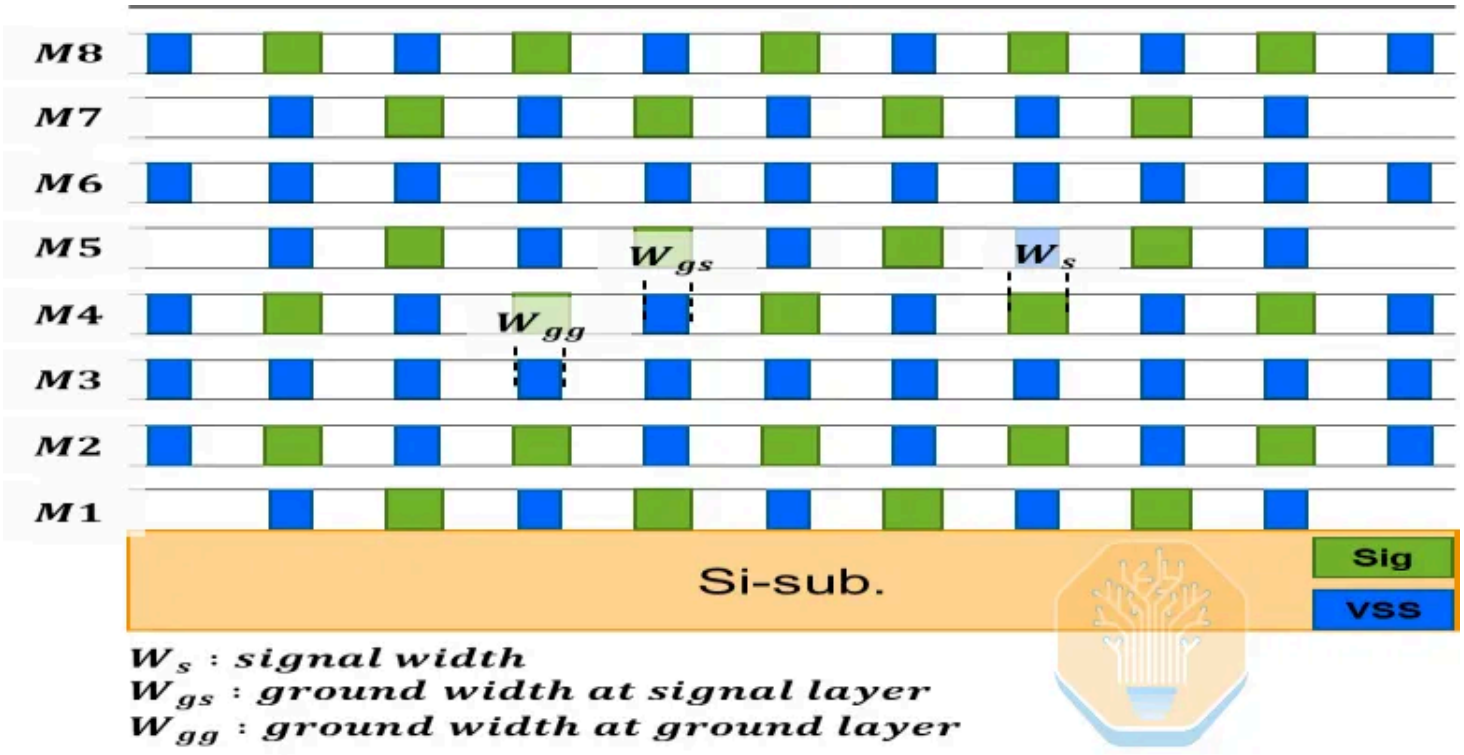
Samsung HBM Interposer 三星 HBM 中介层



Interposer and power scaling from HBM2 to HBM4E. Source: Samsung, "Advanced SI/PI Interposer Design Solution Enabling High-Performance HBM4e at up to 12Gbps," ECTC 2026

从 HBM2 到 HBM4E 的中介层与功耗扩展。来源：三星，《支持高达 12Gbps 高性能 HBM4e 的先进 SI/PI 中介层设计方案》，ECTC 2026

Samsung also presented its solution for HBM4E on interposers. HBM4E pushes data rates to 12 Gb/s and above, and doubles the number of I/O pins, increasing routing complexity. HBM4E could require 2× the interposer layers versus HBM3E and 5× as much as HBM2. Due to the increased I/O count and higher data rates, power consumption is also expected to increase by 86% versus HBM3E, and 5.6 times versus HBM2.



8-layer silicon interposer stack-up for HBM4E with two-signal / one-ground routing. Source: Samsung, “Advanced SI/PI Interposer Design Solution Enabling High-Performance HBM4e at up to 12Gbps,” ECTC 2026

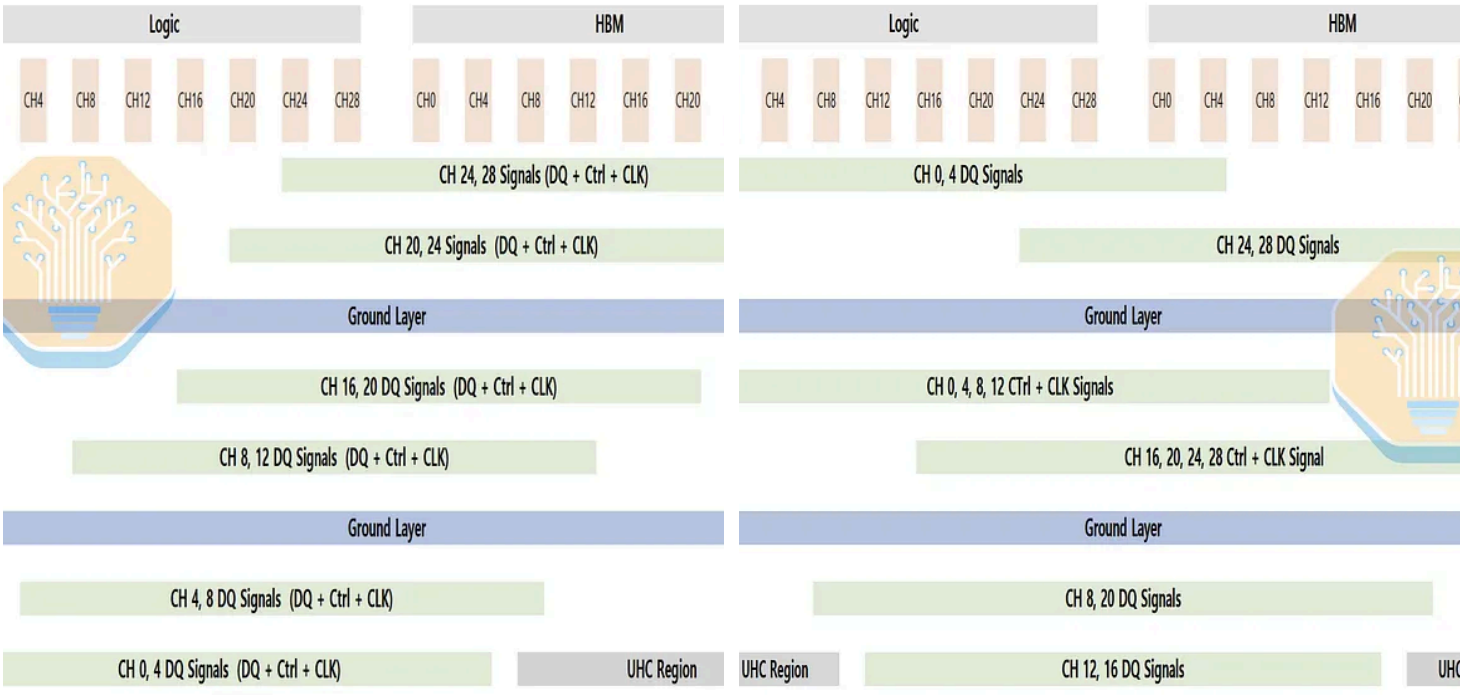
用于 HBM4E 的 8 层硅中介层叠层，采用双信号/单地线布线。来源：三星，《高性能 HBM4e 先进 SI/PI 中介层设计解决方案，速率高达 12Gbps》，ECTC 2026

Samsung proposed an 8-layer silicon interposer which it says reduces the layer count by 20% versus the estimated requirement. The interposer stack uses a repeated two-signal / one-ground staggered arrangement to shield high-speed signals, with 75% of the layers being allocated for signal routing.

三星提出了一种 8 层硅中介层，据称与预估需求相比减少了 20%的层数。该中介层堆叠采用重复的两信号/一地交错布局来屏蔽高速信号，其中 75%的层数被分配给信号布线。

Another key aspect of the interposer is the ultra-high-density capacitors (UHCs). Samsung does not spell out the exact capacitor structure, but they are likely similar to Intel EMIB-T MIM capacitors or TSMC CoWoS DTCs. The UHCs can only be placed on the M1 layer, which is also heavily used for signal routing. Thus, the available area is limited.

中介层另一个关键方面是超高密度电容器（UHC）。三星未明确说明电容器具体结构，但很可能与英特尔 EMIB-T 的 MIM 电容器或台积电 CoWoS 的 DTC 类似。由于 UHC 只能布置在同样大量用于信号布线的 M1 层上，因此可用面积十分有限。



Unbalanced versus balanced capacitor layout in 8-layer silicon interposer. Source: Samsung, “Advanced SI/PI Interposer Design Solution Enabling High-Performance HBM4e at up to 12Gbps,” ECTC 2026

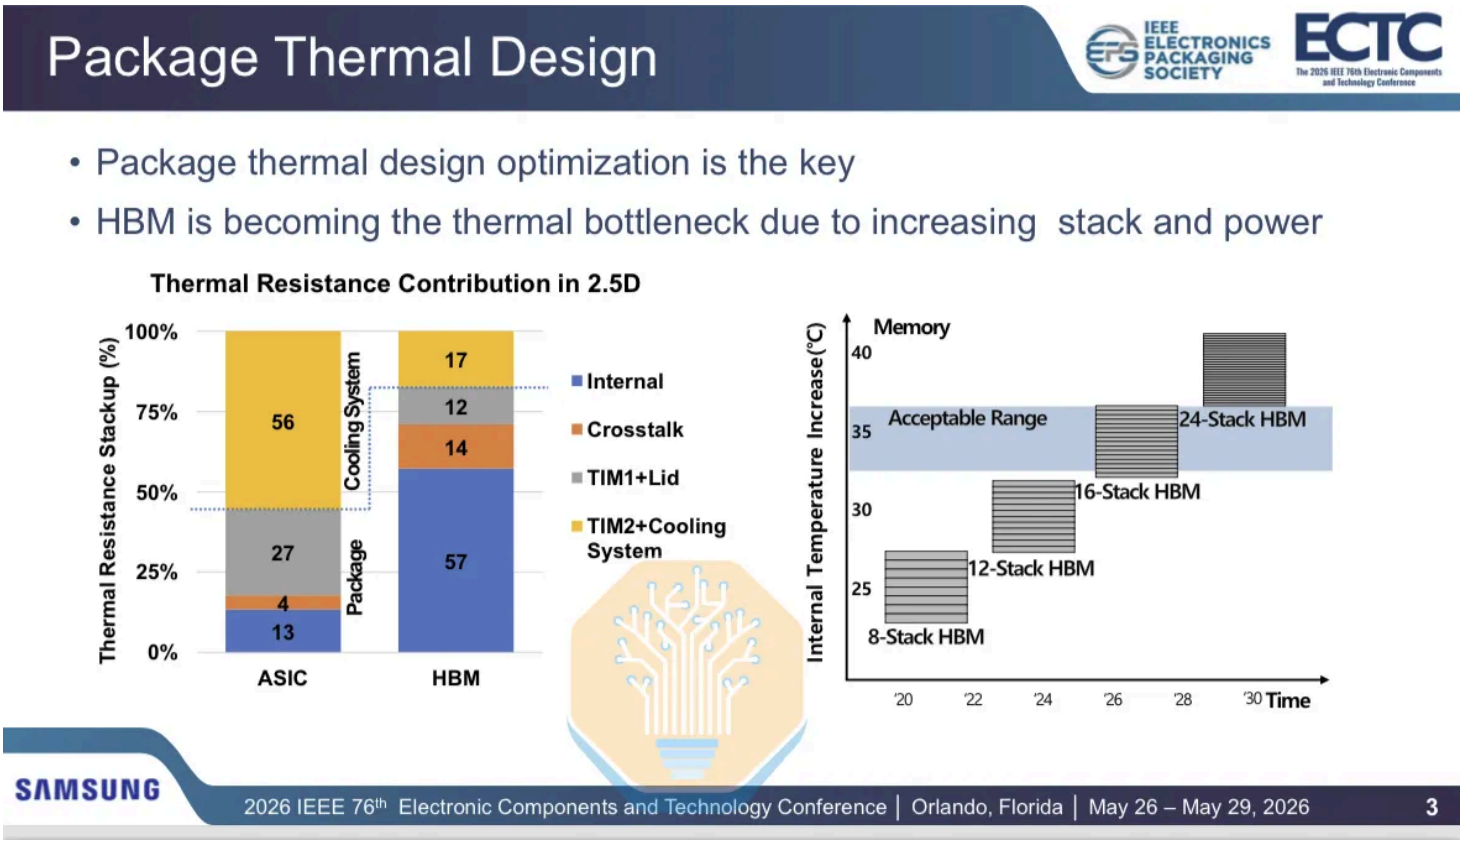
8 层硅中介层中不平衡与平衡电容布局对比。来源：三星，“支持高达 12Gbps HBM4e 的高性能 SI/PI 中介层设计解决方案”，ECTC 2026

If routing is unbalanced, the capacitors get pushed to one side of the interface, creating uneven PDN behavior between the logic and HBM sides. Samsung's layout redistributes the routing across M1 and the other layers so that UHCs can be placed more evenly across the entire interface. This reduces PDN impedance and voltage noise while keeping routing density manageable.

若布线失衡，电容会被推至接口一侧，导致逻辑端与 HBM 端之间的 PDN 性能不均。三星的布局方案通过将布线重新分配至 M1 及其他层，使超高电容（UHC）得以更均匀地分布于整个接口。此举在控制布线密度的同时，降低了 PDN 阻抗与电压噪声。

Samsung HBM Hybrid Bonding Thermals

三星 HBM 混合键合热管理

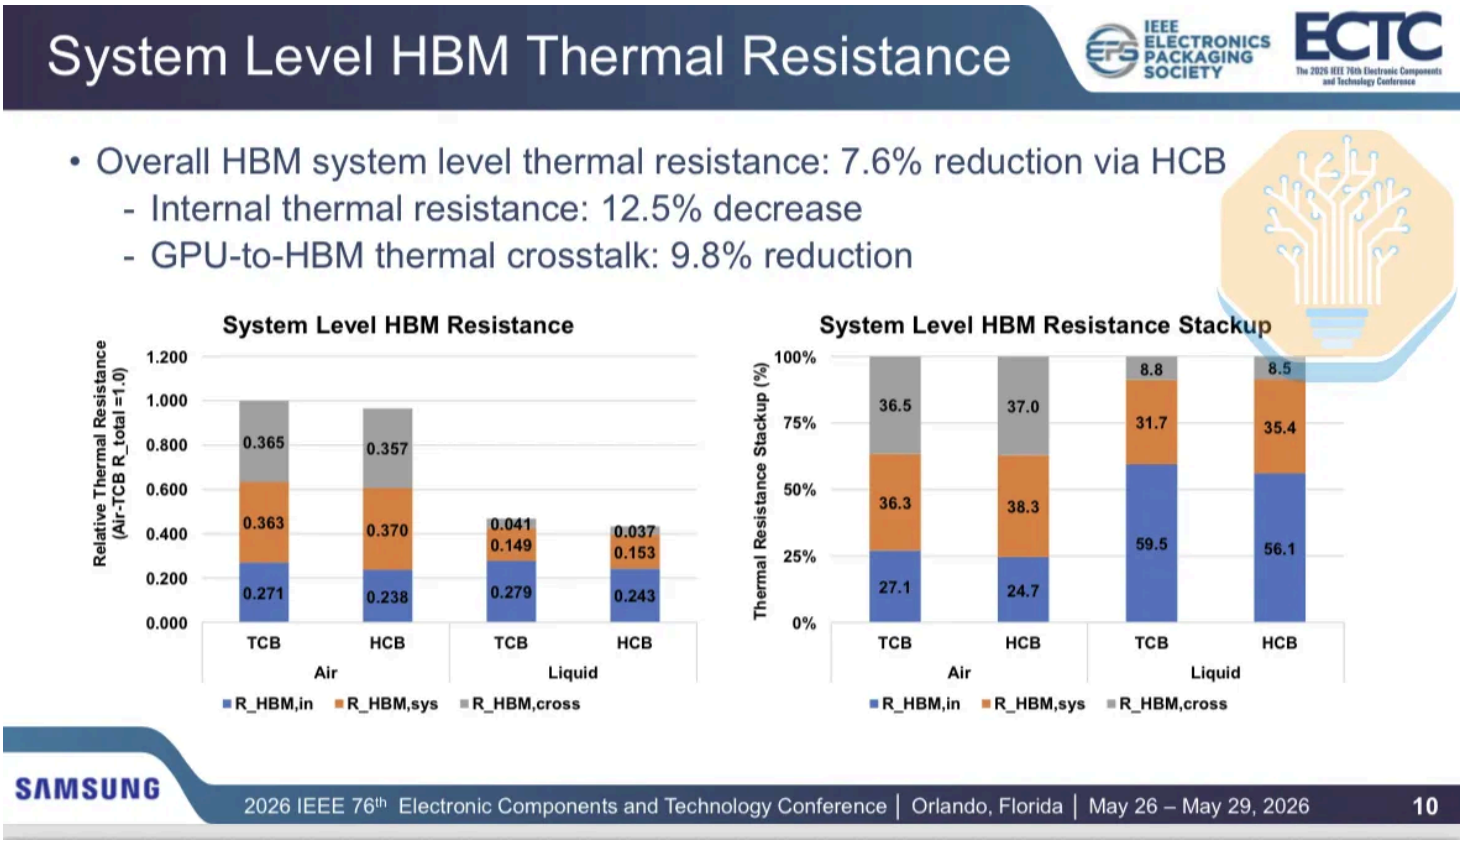


2.5D package thermal resistance contribution. Source: Samsung, “System-Level Thermal Validation of 2.5D Packages in GPU Servers: Impact of TCB vs HCB HBM Platforms,” ECTC 2026

2.5D 封装热阻贡献。来源：三星，“GPU 服务器中 2.5D 封装系统级热验证：TCB 与 HCB HBM 平台的影响”，ECTC 2026

Samsung also discussed HBM thermals, particularly with hybrid bonding. HBM stacks are getting faster and taller, while the logic die underneath is also consuming more power. With 16-hi HBM, the thermal resistance is still acceptable, but with future generations moving to 20-hi and 24-hi HBM, new approaches are needed.

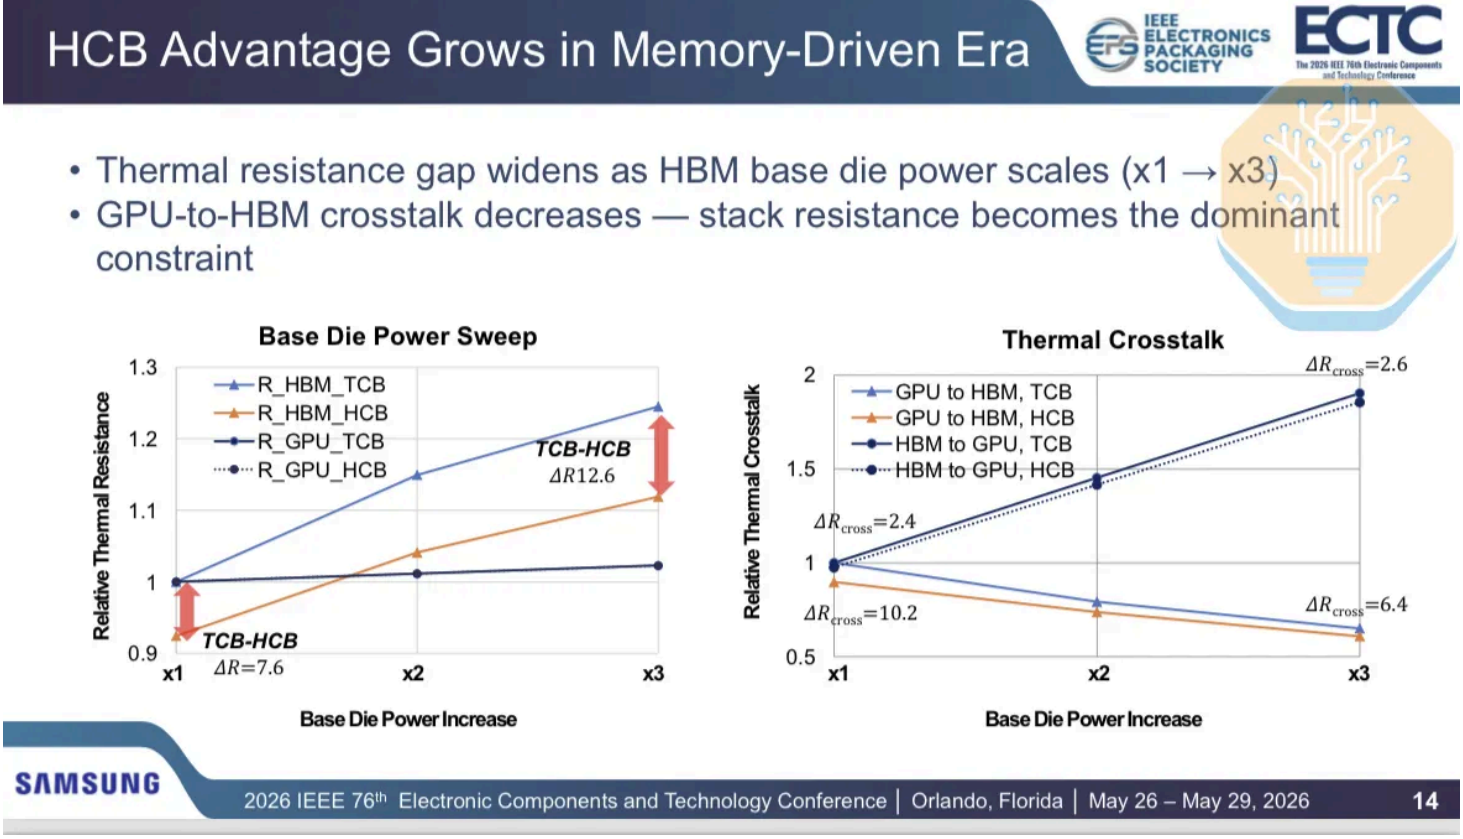
三星还讨论了 HBM 的散热问题，尤其是在混合键合方面。HBM 堆叠正变得更快更高，同时底部的逻辑芯片功耗也在增加。对于 16 层 HBM，热阻仍在可接受范围内，但随着未来几代产品升级至 20 层和 24 层 HBM，需要采用新的解决方案。



Samsung HBM thermal resistance with TCB and HCB under air and liquid cooling.
Source: Samsung, "System-Level Thermal Validation of 2.5D Packages in GPU Servers: Impact of TCB vs HCB HBM Platforms," ECTC 2026

Samsung compared thermal compression bonding (TCB) and hybrid copper bonding (HCB) for HBM on 2.5D GPU packages with 2 GPU dies and 8 HBM stacks, similar to Nvidia Blackwell. The internal HBM thermal resistance drops by 12.2% with air cooling and 12.9% with liquid cooling. Total HBM thermal resistance drops by 3.5% with air cooling and 7.7% with liquid cooling.

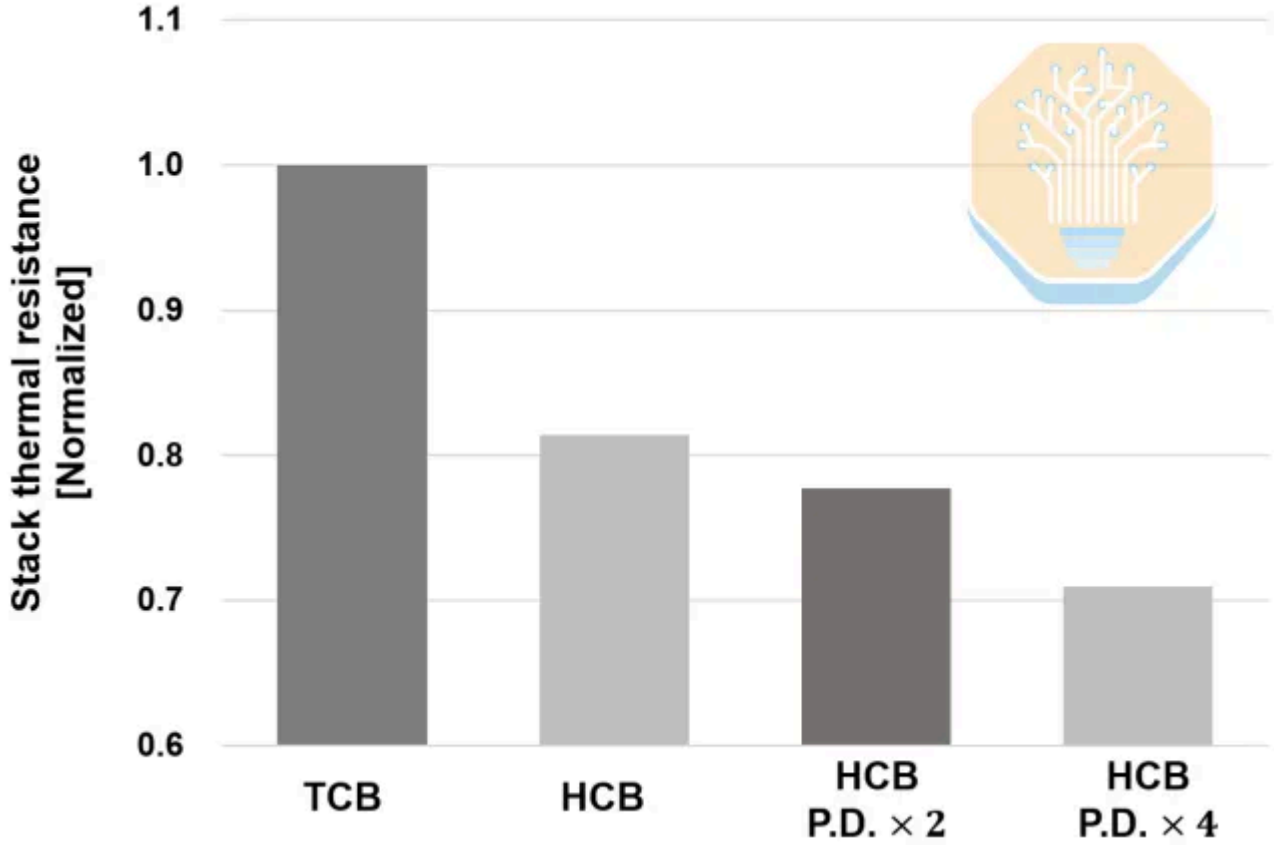
The improvement is uneven as HCB only addresses part of the thermal network. Samsung separates the path into internal resistance, system-level resistance, and GPU-to-HBM crosstalk. The internal resistance and crosstalk drop by ~12.5% and ~9.8% respectively, but the system-level resistance, including thermal interface materials and cooling, increases by ~2.3%.



HBM thermal resistance as base die power increases. Source: Samsung, "System-Level Thermal Validation of 2.5D Packages in GPU Servers: Impact of TCB vs HCB HBM Platforms," ECTC 2026

As more power moves into the HBM base die, such as in memory-bound workloads, the thermal bottleneck shifts. This is especially relevant for custom HBM, where the memory controller and more logic move into the base die. GPU-to-HBM crosstalk becomes a smaller share of the total thermal resistance, falling from 13% at 1× base-die power to 5% at 3× base-die power.

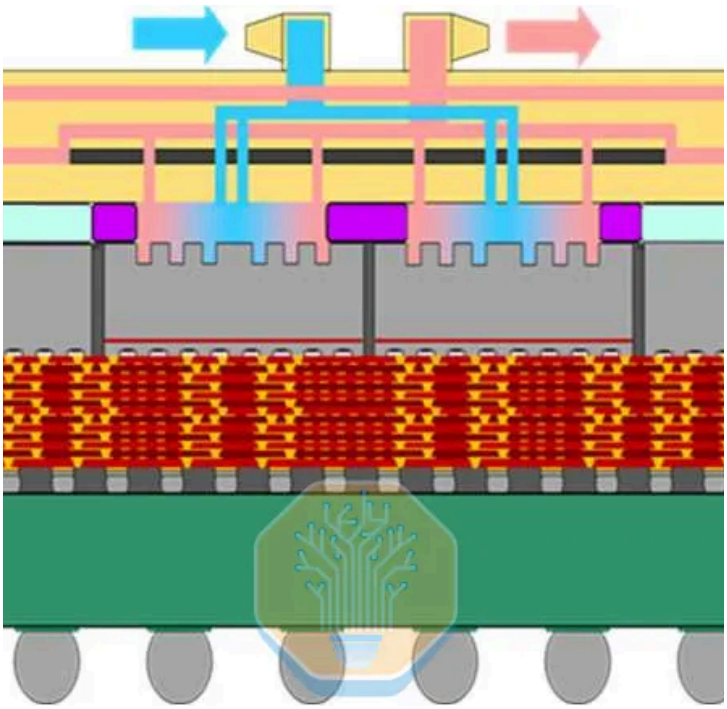
According to Samsung, moving to HCB could allow for higher inlet temperatures or higher package power. In its estimate, moving to HCB allows inlet temperatures to rise by 1-2 °C at constant package power, or package power to increase by ~4% at constant temperature. Samsung also estimates cooling power would fall by ~7%.

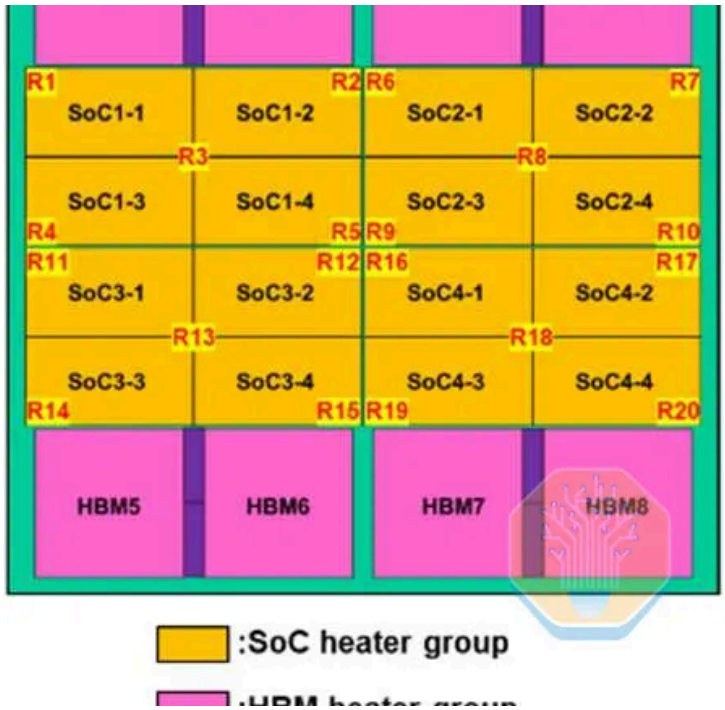


Stack-level thermal resistance with TCB and HCB with varying pad density (P.D.) in HCB. Source: Samsung, "System-Level Thermal Characterization of Hybrid Cu Bonded HBM on 2.5D Advanced Packaging," ECTC 2026

Samsung also isolated the impact of HCB at the stack level. The improvement is much larger here: Baseline HCB reduces stack thermal resistance by ~19% versus TCB. Increasing the HCB pad count raises the reduction to 22.3% at 2× pad density and 29.1% at 4× pad density.

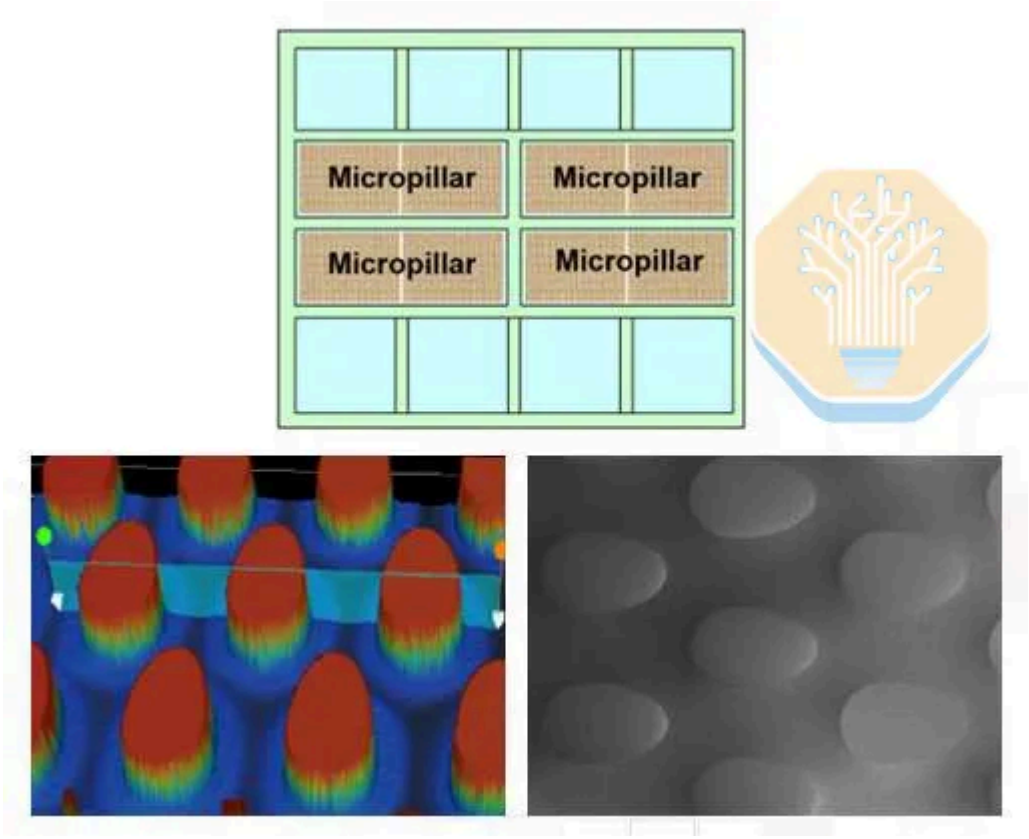
Microfluidic Cooling





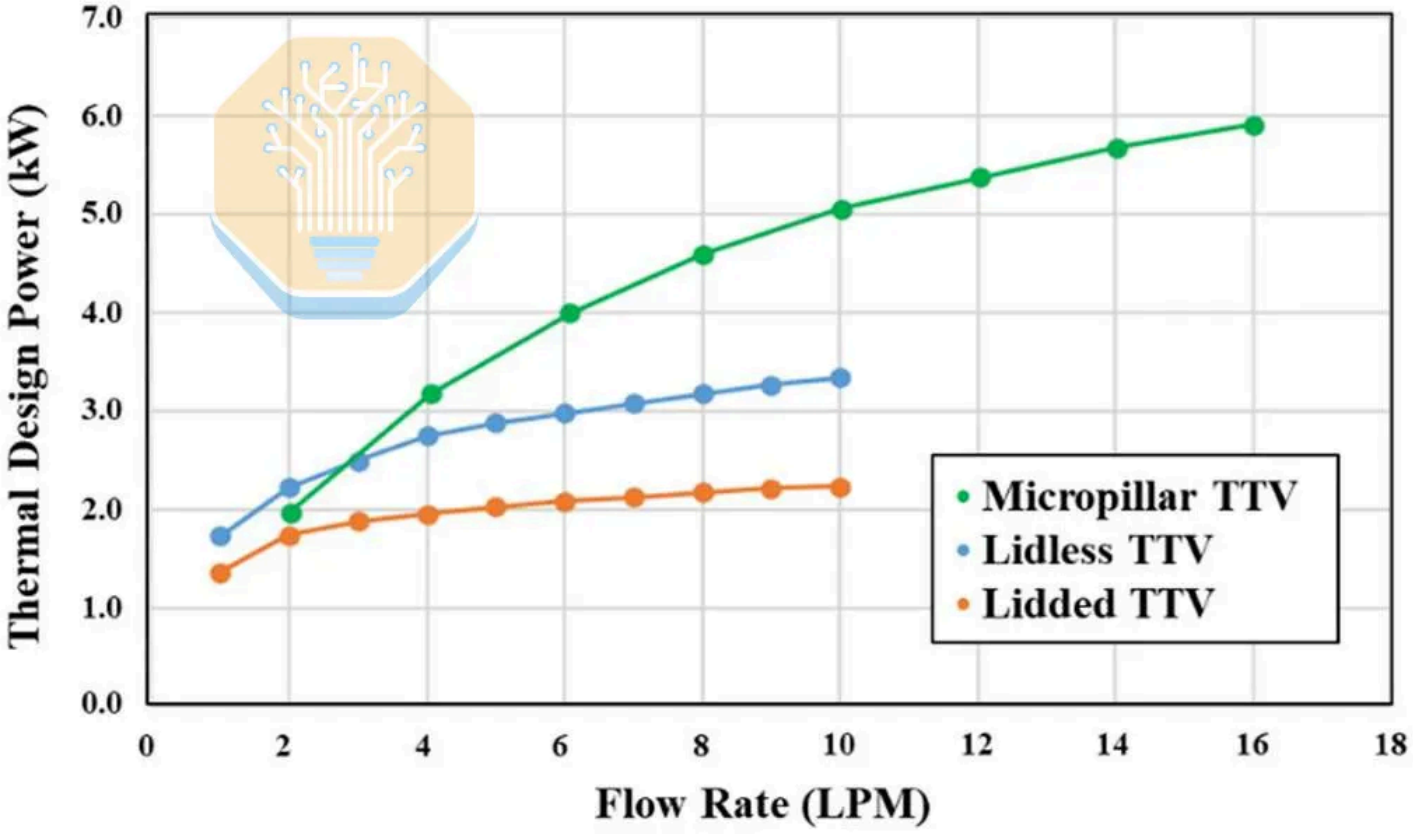
Direct-to-silicon cooling concept and CoWoS-R test vehicle. Source: TSMC, “Process Development and Thermal Characterization of Micropillar Direct-to-Silicon Liquid Cooling Solution on CoWoS-R Platform,” ECTC 2026

TSMC showed direct-to-silicon cooling on a large GPU-like test vehicle on CoWoS-R. CoWoS-R differs from CoWoS-S by using an organic rather than a silicon interposer. CoWoS-R was chosen as it provides better warpage tolerance and process compatibility. The test vehicle uses a 3.3× reticle interposer with 4 SoC dies and 8 HBM stacks. Each SoC die consists of 4 SoC heater groups, which together cover roughly half the interposer area.



Silicon micropillar formation on the backside of the SoC dies for heat dissipation. Source: TSMC, “Process Development and Thermal Characterization of Micropillar Direct-to-Silicon Liquid Cooling Solution on CoWoS-R Platform,” ECTC 2026

TSMC compared three approaches: a conventional lidded cold plate package, a lidless cold plate package, and its micropillar direct-to-silicon design. The lidded and lidless cases still use traditional cold plates and thermal interface material (TIM). For the last variant, micropillars were formed directly onto the backside of the SoC dies.



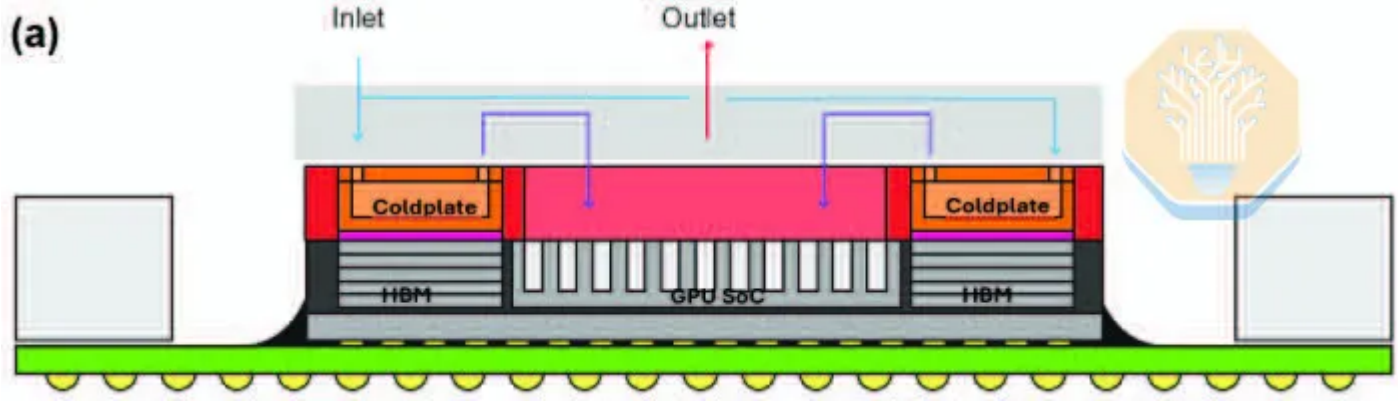
Thermal characterization of test vehicle with and without lid, and with micropillars. Source: TSMC, "Process Development and Thermal Characterization of Micropillar Direct-to-Silicon Liquid Cooling Solution on CoWoS-R Platform," ECTC 2026

With conventional cooling at 1-2 liters per minute (LPM), a lidded package dissipated 1.9-2.3 kW while the lidless package dissipated 2.5-3.0 kW, using relatively warm 40 °C DI water. Both solutions saturate beyond 4 LPM because TIM becomes the bottleneck.

The micropillar test vehicle matched the lidless cold plate result at 2 LPM, then pulled ahead at higher flow rates, dissipating 4 kW at 4 LPM and 5.3 kW at 8 LPM. Across the full test vehicle, TSMC reported uniform power dissipation above 5 kW. The micropillars brought the liquid coolant into much closer proximity to the heat source, facilitating this improvement.

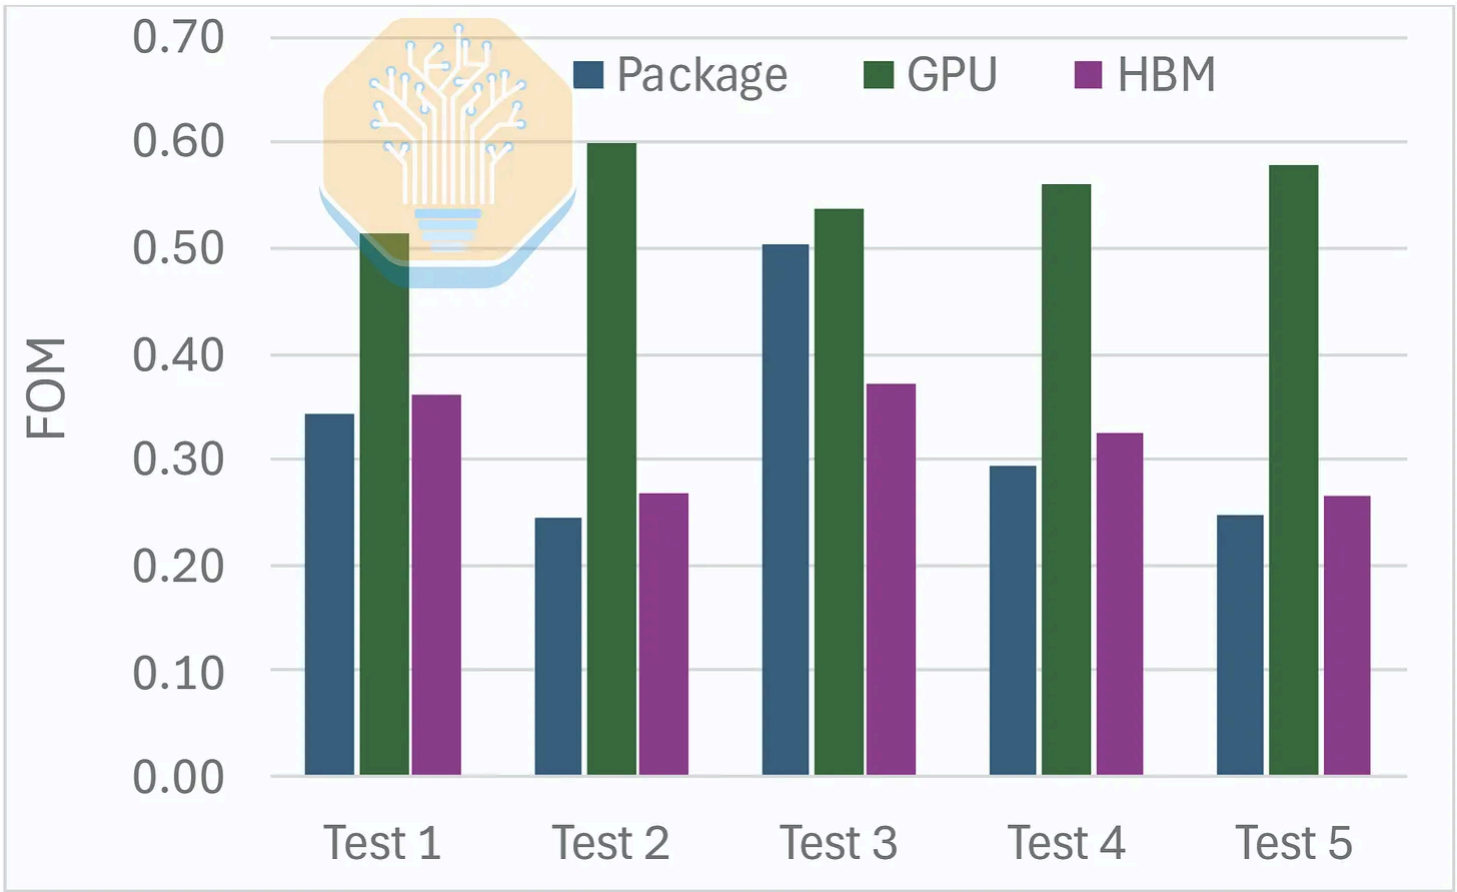
However, micropillars are not without downsides. TSMC had to form the micropillars after chip-on-wafer (CoW) processing without damaging the CoWoS-R structure and develop new sealant materials to keep the coolant contained despite package warpage and thermal expansion mismatch. The test vehicle passed moisture sensitivity level 4 (MSL4) with no helium leakage or sealant delamination.

Unfortunately, this was one of only 3 papers by TSMC at ECTC as compared with Intel's 12 and Samsung's 11.



Microsoft's direct-to-silicon microfluidic cooling assembly for Nvidia GH200. Source: Microsoft, "Direct to Silicon Microfluidic Cooling for Datacenters," ECTC 2026

Microsoft's approach differs from TSMC's in the cooling structure. While TSMC used silicon micropillars, Microsoft used straight microchannels etched into the GPU silicon. Instead of using a thermal test vehicle, Microsoft tested on an actual Nvidia GH200 GPU. This likely allowed Microsoft to capture real thermal distribution and hotspots more accurately. Microsoft tested a variety of workloads on the GPU such as HPCG and HPL, each with different compute and memory stress characteristics.



Microsoft GH200 microfluidic cooling results across GPU workloads. Source: Microsoft, "Direct to Silicon Microfluidic Cooling for Datacenters," ECTC 2026

微软 GH200 微流控冷却在不同 GPU 工作负载下的测试结果。来源：微软，《面向数据中心的直接硅基微流控冷却》，ECTC 2026

Across these workloads, Microsoft reports 51-60% lower junction-to-inlet thermal resistance for the GPU at a 1 LPM flow rate. The HBM improved less, at only 27-37%, because it was still cooled through a cold plate and TIM. Overall, this results in a 50% reduction in thermal resistance for the package.

在这些工作负载下，微软报告称，在 1 LPM 流速下，GPU 的结到入口热阻降低了 51-60%。HBM 的改善幅度较小，仅为 27-37%，因为其仍通过冷板和导热界面材料（TIM）进行冷却。总体而言，这使封装的热阻降低了 50%。

Microsoft also showed some preliminary reliability data. While thermal performance is important, datacenter deployments also require high reliability and low downtime. Over a period of 6 months, Microsoft recorded only 9 potential clogging events across ~4370 observations. The rate declined over time, suggesting early instability after installation, followed by a more stable operating period. Even after 6 months, there was no measurable silicon erosion in the microchannels. At the node level, the GH200 successfully completed 3 weeks of repeated benchmarking followed by a 1-week continuous run at a stable package power. Microsoft is still testing cluster-level mean time between failures (MTBF) and availability.

微软还展示了一些初步的可靠性数据。虽然热性能很重要，但数据中心部署同样要求高可靠性和低停机时间。在 6 个月的时间里，微软在~4370 次观测中仅记录了 9 次潜在的堵塞事件。随着时间的推移，堵塞率逐渐下降，这表明安装初期存在不稳定性，随后进入了更稳定的运行阶段。即使在 6 个月后，微通道中也没有出现可测量的硅侵蚀现象。在节点层面，GH200 成功完成了 3 周的重复基准测试，随后在稳定的封装功率下连续运行了 1 周。微软仍在测试集群级别的平均故障间隔时间（MTBF）和可用性。

SemiAnalysis is a reader-supported publication. To receive new posts and support our work, consider becoming a free or paid subscriber.

SemiAnalysis 是一份由读者支持的刊物。如需接收新文章并支持我们的工作，可考虑成为免费或付费订阅用户。

Marvell Optical Interconnects

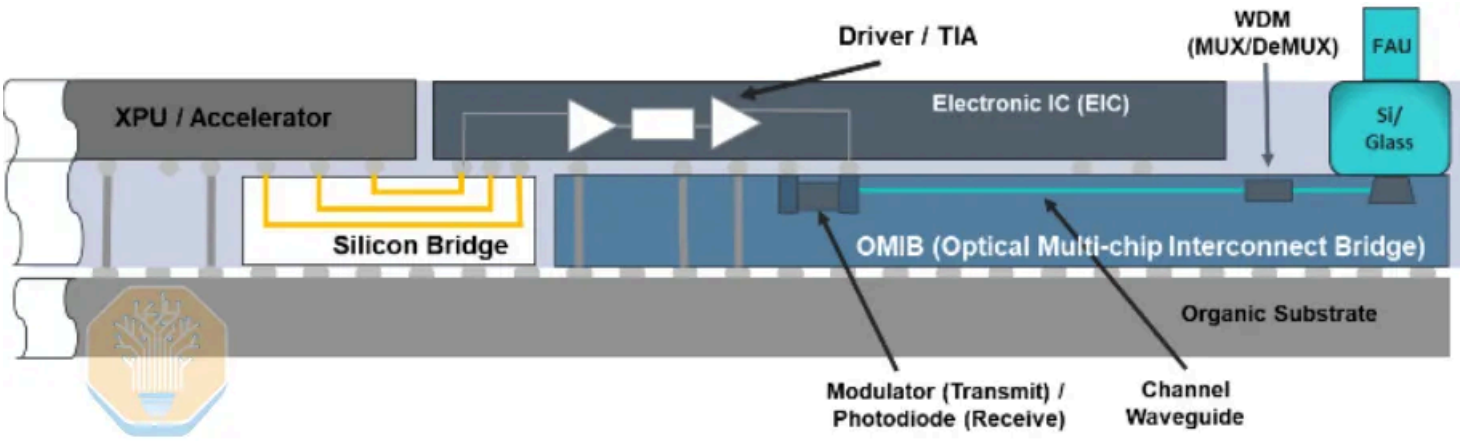
Marvell 光学互连

Another prevalent theme at ECTC was optical interconnects and co-packaged optics. Marvell had several papers in this area, primarily on its Optical Multi-Chip Interconnect Bridge (OMIB) and Photonic Fabric, both of which it acquired through its purchase of Celestial AI earlier this year.

ECTC 的另一个热门主题是光学互连和共封装光学器件。Marvell 在该领域发表了多篇论文，主要涉及其光学多芯片互连桥（OMIB）和光子结构，这两项技术均通过今年早些时候收购 Celestial AI 获得。

While photonic interposers have been discussed for a few years, Marvell showed off a more localized version at ECTC. Fabricating a multi-reticle photonic interposer will likely prove challenging for yields, especially due to reticle stitching. Marvell also noted that a photonic interposer may lack features available in conventional silicon interposers such as high-density capacitors.

尽管光子中介层已讨论多年，但 Marvell 在 ECTC 上展示了一种更为本地化的方案。制造多掩模版光子中介层在良率上将面临挑战，尤其受限于掩模版拼接技术。Marvell 同时指出，光子中介层可能缺乏传统硅中介层的高密度电容器等特性。



OMIB package concept with EIC stacked on PIC and an electrical bridge embedded in an organic interposer. Source: Marvell, “Optical Multi-Chip Interconnect Bridge (OMIB) Interposer Demonstration to Enable High-density Photonic Interconnects for High-Performance Computing Applications,” ECTC 2026

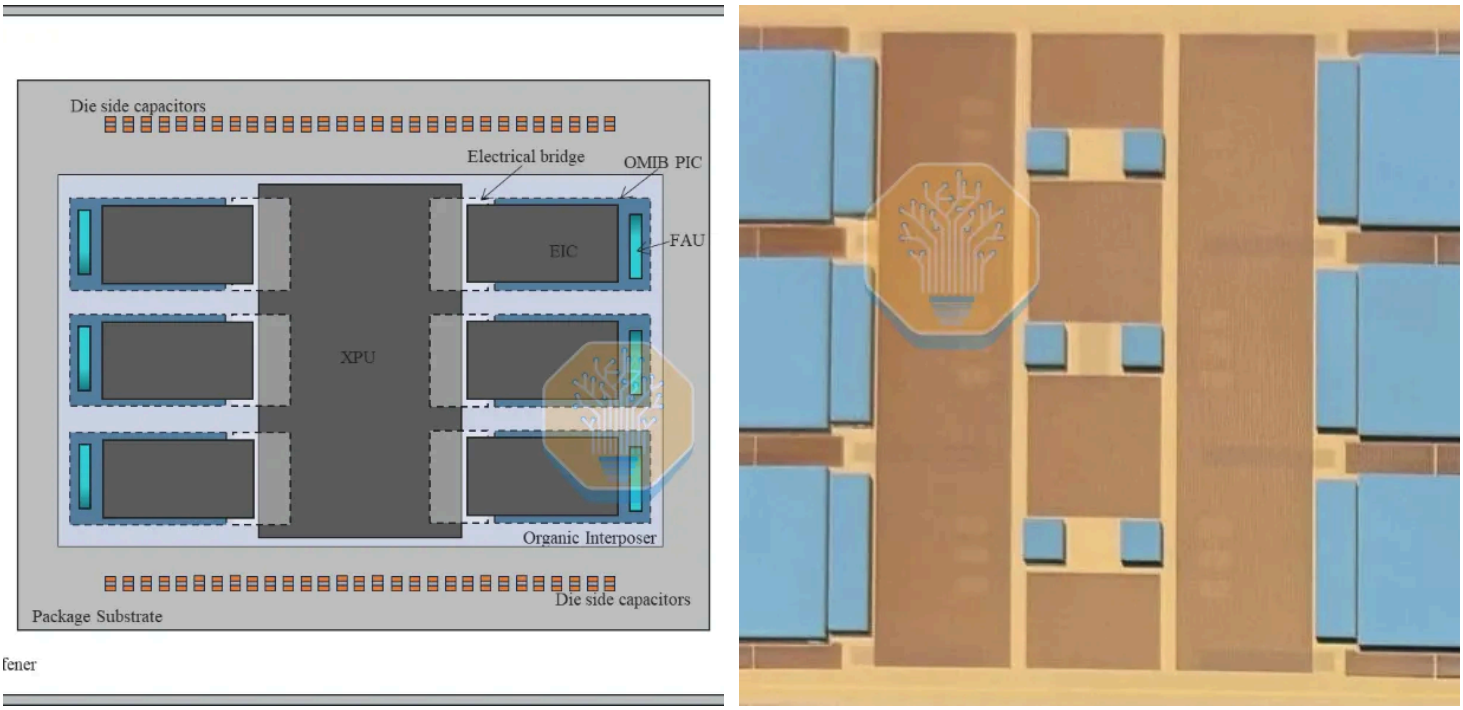
OMIB 封装概念：EIC 堆叠于 PIC 之上，电桥嵌入有机中介层。来源：Marvell，《面向高性能计算应用的高密度光子互连光学多芯片互连桥（OMIB）中介层演示》，ECTC 2026

Marvell's approach has a photonic integrated circuit (PIC) embedded in the organic RDL interposer only where it is needed. In other regions where optical interconnects are not required, an electrical bridge can be used instead.

Marvell 的方案是将光子集成电路（PIC）仅嵌入需要光互连的有机 RDL 中介层区域。在无需光互连的其他区域，则可采用电桥替代。

As the PIC is embedded in the RDL, its grating couplers would normally be obstructed after overmolding. Marvell places a silicon/glass optical block over the grating region before molding to maintain an optical path to the top, where the fiber array unit (FAU) can be attached.

由于 PIC 嵌入 RDL 后，其光栅耦合器在塑封后会受到遮挡。Marvell 通过在塑封前于光栅区域上方放置硅/玻璃光学模块，维持通向顶部的光路，以便安装光纤阵列单元（FAU）。



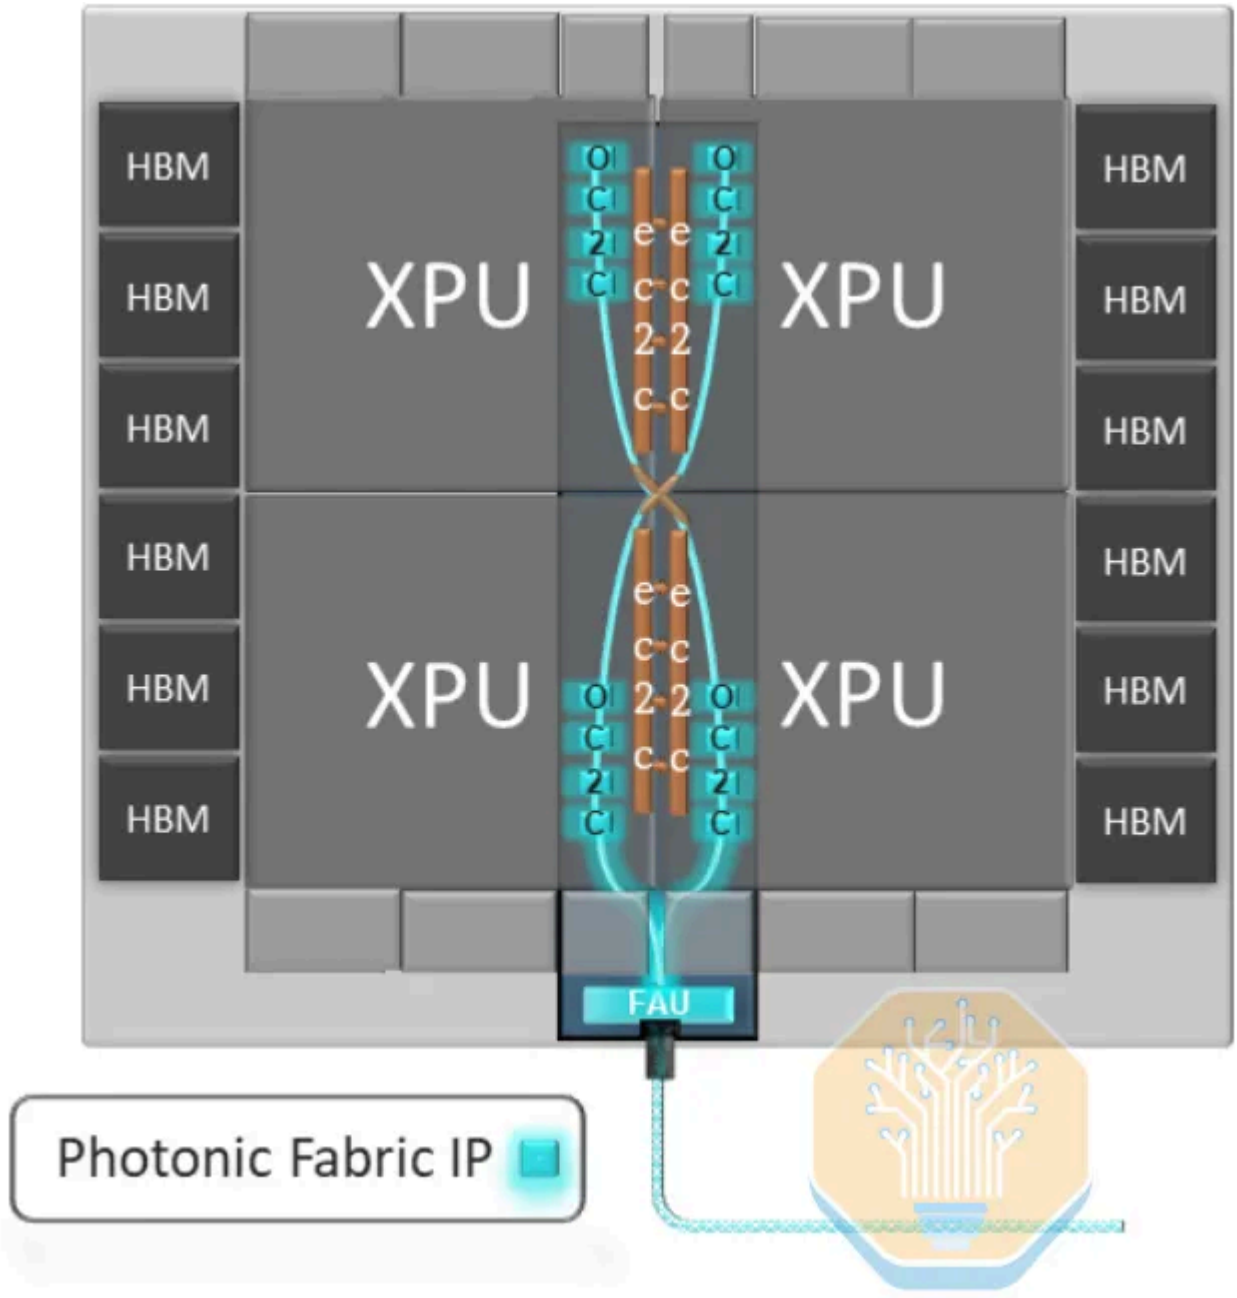
OMIB test vehicle floorplan and RDL interposer after PIC, electrical bridge, and DTC attach. Source: Marvell, "Optical Multi-Chip Interconnect Bridge (OMIB) Interposer Demonstration to Enable High-density Photonic Interconnects for High-Performance Computing Applications," ECTC 2026

OMIB 测试载具平面图及 RDL 中介层（已安装 PIC、电桥及 DTC）。来源：Marvell，《用于高性能计算应用的光学多芯片互连桥（OMIB）中介层演示》，ECTC 2026

Marvell's OMIB test vehicle has one primary XPU die and six EIC dies on top. Embedded in the interposer are six PICs (big blue rectangles), six electrical bridges (thin blue rectangles) and 12 DTC dies (small blue squares). The roughly 2×-reticle RDL

interposer uses four layers at 2/2 µm L/S.

Marvell 的 OMIB 测试载具顶部包含一颗主 XPU 芯片和六颗 EIC 芯片。嵌入中介层的有六个 PIC（蓝色大矩形）、六颗电桥（蓝色细矩形）及 12 颗 DTC 芯片（蓝色小正方形）。该约 2 倍掩模版的 RDL 中介层采用 4 层布线，线宽/线距为 2/2 微米。



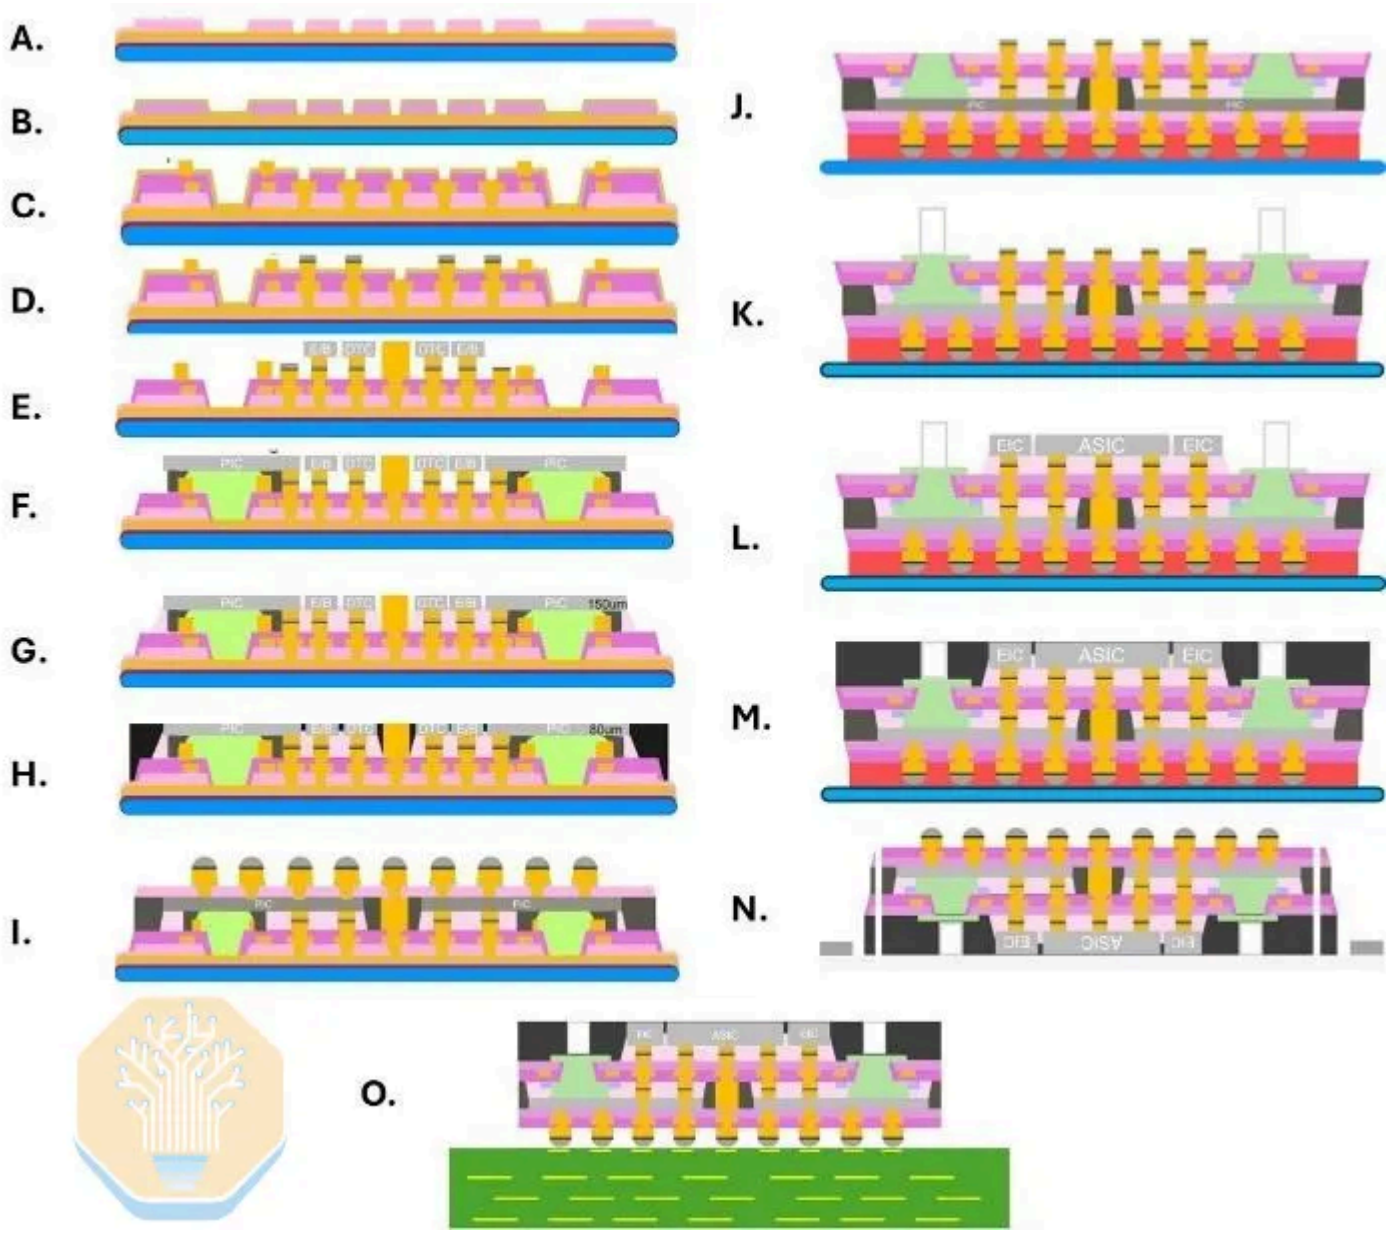
Concept of OMIB with XPU-XPU interconnect and external FAU. Source: Marvell, "Optical Multi-Chip Interconnect Bridge (OMIB) Interposer Demonstration to Enable High-density Photonic Interconnects for High-Performance Computing Applications," ECTC 2026

OMIB 概念示意图：XPU 间互连及外部 FAU 连接。来源：Marvell，《用于高性能计算应用的光学多芯片互连桥（OMIB）中介层演示》，ECTC 2026

Marvell also showed off a conceptual multi-die XPU with optical chip-to-chip interconnects to reduce latency and hop count. Marvell claims that OMIB removes the shoreline limitations, since the same bridge can route both on-package die-to-die

links and external optical interconnects. Marvell cites a bandwidth density of 1.8 Tbps/mm^2 with this approach.

Marvell 还展示了一种采用光学芯片间互连的多芯片 XPU 概念方案，旨在降低延迟和跳数。Marvell 声称 OMIB 消除了海岸线限制，因为同一桥接器既可路由封装内芯片间链路，也可路由外部光学互连。据 Marvell 数据，该方案可实现 1.8 Tbps/mm^{2}的带宽密度。



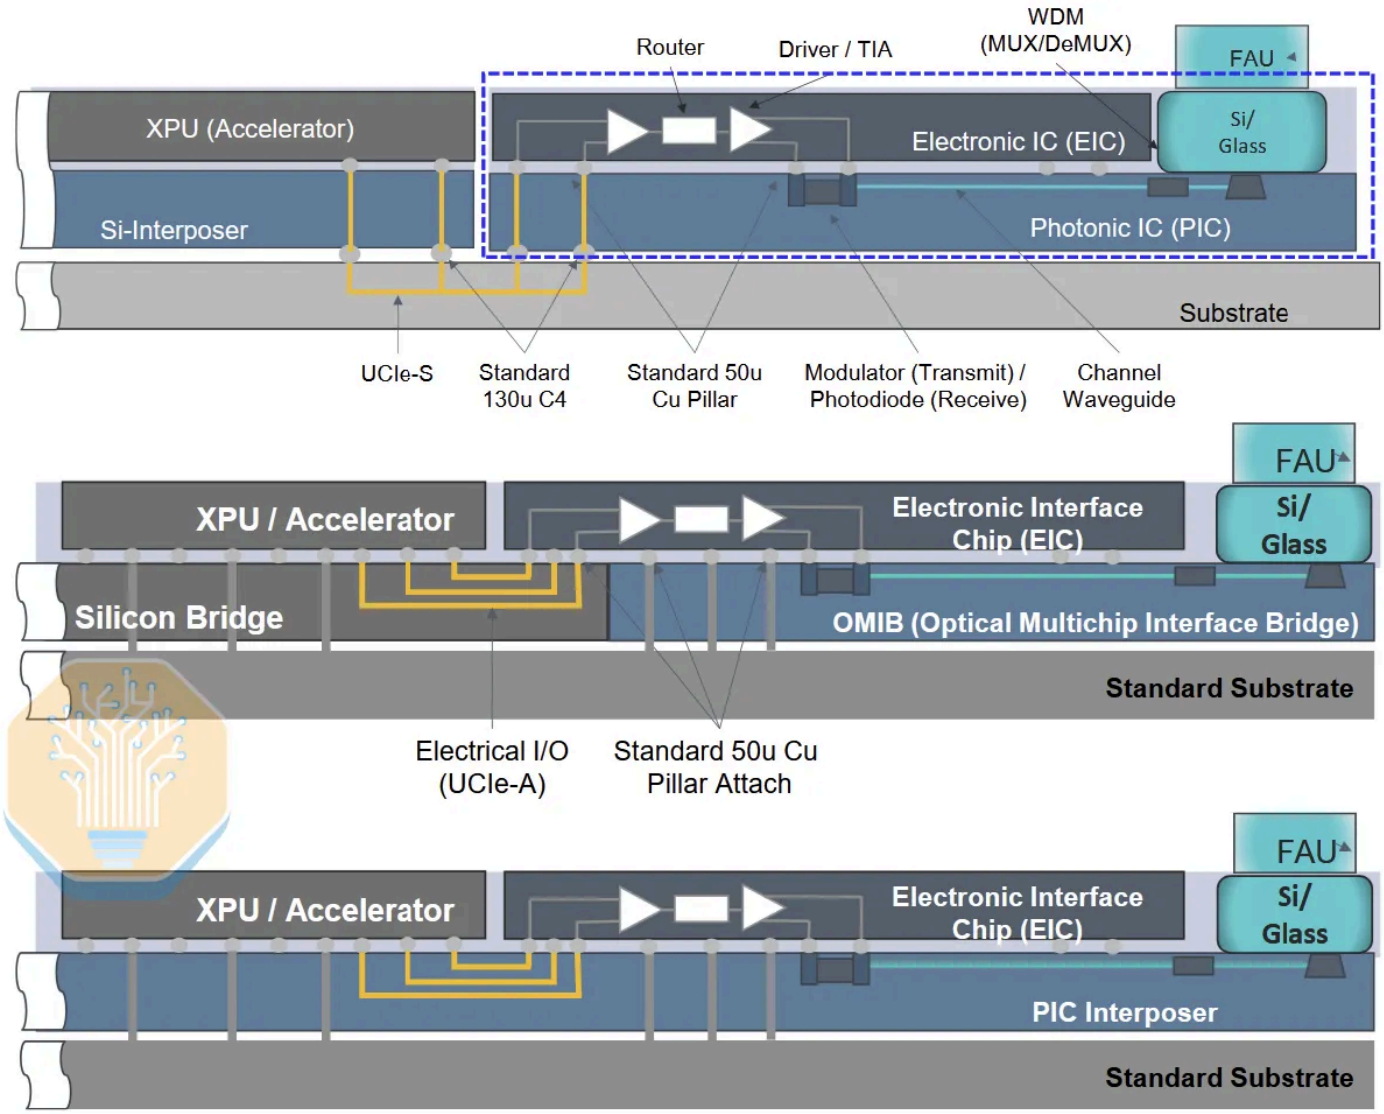
Process flow of 2.5D OMIB integration. Source: Marvell, “Optical Multi-Chip Interconnect Bridge (OMIB) Interposer Demonstration to Enable High-density Photonic Interconnects for High-Performance Computing Applications,” ECTC 2026

2.5D OMIB 集成的工艺流程。来源：Marvell，《面向高性能计算应用的光学多芯片互连桥接 (OMIB) 中介层演示，以实现高密度光子互连》，ECTC 2026

The process flow shown by Marvell is chip-last, similar to TSMC’s CoWoS-L. Marvell builds an organic RDL interposer with embedded bridges, OMIB PICs, DTCs, and

other components, with C4 bumps connecting it to the package substrate. Bridge TSVs and tall copper pillars connect to the RDL, and the ASIC dies and EICs are attached last.

马维尔展示的工艺流程属于后置芯片（chip-last）方案，与台积电的 CoWoS-L 技术类似。马维尔构建了一个有机 RDL 中介层，其中嵌入了桥接层、OMIB 光子集成电路、DTC 及其他组件，并通过 C4 凸点与封装基板相连。桥接硅通孔（TSV）和高铜柱连接至 RDL，ASIC 芯片和电子集成电路（EIC）则最后完成贴装。



Package concepts with optical engine on substrate, OMIB and photonic interposer. Source: Marvell, "Photonic Fabric™ Interconnect for a Scale-up Network Solution in Accelerated Computing," ECTC 2026

基板上的光学引擎、OMIB 及光子中介层封装方案示意图。来源：马维尔，《面向加速计算集群网络解决方案的 Photonic Fabric™互连技术》，ECTC 2026

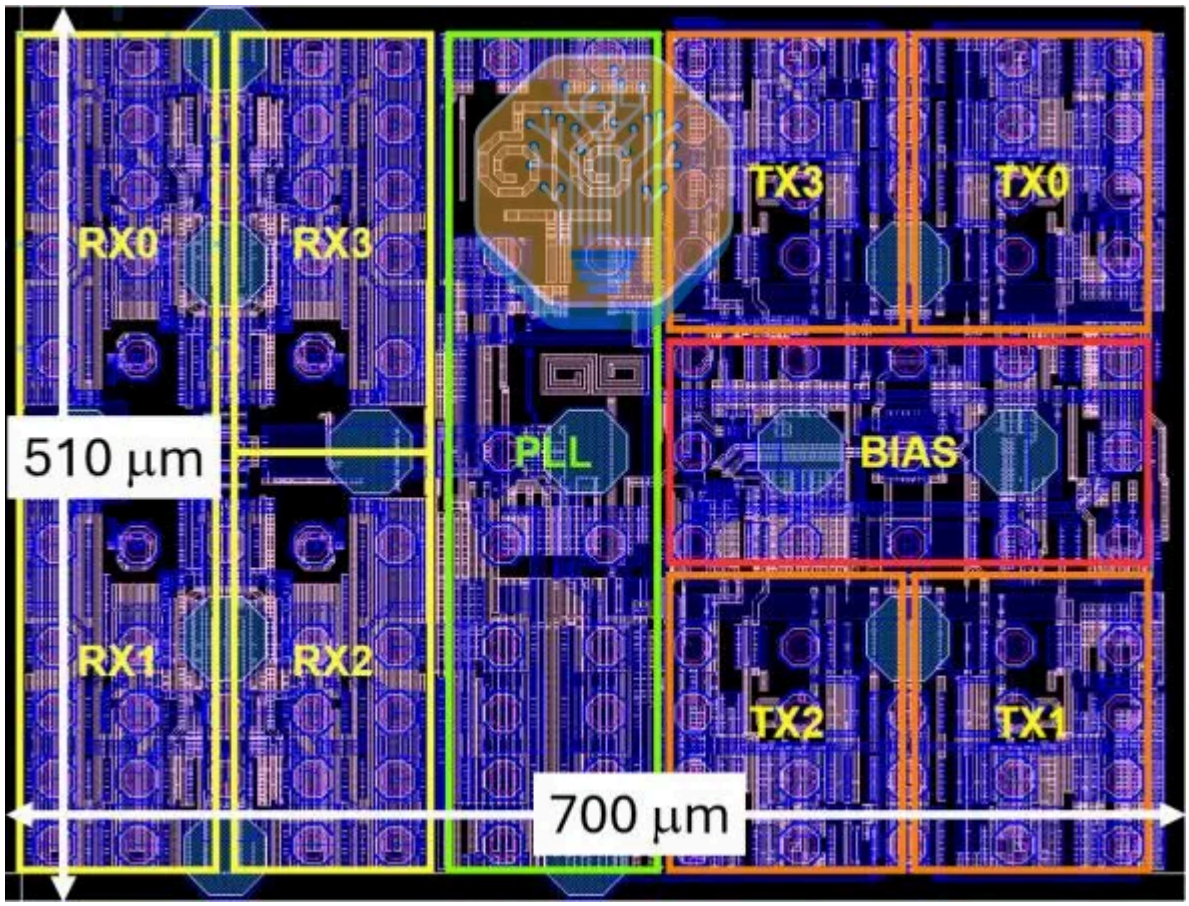
In the near term, vertically stacked optical engines like TSMC's COUPE are more achievable than OMIB-style connections or a full photonic interposer. Marvell connects the EIC and PIC using microbumps at 50 µm pitch, then mounts the

resulting engine to either the package substrate or an interposer.

在短期内，像台积电 COUPE 这样的垂直堆叠光学引擎，比 OMIB 式连接或全光子中介层更具可行性。马维尔采用 50 微米间距的微凸点连接 EIC 和 PIC，然后将组装好的引擎安装至封装基板或中介层上。

The substrate configuration can use a UCIe-S-like parallel bus at a coarse 130 µm C4 pitch, while the interposer configuration can use a UCIe-A interface at a tighter ~40-45 µm pitch. Marvell favors the substrate approach due to its simplicity and better thermal isolation.

基板配置可采用间距为 130 微米的粗间距 C4 凸点的 UCIe-S 并行总线，而中介层配置则可使用间距更紧凑（~40-45 微米）的 UCIe-A 接口。马维尔更倾向于基板方案，因其结构简单且易于实现。



Photonic Fabric optical link test chip EIC with four TX-RX pairs. Source: Marvell, "Photonic Fabric™ Interconnect for a Scale-up Network Solution in Accelerated Computing," ECTC 2026

光子互连光学链路测试芯片 EIC，包含四对 TX-RX 通道。来源：Marvell，《面向加速计算扩展网络解决方案的光子互连架构》，ECTC 2026

Marvell tested an optical engine using a 5 nm (likely TSMC N5) EIC, with four 56 Gb/s TX-RX pairs for 224 Gb/s in each direction. The design uses electro-absorption modulators (EAMs) instead of the micro-ring modulators (MRMs) preferred by other companies, citing better thermal stability and a wider operating wavelength range. While these advantages are real, we believe that EAMs will prove difficult to manufacture at scale.

Marvell 使用 5 纳米（推测为台积电 N5 工艺）EIC 测试了一款光学引擎，集成四对 56 Gb/s TX-RX 通道，实现每方向 224 Gb/s 传输速率。该设计采用电吸收调制器（EAM）替代其他企业偏好的微环调制器（MRM），其优势在于更优的热稳定性和更宽的工作波长范围。尽管这些优势真实存在，但我们认为 EAM 在规模化制造中仍将面临挑战。



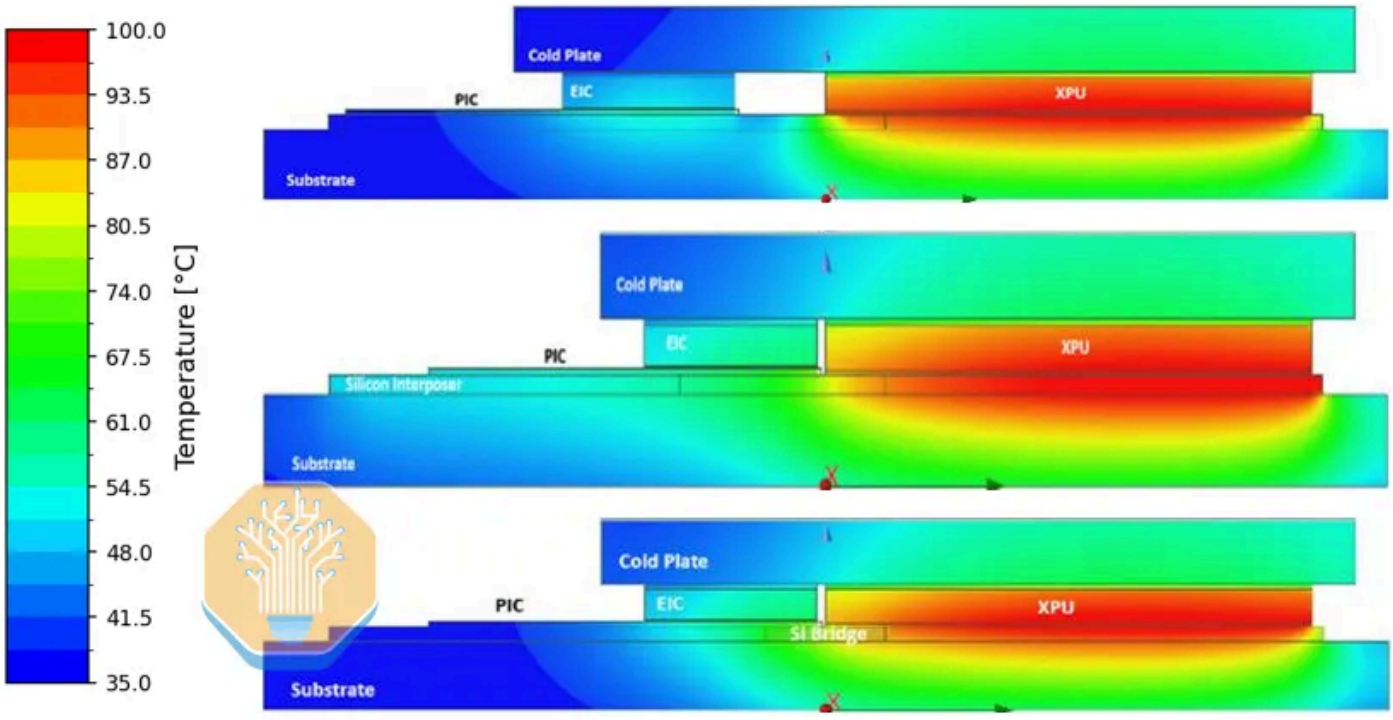
Co-Packaged Optics (CPO) Book – Scaling with Light for the Next Wave of Interconnect

共封装光学（CPO）手册——以光为媒介实现互连新突破

DYLAN PATEL, DANIEL NISHBALL, AND 4 OTHERS

作者：DYLAN PATEL、DANIEL NISHBALL 及其他四位合作者 · 1月2日 1 月 2 日

Read full story 阅读完整报道 →



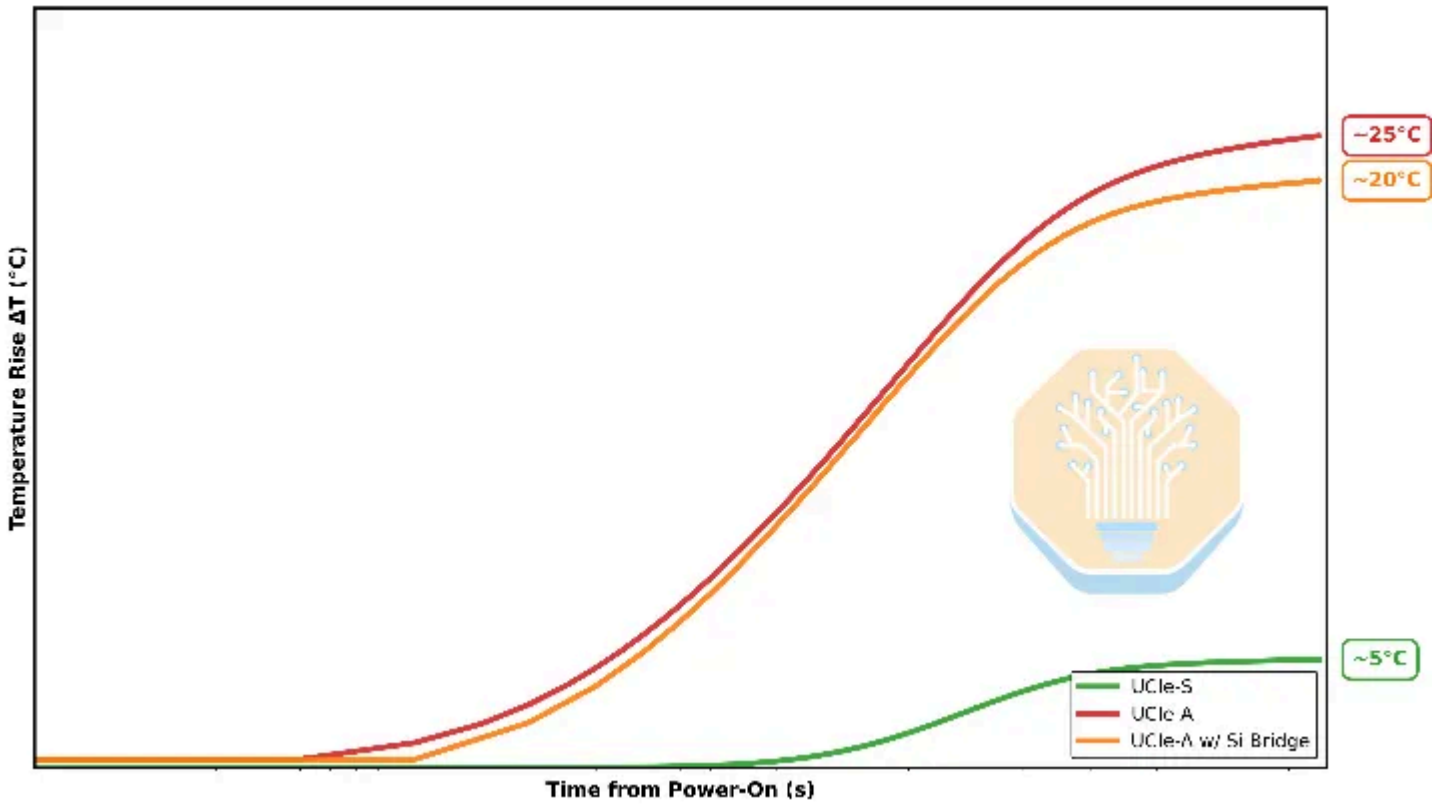
PIC steady-state temperature gradient under full XPU power on substrate, silicon interposer and silicon bridge. Source: Marvell, "Photonic Fabric™ Chiplets for Co-

Packaged Optics in AI Data Centers," ECTC 2026

在 XPU 满功率运行条件下，PIC（光子集成电路）在基板、硅中介层和硅桥上的稳态温度梯度。来源：Marvell，《Photonic Fabric™ 芯粒用于 AI 数据中心共封装光学》，ECTC 2026

Marvell also compared the thermal characteristics of an optical engine connected via UCIe-S on substrate and UCIe-A both on a silicon interposer and over a silicon bridge. Under full XPU load, the PIC temperature rose by less than 5 °C on the substrate, versus ~25 °C on the interposer and ~20 °C with the bridge. The organic substrate's low thermal conductivity and relatively large millimeter-scale air gap isolate the PIC. In both UCIe-A configurations, the fine-pitch silicon close to the XPU provides a low-resistance thermal path. There is also a much smaller air gap, at only ~100 µm, which contributes to the heating of the PIC.

Marvell 还比较了通过 UCIe-S 连接在基板上的光引擎，以及通过 UCIe-A 连接在硅中介层和硅桥上的光引擎的热特性。在 XPU 满载时，PIC 在基板上的温度上升不到 5°C，而在中介层上为 ~25 °C，在硅桥上为 ~20 °C。有机基板的热导率低且空气间隙较大（毫米级），从而隔离了 PIC。在两种 UCIe-A 配置中，靠近 XPU 的细间距硅结构提供了低热阻路径。此外，空气间隙也小得多，仅为 ~100 µm，这导致了 PIC 的加热。



PIC transient temperature rise on substrate, silicon interposer and silicon bridge. Source: Marvell, "Photonic Fabric™ Chiplets for Co-Packaged Optics in AI Data

Centers," ECTC 2026

PIC 在基板、硅中介层和硅桥上的瞬态温升。来源：Marvell，《Photonic Fabric™ 芯粒用于 AI 数据中心共封装光学》，ECTC 2026

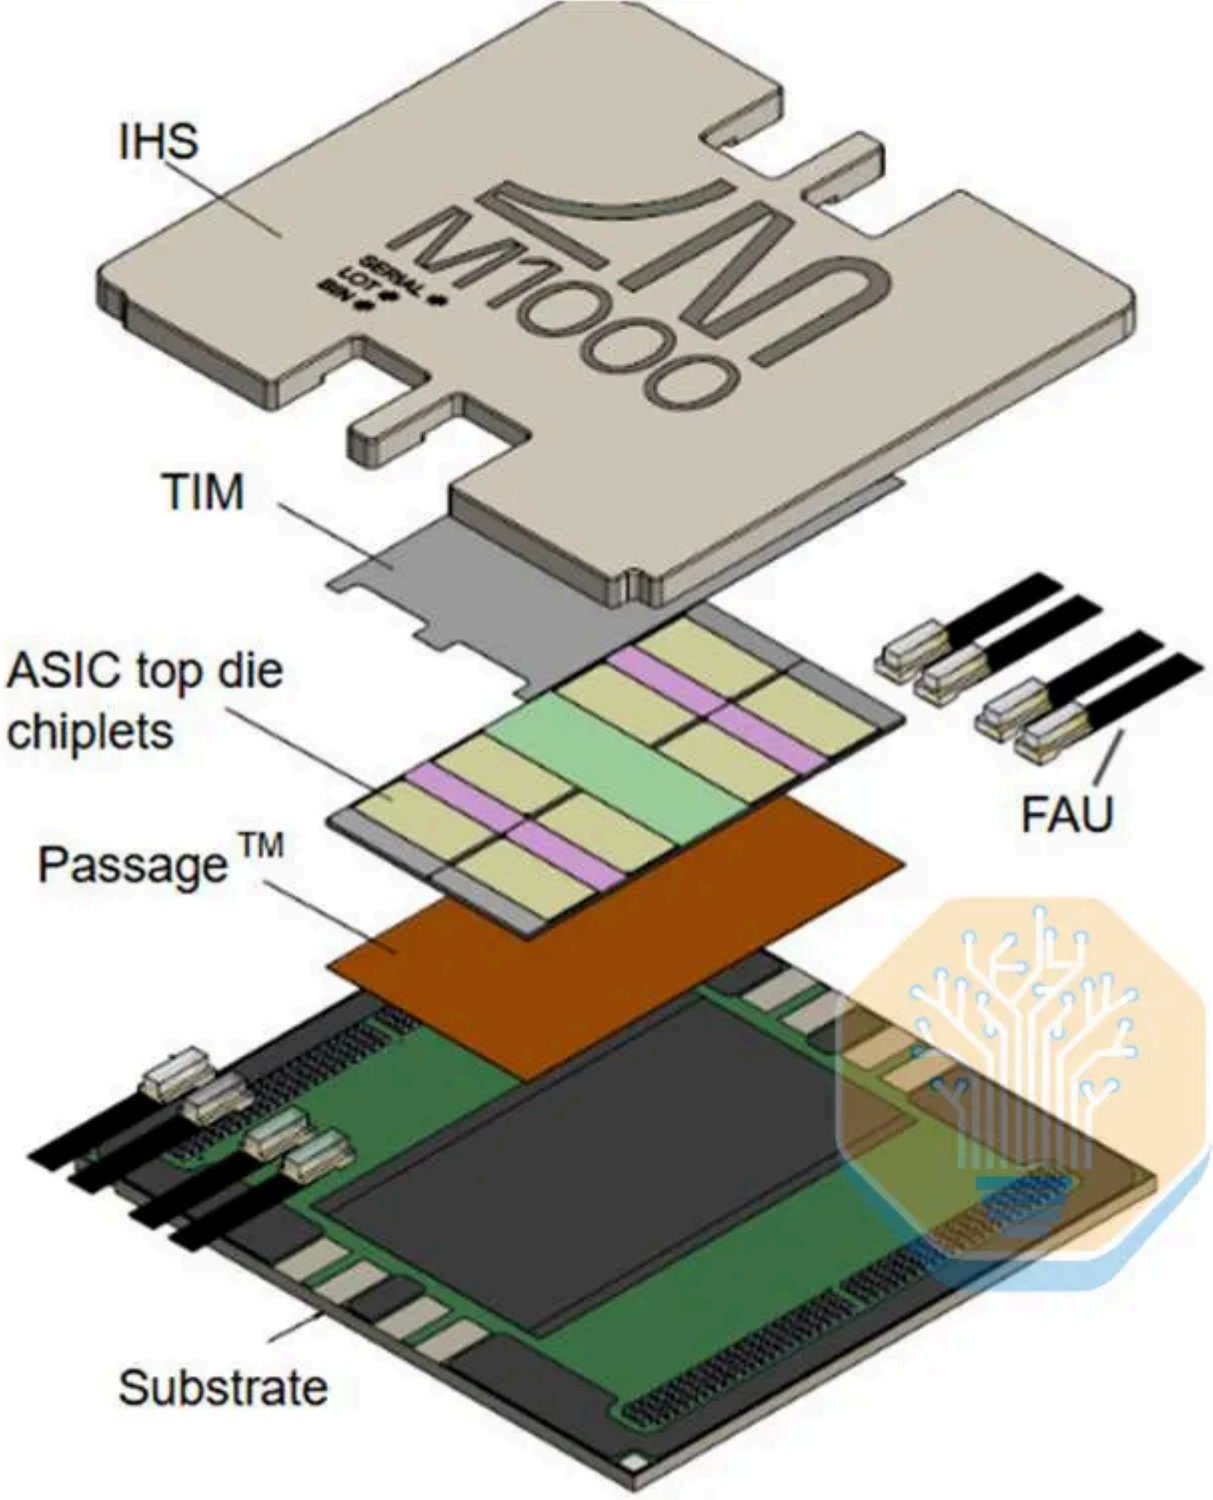
The thermal transients occur within ~30 ms of an XPU power-state change. The PIC heats at ~10 °C/s on the organic substrate, versus ~100 °C/s with the bridge and ~120 °C/s on the interposer. Marvell argues that the EAM bias voltage can be adjusted electronically fast enough to track these changes, while ring modulators require heater-and-feedback loops, which are constrained by slower time constants.

热瞬态发生在 XPU 功率状态变化的~30 毫秒内。PIC 在有机基板上的加热速率为~10°C/秒，而在桥接片上为~100°C/秒，在硅中介层上为~120°C/秒。Marvell 认为，EAM 偏置电压可通过电子方式快速调整以跟踪这些变化，而环形调制器则需要加热器与反馈回路，受限于较慢的时间常数。

Lightmatter Passage M1000

Lightmatter has previously detailed the Passage M1000's architecture and optical interconnects, and has outlined its manufacturing approach. At ECTC, it provided a much deeper look at the assembly process, fiber attachment, and packaging results for integrating the multi-reticle photonic interposer with ASIC chiplets.

Lightmatter 此前已详细介绍了 Passage M1000 的架构和光学互连方案，并概述了其制造方法。在 ECTC 上，该公司对多光罩光子中介层与 ASIC 芯粒集成的组装工艺、光纤贴装及封装成果进行了更深入的展示。



Passage M1000 schematic. Source: Lightmatter, “Advancing Interconnect Performance and Reliability with Innovations in 3D Photonic Integration Packaging and Fiber Coupling,” ECTC 2026

Passage M1000 示意图。来源：Lightmatter，《通过 3D 光子集成封装与光纤耦合创新提升互连性能与可靠性》，ECTC 2026

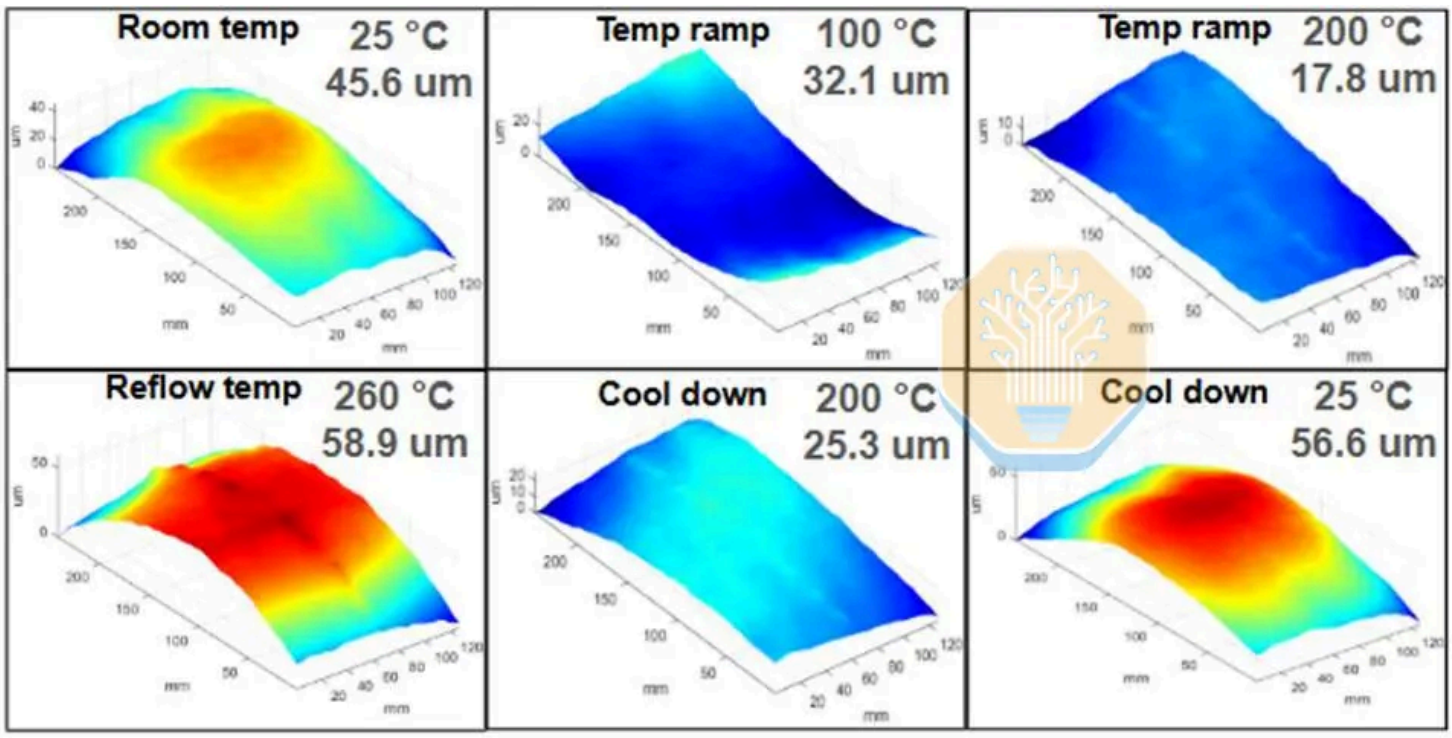
The test vehicle uses chip-on-wafer assembly to attach 15 ASIC chiplets to a four-tile M1000 interposer. We estimate the interposer measures ~2100 mm^2, about half of the 4000 mm^2 eight-tile configuration shown at Hot Chips 2025. The M1000 is highly

configurable to different sizes and tile counts depending on the needs of the product.

测试载具采用芯片-晶圆组装技术，将 15 个 ASIC 芯粒贴装到四片式 M1000 中介层上。我们估算该中介层尺寸为~2100 mm^2，约为 Hot Chips 2025 上展示的 4000 mm^2 八片式配置的一半。M1000 可根据产品需求灵活配置不同尺寸和片数。

Each of the four tiles has 32 optical waveguides at 127 µm pitch. Electrical signals and power travel from the substrate to the ASIC chips through the C4 bumps at ~176 µm pitch, two backside RDL layers, and 126 µm deep TSVs with a 10 µm width before reaching the ASIC chiplets through microbumps. The ~2100 mm^2 interposer occupies less than one-third of the 7200 mm^2 organic substrate. This ratio is closer to the logic area to package size ratio than the package-spanning interposer area of an equivalent silicon or organic RDL solution. It is unclear whether this is a consequence of the optical architecture or a design choice specific to this test vehicle.

每个中介层包含 32 根光学波导，间距为 127 微米。电信号与电源通过间距~176 微米的 C4 凸点、两层背面重分布层、深度 126 微米（宽度 10 微米）的硅通孔从基板传输至 ASIC 芯片，再经由微凸点到达 ASIC 芯粒。~2100 mm^2的中介层仅占 7200 mm^2有机基板的三分之一以下。这一比例更接近逻辑面积与封装尺寸的比值，而非同等硅基或有机 RDL 方案中横跨封装的中介层面积。目前尚不清楚这是光学架构的必然结果，还是该测试载具特有的设计选择。



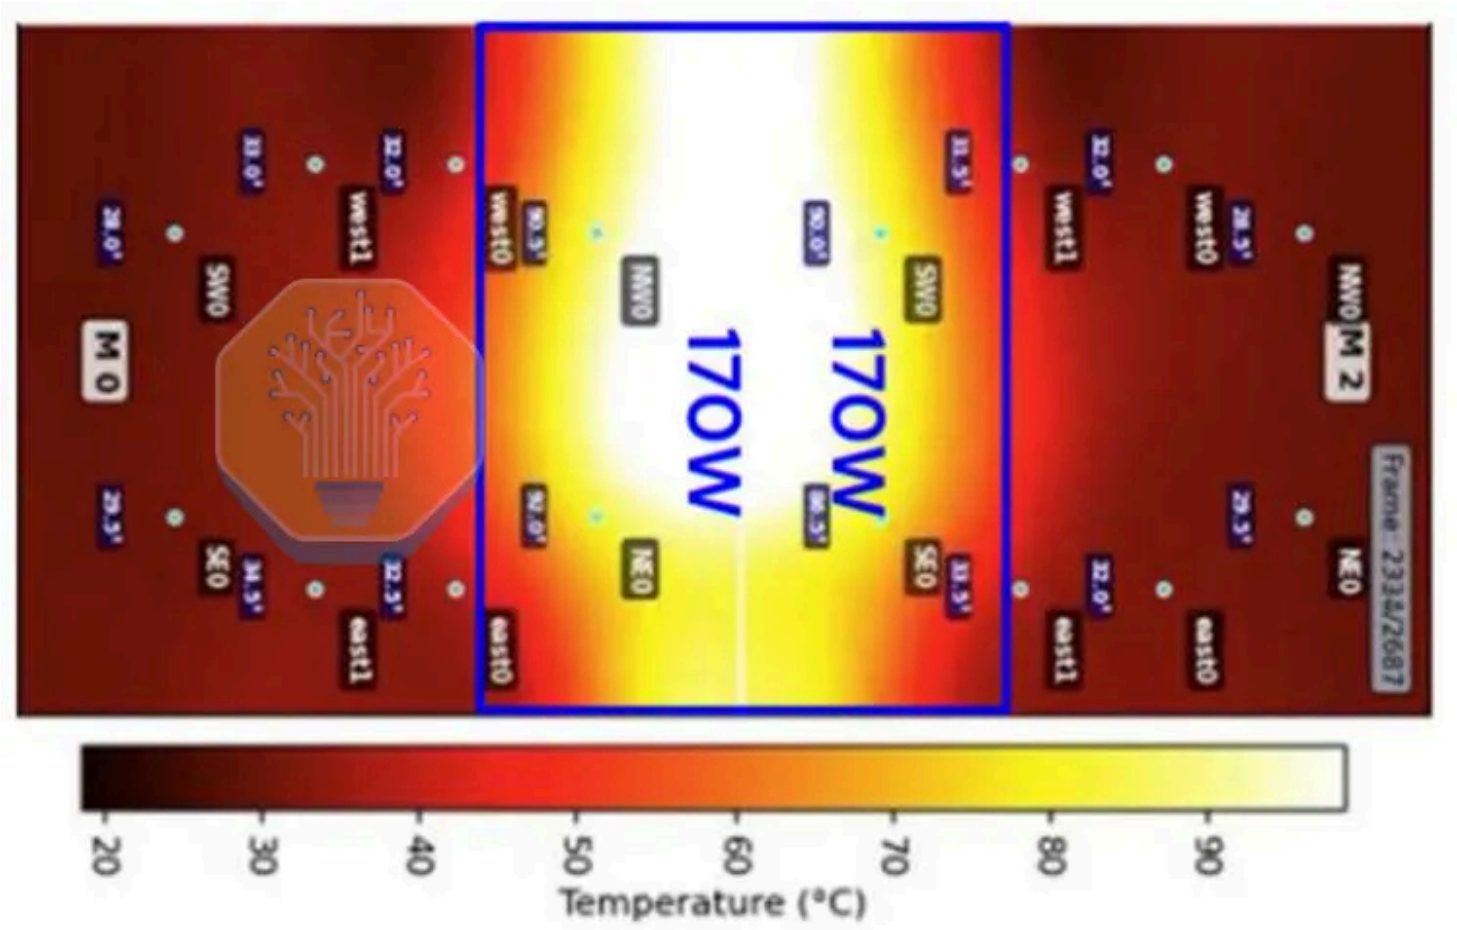
Passage M1000 package warpage during temperature ramp-up and cool-down.
Source: Lightmatter, “Advancing Interconnect Performance and Reliability with

Innovations in 3D Photonic Integration Packaging and Fiber Coupling," ECTC 2026

M1000 封装在升温与降温过程中的翘曲情况。来源：Lightmatter，《通过 3D 光子集成封装与光纤耦合创新提升互连性能与可靠性》，ECTC 2026

Attaching a silicon interposer of this size to an organic substrate creates severe warpage. The module reached ~59 µm of warpage at the 260 °C reflow temperature, and ~56 µm after cooling back down to room temperature. With a 118 µm-thick interposer and C4 bumps at ~176 µm pitch, this is enough to compromise joint formation. Lightmatter used a magnetic fixture to hold the substrate flat during attachment and reported >95% electrical assembly yield, with healthy microbump and C4 joints across the package.

将如此尺寸的硅中介层贴装至有机基板会引发严重翘曲。该模块在 260°C 回流焊温度下产生~59 微米翘曲，冷却至室温后仍残留~56 微米翘曲。采用 118 微米厚中介层及 ~176 微米间距的 C4 凸点，此翘曲量足以影响焊点成型。Lightmatter 通过磁性夹具在贴装过程中保持基板平整，最终实现>95%的电气组装良率，整个封装体上的微凸点与 C4 焊点均表现良好。



Thermal map of the test chip at 170 W per quadrant. Source: Lightmatter, "Advancing Interconnect Performance and Reliability with Innovations in 3D

Photonic Integration Packaging and Fiber Coupling," ECTC 2026

测试芯片在每象限 170 瓦功率下的热分布图。来源：Lightmatter，《通过 3D 光子集成封装与光纤耦合创新提升互连性能与可靠性》，ECTC 2026

Lightmatter used a thermal test chip with four independently powered quadrants, each dissipating 170 W. This resulted in a power density of 1.47 W/mm^2 across the 369 mm^2 active area. At this power, the photonic interposer reached approximately 100 °C using a 25 °C coolant flowing at 1.8 LPM/kW. This validates cooling 680 W from a concentrated test-chip area in a package designed for more than 900 W across nearly three reticles of ASIC silicon.

Lightmatter 使用了一个热测试芯片，该芯片包含四个独立供电的象限，每个象限功耗为 170 瓦。这使得 369 平方毫米有源区域内的功率密度达到 1.47 瓦/平方毫米。在此功率下，光子中介层在 25°C 冷却液以 1.8 升/分钟/千瓦的流速作用下，温度升至约 100°C。这验证了在专为超过 900 瓦功耗设计、覆盖近三个光罩尺寸 ASIC 芯片的封装中，能够从集中式测试芯片区域有效冷却 680 瓦热量。

Other Highlights 其他亮点

Hybrid Bonding 混合键合

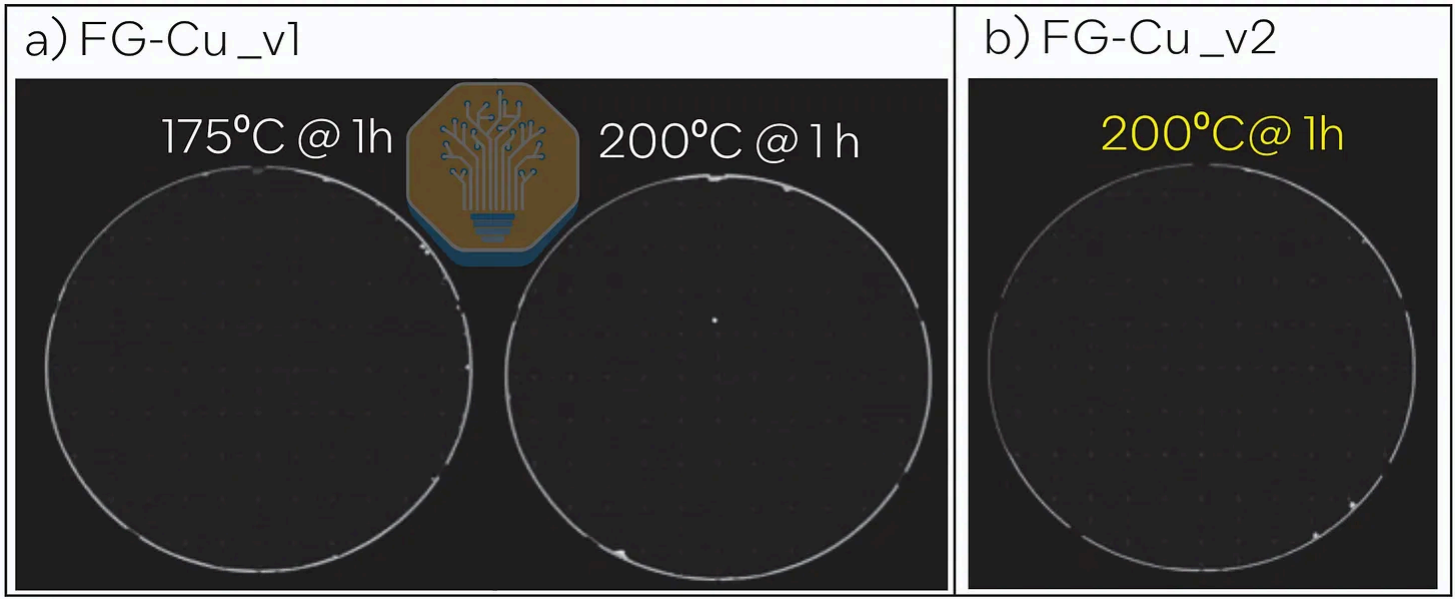
Hybrid copper bonding still delivers the finest pitch and highest I/O density for HPC applications. The open problem is keeping the interface extremely flat and clean while reducing the bonding temperature. Two material approaches stood out this year.

混合铜键合技术仍能为高性能计算应用提供最精细的间距和最高的 I/O 密度。当前面临的开放性问题是如何在降低键合温度的同时，保持界面的极度平坦与洁净。今年有两类材料方案脱颖而出。

The first uses organic dielectrics, whose mechanical compliance increases tolerance to particles and surface roughness while reducing bonding stress. Mitsui Chemicals and ASE demonstrated pressure-less Cu/polymer bonding at 200 °C and 10 µm pitch. TOK and NYCU demonstrated a 10-second bonding process at 150 °C. Samples bonded at

200 °C maintained stable resistance through reliability testing.

第一种方法采用有机介电材料，其机械柔顺性可提升对颗粒和表面粗糙度的容忍度，同时降低键合应力。三井化学与日月光半导体展示了在 200°C、10 微米间距下实现的无压铜/聚合物键合技术。东京应化工业与阳明交通大学则演示了 150°C 条件下仅需 10 秒的键合工艺。经 200°C 键合的样品在可靠性测试中保持稳定的电阻值。



Void-free C-SAM bonding maps for fine-grain Cu annealed at 175 °C and 200 °C. Source: Intel, “Enabling Ultra Low Temperature Hybrid Bonding for D2W Scaling,” ECTC 2026

175°C 与 200°C 退火条件下细晶铜的无空洞 C-SAM 键合图谱。来源：英特尔，《面向晶圆级键合尺寸缩放的超低温混合键合技术》，ECTC 2026

The second approach uses fine-grain copper. Its higher grain-boundary density accelerates copper diffusion at lower temperatures, and subsequent grain growth increases conductivity. Intel combined fine-grain copper with a low-temperature dielectric stack, achieving uniform wafer bonding after 175 °C and 200 °C anneals. Electrical yield was around 60% in two of three samples, although Intel described these results as a lower bound due to test vehicle and probing limitations. The experiments used wafer-to-wafer (W2W) test vehicles rather than the die-to-wafer (D2W) process targeted by the technology.

第二种方法采用细晶铜。其高晶界密度可在低温下加速铜扩散，后续晶粒生长则能提升导电率。英特尔将细晶铜与低温介电材料堆叠相结合，在 175°C 和 200°C 退火后实现了均匀的晶圆键合。三个样品中有两个的电学良率达到约 60%，但英特尔指出，受测试载体和探针限制，该结果仅为良率下限。实验采用晶圆对晶圆（W2W）测试载体，而非该技术目标应用的芯片对晶圆（D2W）工艺。

The most aggressive pitch came from Applied Materials and EV Group, which demonstrated 450 nm pitch wafer-to-wafer bonding at 98% yield across a chain of 20 million links. Failure analysis associated open links with carbon-rich benzotriazole (BTA) residue at the copper interface. A PVD TaN/Ta barrier stack significantly improved yield. CEA-Leti separately achieved greater than 97% yield after a 100 °C anneal without plasma activation.

应用材料公司与 EV 集团展示了最激进的间距方案，在 2000 万个互连链路上实现了 450 纳米间距的晶圆对晶圆键合，良率达 98%。失效分析显示，开路故障与铜界面处富碳苯并三唑（BTA）残留物有关。采用 PVD TaN/Ta 阻挡层堆叠可显著提升良率。法国原子能委员会电子与信息技术实验室（CEA-Leti）则在不经等离子体激活的条件下，经过 100°C 退火后实现了超过 97%的良率。

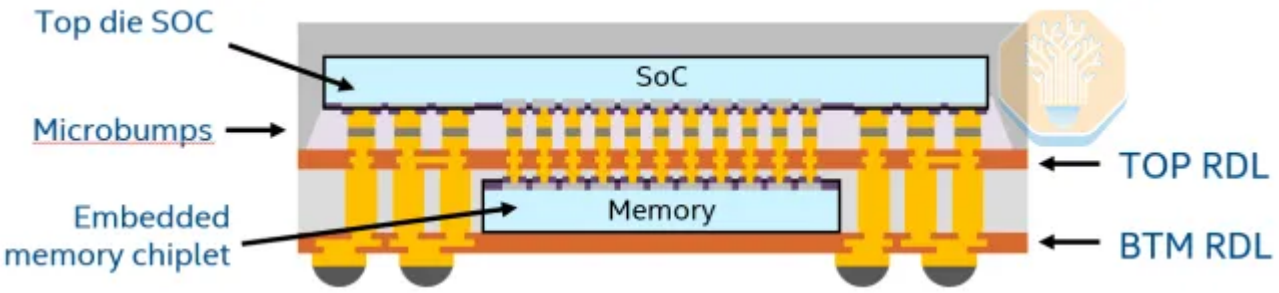
Together, these results show that reducing pitch and bonding temperature requires the copper, dielectric, CMP, surface preparation, and annealing to be co-optimized to achieve hybrid bonding with low warpage and no cracking. We expect continued refinement by material suppliers and equipment vendors to improve post-bond yield from 2027 onward.

这些结果表明，缩小间距和降低键合温度需要铜材料、介电层、CMP 工艺、表面处理与退火工艺协同优化，方能实现低翘曲、无裂纹的混合键合。预计从 2027 年起，材料供应商和设备制造商将持续改进工艺，以提升键合后良率。

Interposer Alternatives 中介层替代方案

As package sizes outgrow the practical limits of circular silicon interposers, more vendors are proposing interposer-less integration schemes beyond Intel's EMIB-T.

随着封装尺寸突破圆形硅中介层的实际极限，越来越多的供应商开始提出超越英特尔 EMIB-T 的无中介层集成方案。



Schematic cross-section of FO-EB with embedded memory chiplet. Source: Intel and SPIL, "3D Integration of an SRAM chiplet in Fan-Out Embedded Bridge

Platform Achieving Low Energy Read/Write," ECTC 2026

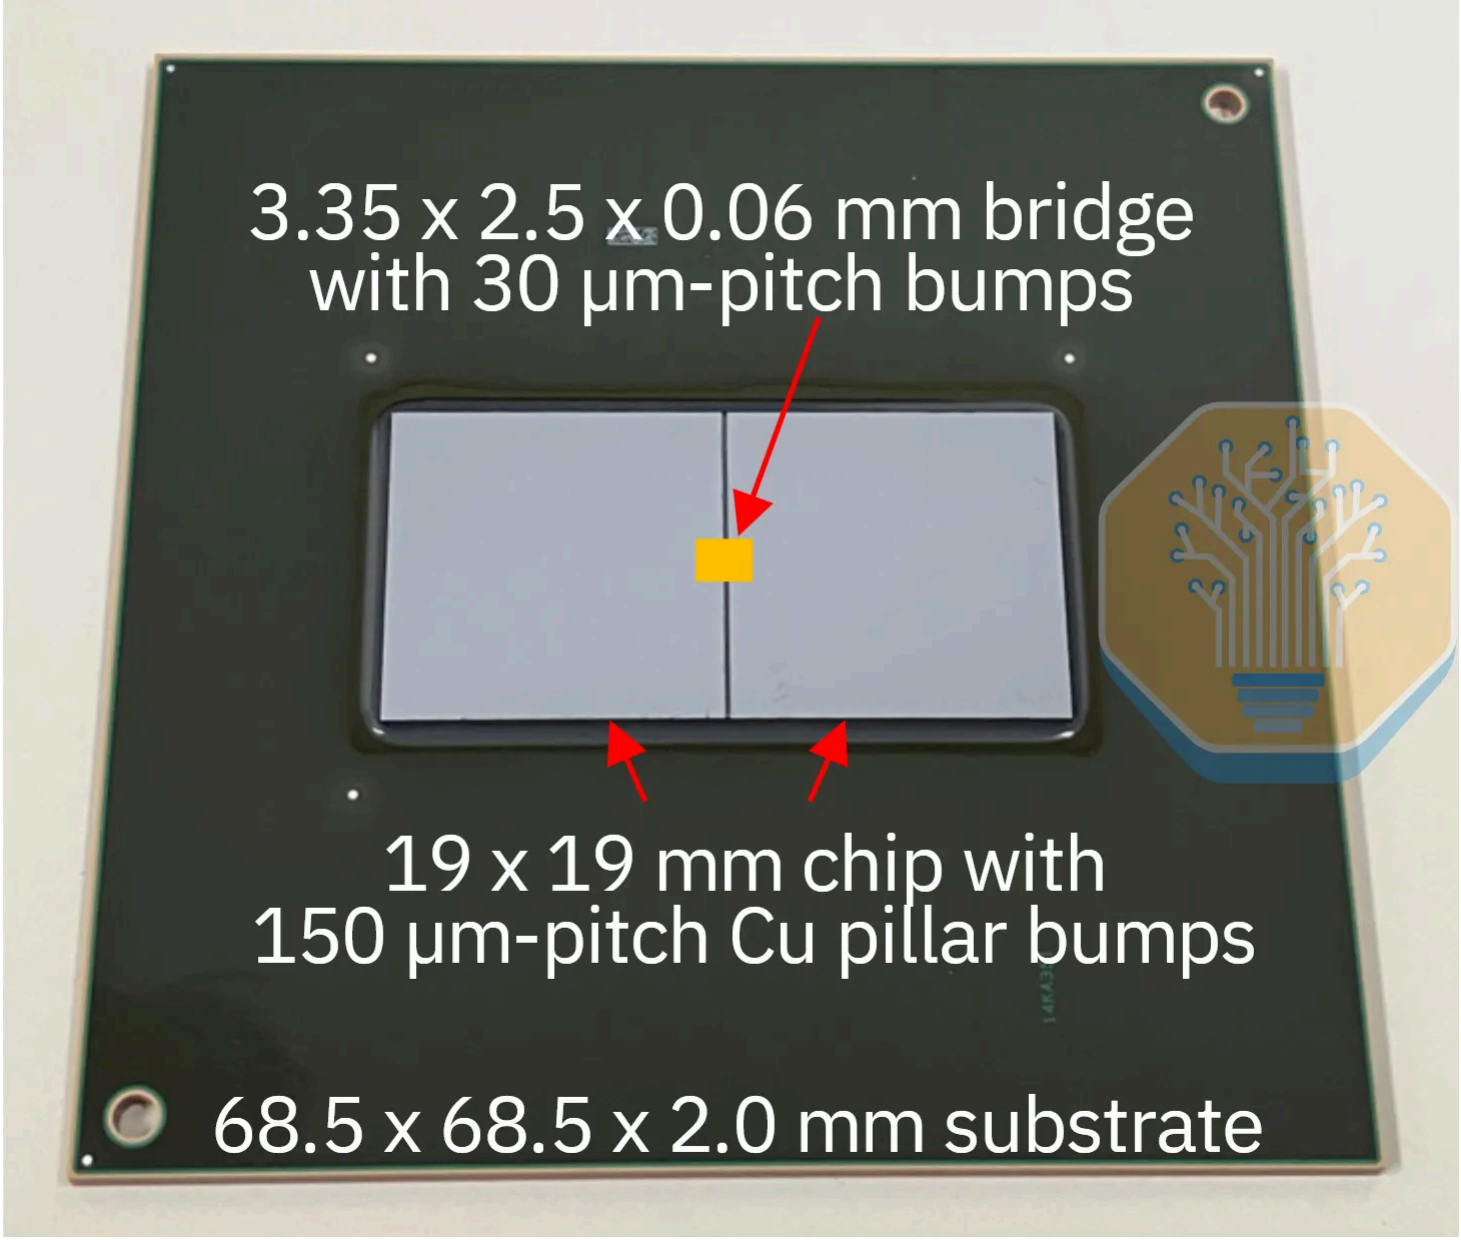
嵌入式存储芯片的扇形嵌入式桥接封装横截面示意图。来源：英特尔与硅品，"在扇形嵌入式桥接平台中实现 SRAM 芯片三维集成，达到低能耗读写性能"，ECTC 2026

Intel and SPIL placed an SRAM chiplet in the embedded-bridge layer of a fan-out embedded bridge (FO-EB) package, connecting it vertically to the logic die through 25 µm pitch microbumps. The test chip reached more than 265 GB/s/mm^2 at 0.24 pJ/b.

英特尔与硅品将一颗 SRAM 芯片植入扇形嵌入式桥接封装的桥接层，通过间距 25 微米的微凸点垂直连接至逻辑芯片。测试芯片在 0.24 皮焦/位的能耗下实现了超过 265 吉字节/秒/平方毫米的带宽密度。

Panel-scale organic interposers provide another route around the size limits of silicon. Resonac demonstrated the individual process modules for a dry-film embedded-bridge interposer on a 320 mm × 320 mm panel, including 5 µm microvias and 2/2 µm L/S. ASE fabricated an RDL on a 600 mm × 600 mm panel, then divided it into four 300 mm × 300 mm panels for assembly with existing equipment. Its test vehicle used three RDL layers at 5/8 µm L/S.

面板级有机中介层为突破硅材料尺寸限制提供了另一条路径。瑞萨电子展示了基于 320 毫米×320 毫米面板的干膜嵌入式桥接中介层各工艺模块，包含 5 微米微通孔和 2/2 微米线宽/线距。日月光半导体在 600 毫米×600 毫米面板上制作了再分布层，随后将其分割为四块 300 毫米×300 毫米的面板，以便使用现有设备进行组装。其测试载具采用三层再分布层，线宽/线距为 5/8 微米。



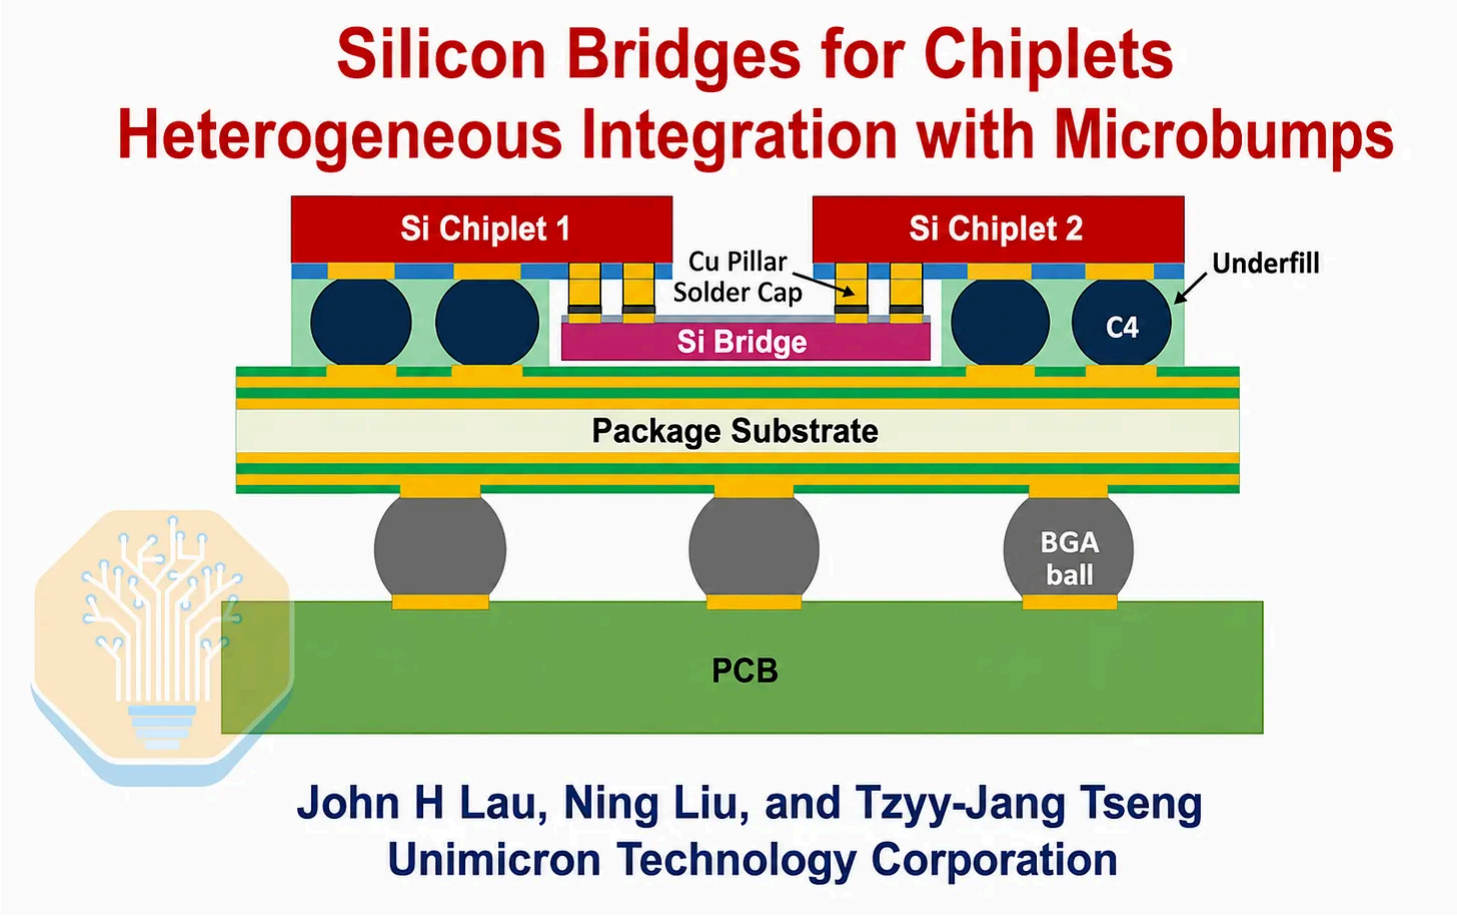
Direct Bridge Multi-die with a 30 µm pitch silicon bridge. Source: IBM, "Direct Bridge Multi-die (DBrM) Package: A Novel Silicon Bridge Chiplet Packaging Technology Using Die-Edge Gluing Technique for Chip Reconstitution," ECTC 2026

采用 30 微米间距硅桥的直接桥接多芯片封装。来源：IBM，《直接桥接多芯片（DBrM）封装：一种利用芯片边缘粘合技术实现芯片重构的新型硅桥芯片封装技术》，ECTC 2026

IBM took a more localized approach with Direct Bridge Multi-die (DBrM). Chiplets are first joined along their edges, creating a mechanically rigid subassembly around a 30 µm pitch silicon bridge. The subassembly withstood more than 30 N in bending tests, versus 0.2 N for IBM's previous underfill-only structure. This is a positive result

for reducing warpage.

IBM 采用更局部化的方法开发了直接桥接多芯片（DBrM）技术。芯片首先沿边缘进行键合，在 30 微米间距的硅桥周围形成一个机械刚性组件。该组件在弯曲测试中承受了超过 30 牛顿的力，而 IBM 此前仅使用底部填充的结构仅能承受 0.2 牛顿。这对减少翘曲而言是一个积极成果。



Chiplets connected through a silicon bridge without an interposer. Source: Unimicron, “Silicon Bridges for Chiplets Heterogeneous Integration with Microbumps,” ECTC 2026

通过硅桥连接的芯片，无需中介层。来源：欣兴电子，《面向带微凸点的芯片异构集成的硅桥技术》，ECTC 2026

Unimicron modeled an even simpler structure requiring neither an interposer nor embedded bridges. Two chiplets connect through a thin silicon bridge mounted beneath them while attaching directly to the substrate. Unimicron’s simulations show that underfill between the chiplets and bridge is necessary to control microbump strain.

欣兴电子模拟了一种更简单的结构，既不需要中介层，也不需要嵌入式桥接。两个芯片通过安装在其下方的薄硅桥进行连接，并直接附着在基板上。欣兴电子的仿真结果表明，芯片与硅桥之间的底部填充对于控制微凸点应变至关重要。

While TSMC's CoWoS-R and CoWoS-L remain limited by the circular wafer on which their RDLs are fabricated, constraining package size and wafer utilization, these alternatives shift integration to panel-level or reconstituted formats or eliminate the interposer entirely. We expect similar architectures to begin appearing in ASICs over the next several years.

尽管台积电的 CoWoS-R 和 CoWoS-L 仍受限于其 RDL 制造所用的圆形晶圆，导致封装尺寸和晶圆利用率受限，但这些替代方案将集成方式转向面板级或重组格式，或者完全省去中介层。我们预计未来几年内，类似的架构将开始出现在 ASIC 中。

Thermal Interface Materials

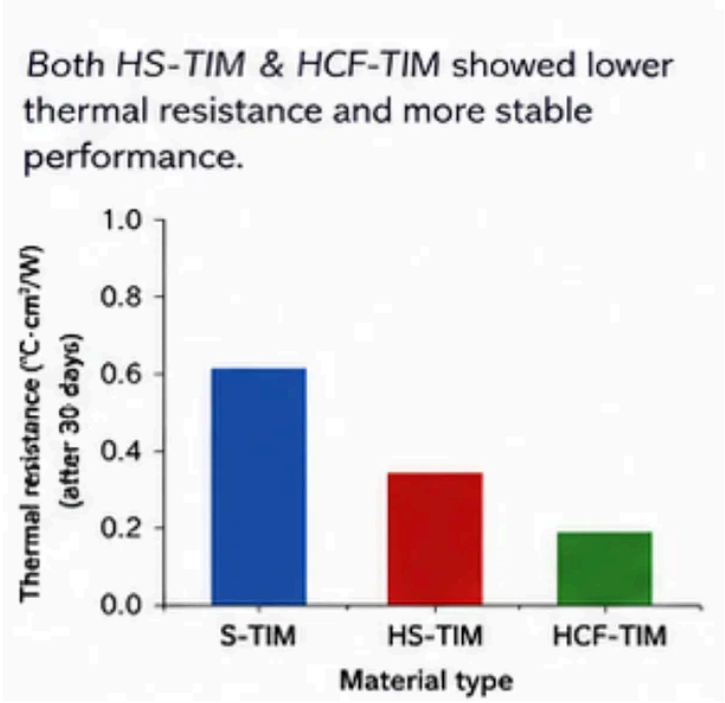
热界面材料

In advanced thermal solutions, TSMC and its partners continue to dominate the show. TSMC's direct-to-silicon cooling removes TIM1 entirely, but most near-term systems will still need better materials between the silicon and heat spreader. TSMC's OSAT partner, SPIL, tested gallium-based liquid-metal (LM) composites, a silicone-based HS-TIM and a carbon fiber HCF-TIM, in a 55 mm × 55 mm FO-EB package. Their measured conductivities were 5.7 W/m·K and 10 W/m·K respectively, and both produced lower thermal resistance than a commercial silicone TIM at 4 W/m·K.

在先进散热方案方面，台积电及其合作伙伴持续主导展会。台积电的直触硅散热技术完全省去了 TIM1，但大多数近期系统仍需要在硅芯片与散热盖之间使用更优质的材料。台积电的封测合作伙伴 SPIL 在 55mm×55mm 的 FO-EB 封装中测试了镓基液态金属复合材料、硅基 HS-TIM 和碳纤维 HCF-TIM。经测量，后两种材料的导热系数分别为 5.7 W/m·K 和 10 W/m·K，两者均比商用硅基 TIM 的 4 W/m·K 产生了更低的热阻。

Properties of selected TIMs

Traditional S-TIM	HS-TIM	
Al + ZnO	LM alloys + Al_2O_3	+
4 W/mK	5.7 W/mK	



Thermal performance and reliability of liquid-metal HS-TIM and HCF-TIM. Source: SPIL, "Reliability Assessment of Liquid Metal Alloy Thermal Interface Materials in Fan-Out Embedded Bridge (FO-EB) Packages," ECTC 2026

液态金属 HS-TIM 与 HCF-TIM 的热性能与可靠性。来源：SPIL，“扇出嵌入式桥接（FO-EB）封装中液态金属合金热界面材料的可靠性评估”，ECTC 2026

The two TIMs exhibited drastically different reliability outcomes, with HCF-TIM maintaining 95% coverage after 1000 hours at 150 °C while HS-TIM fell to only 75% as its silicone matrix hardened and partially delaminated. Both LM-based TIMs reduced thermal resistance relative to conventional S-TIM, with HCF-TIM delivering the best performance and reliability.

两种 TIM 展现出截然不同的可靠性结果：HCF-TIM 在 150°C 下经过 1000 小时后仍保持 95%的覆盖率，而 HS-TIM 的硅胶基体因硬化并部分分层，覆盖率降至仅 75%。两种液态金属基 TIM 相比传统 S-TIM 均降低了热阻，其中 HCF-TIM 在性能和可靠性方面表现最优。

Purdue, the University of Aveiro, and UCLA took a different approach, embedding Cu/Sn microbumps in nanocrystalline diamond. The resulting interconnect layer reached an effective in-plane thermal conductivity of 500 to 600 W/m·K, around twenty times that of conventional microbumps with underfill. This is not a TIM1 replacement, but a way to spread heat laterally through the interconnect layer of a 3D stack. The process remains early, with the paper only using a single-sided test

structure rather than an assembled 3D stack.

普渡大学、阿威罗大学与加州大学洛杉矶分校另辟蹊径，将铜/锡微凸点嵌入纳米晶金刚石中。所制成的互连层有效面内热导率达到 500 至 600 W/m·K，约为传统带底部填充胶的微凸点的二十倍。这并非取代 TIM1，而是通过 3D 堆叠的互连层实现横向热量扩散。该技术仍处于早期阶段，论文仅采用单侧测试结构，尚未完成完整的 3D 堆叠组装。

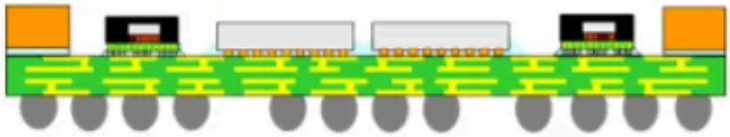
On the other hand, we saw little discussion on SiC as TIM1 or part of thermal solutions, implying a relatively long way till maturity.

另一方面，关于碳化硅作为 TIM1 或热解决方案组成部分的讨论鲜有耳闻，这意味着该技术距离成熟仍有一段相当长的路要走。

Glass Substrates 玻璃基板

Glass's aurora dimmed somewhat this year, with fewer innovative papers presented at ECTC. However, the unsolved problem is still SeWaRe, the lateral crack that begins at a diced glass edge under RDL stress. Georgia Tech characterized the failure experimentally. Corning used finite element analysis (FEA), peridynamics, and analytical fracture mechanics to model its propagation, showing that stiff copper layers drove cracks towards the glass midplane, while compliant polymer layers changed the crack path. Corning also found that low-CTE polymers combined with appropriate glass selection could reduce the failure risk.

玻璃基板的势头今年有所减弱，在 ECTC 上展示的创新论文数量减少。然而，悬而未决的问题仍是 SeWaRe——即在 RDL 应力下从切割玻璃边缘开始的横向裂纹。佐治亚理工学院通过实验对这类失效进行了表征。康宁公司运用有限元分析、近场动力学及分析断裂力学对其扩展过程进行建模，结果表明刚性铜层会将裂纹引向玻璃中平面，而柔性聚合物层则会改变裂纹路径。康宁公司还发现，低热膨胀系数聚合物与合适的玻璃搭配使用，可以降低失效风险。

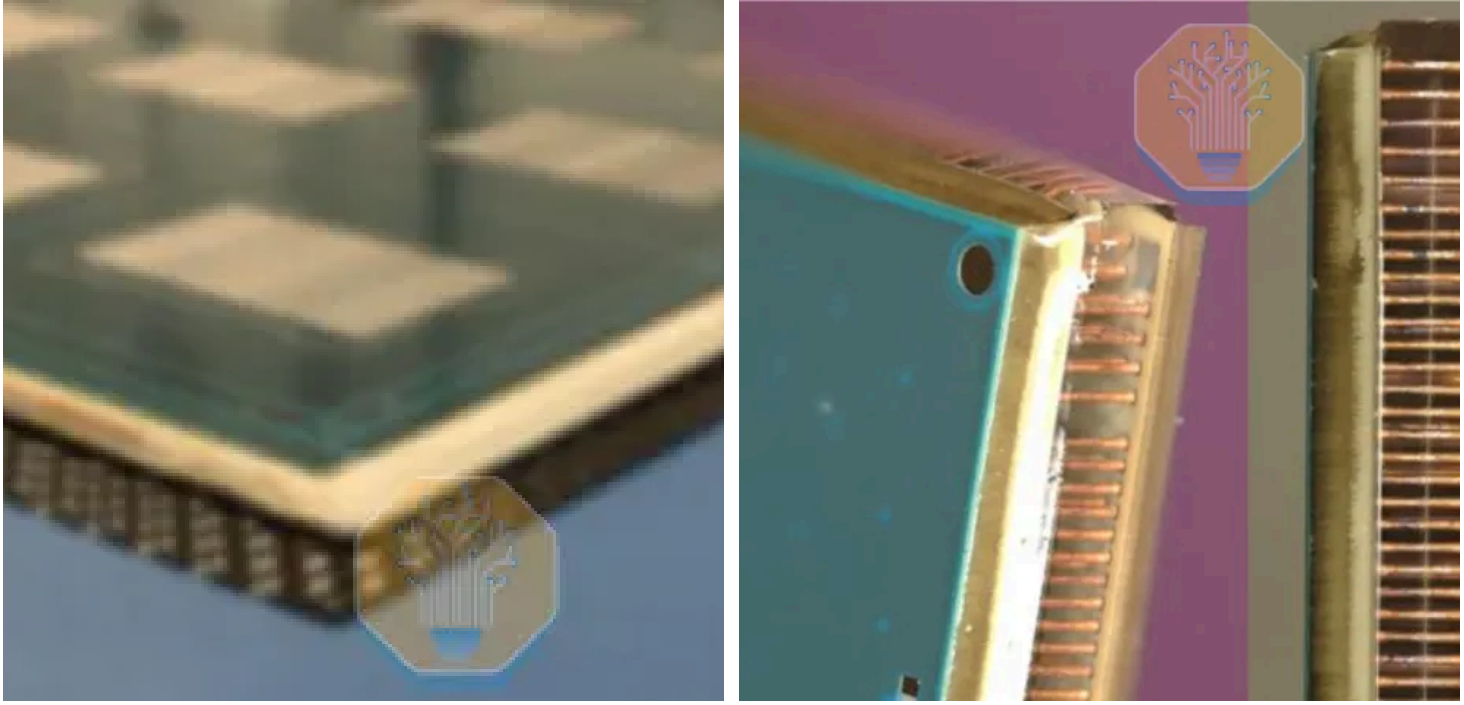
Category	Items	Organic TV	Glass TV
PKG	Configuration		
	Size	74x74mm	
	Lid Size/ Thickness	73.4x73.4mm/ 2.5mmT	
Die	Size	25x25mm (X2)	
	Bump Pitch	130um	
Memory	Size	15x10mm (X6)	
	BGA Pitch	0.4mm	
Substrate	Layer #	14L(6-2-6)	8L(3-2-3)
	Core / Thickness	E705G / 1.2mm	Glass / 0.6mm
	PTH/TGV Size	150um	90um
	ABF Material	GL102	
	Thickness	1.81mm	0.92mm

Comparison of the 74 × 74 mm organic- and glass-core test vehicles. Source: STATS ChipPAC, "Assembly and Reliability Characterization of Glass-Cored Substrate Package for AI/HPC Applications," ECTC 2026

74 × 74 mm 有机核与玻璃核测试载具对比。来源：STATS ChipPAC，《面向 AI/HPC 应用的玻璃核基板封装的组装与可靠性表征》，ECTC 2026

STATS ChipPAC investigated the assembly and reliability of large glass-core packages. Its 74 mm × 74 mm glass-core packages failed every test segment without edge coating, while the edge-coated packages completed assembly and reliability testing without abnormalities. The edge coating also reduced warpage by 33.5% relative to the uncoated glass-core package. Build-up pull-back and edge coating increasingly look like requirements for reliable glass-core substrate assembly.

STATS ChipPAC 研究了大尺寸玻璃核心封装的组装与可靠性。其 74 毫米×74 毫米的玻璃核心封装在未采用边缘涂层的情况下，未能通过任何测试环节，而采用边缘涂层的封装则顺利完成了组装和可靠性测试，未出现异常。与未涂层的玻璃核心封装相比，边缘涂层还将翘曲度降低了 33.5%。积层回缩和边缘涂层越来越像是实现可靠玻璃核心基板组装的两项必备要求。



24-layer (10-2-10) glass-core panel with copper-filled TGVs, embedded EMIB bridges and optical waveguides, with singulated units showing no SeWaRe after thermal-shock testing. Source: Intel, “Glass Core Substrates – Next Generation Advanced Packaging Platform for AI and HPC,” ECTC 2026

24 层（10-2-10）玻璃核心面板，配备铜填充的 TGV、嵌入式 EMIB 桥和光波导，经过热冲击测试后，分割单元未出现 SeWaRe。来源：Intel，“玻璃核心基板——面向 AI 和 HPC 的下一代先进封装平台”，ECTC 2026

On a positive note, Intel demonstrated an industry-first 510 mm × 515 mm, 24-layer glass-core panel with fully copper-filled through-glass vias (TGVs), two embedded EMIB bridges, and optical waveguides co-formed between the TGVs. The large prototype was displayed at Intel’s booth. It was processed on existing organic-substrate lines, while singulated units showed no SeWaRe after thermal shock testing.

积极的一面是，Intel 展示了业界首款 510 毫米×515 毫米、24 层玻璃核心面板，该面板采用完全铜填充的玻璃通孔（TGV），内嵌两个 EMIB 桥，并在 TGV 之间共集成光波导。这一大型原型在 Intel 展台进行了展示。它是在现有的有机基板生产线上制造的，且分割单元在经过热冲击测试后未出现 SeWaRe。

As OSAT adopters, Amkor and STATS ChipPAC measured 30-40% lower substrate-level warpage with a thinner glass core than their organic references, although assembly defects and TGV filling problems show that the process remains immature.

作为 OSAT（外包半导体封装与测试）的采用者，Amkor 和 STATS ChipPAC 测得，采用更薄玻璃基板相比有机基板可降低 30%-40%的基板翘曲，不过组装缺陷和 TGV（玻璃通孔）填充问题表明该工艺仍不成熟。

Glass is making real progress, but this year's data still support manufacturing development rather than high-volume adoption.

玻璃技术确实取得了实质性进展，但今年的数据仍支持的是制造开发阶段，而非大规模量产应用阶段。

RDL Scaling RDL（重分布层）微缩

Finally, RDL L/S continues to shrink even as package sizes grow. The primary driver is UCIe 3.0, which supports speeds of up to 64 GT/s for future ASIC-to-ASIC and ASIC-to-HBM links. These high-speed die-to-die interconnects tighten signal integrity requirements as organic interposers grow larger and denser.

最后，即便封装尺寸持续增大，RDL 的线宽/线距（L/S）仍在不断缩小。其主要驱动力来自 UCIe 3.0 标准——该标准支持未来 ASIC 间及 ASIC 与 HBM（高带宽存储器）互连高达 64 GT/s 的速率。这些高速芯片间互连对信号完整性提出更严苛要求，而有机中介层正变得更大更密。

The roadmap has progressed from 10/10 µm line/space around 2015 to 2/2 µm today, with 1/1 µm emerging as the next target. Reaching the submicron era will require major changes to both RDL routing architectures and manufacturing processes. The process is shifting from semi-additive plating towards damascene for sub-2 µm copper, where CMP planarization and low-shrinkage dielectrics become the key gating steps.

路线图从 2015 年左右的 10/10 微米线宽/间距发展到如今的 2/2 微米，而 1/1 微米正成为下一个目标。进入亚微米时代将对 RDL 布线架构和制造工艺带来重大变革。制程正从半加成法向用于 2 微米以下铜线的镶嵌工艺转变，其中 CMP 平坦化和低收缩率介电材料成为关键制约步骤。

Resonac used polymer damascene and panel CMP to form 2/2 µm wiring on a 320 mm × 320 mm glass panel, including a four-layer via-and-trench structure. imec and Fujifilm pushed damascene to 1/1 µm on 300 mm wafers. Ushio resolved 1.5/1.5 µm L/S over an 18-reticle field without stitching, with 16 exposures covering a full 510 mm × 515 mm panel. Sumitomo Bakelite and Georgia Tech showed a fully imidized liquid dielectric with only 4% cure shrinkage at a relatively low temperature of 200 °C and a

fine 2/2 µm L/S.

Resonac 采用聚合物镶嵌工艺和面板 CMP 技术，在 320 mm × 320 mm 的玻璃面板上形成了 2/2 µm 的布线，包括四层通孔与沟槽结构。imec 与富士胶片将镶嵌工艺推进至 300 mm 晶圆上的 1/1 µm 线宽。Ushio 在 18 个掩模版范围内实现了 1.5/1.5 µm 线宽/间距的无拼接曝光，通过 16 次曝光覆盖完整的 510 mm × 515 mm 面板。住友电木与佐治亚理工学院展示了一种完全亚胺化的液态介电材料，在 200°C 的较低温度下仅产生 4% 的固化收缩率，并实现了精细的 2/2 µm 线宽/间距。

6. UCIe-A IP & CoWoW-R Silicon Correlation

This work correlates at two silicon measured approaches, which are (1) high speed interconnect test keys in S-Param. by VNA; (2) 16 & 32 GT/s eyediagrams on an evaluation board environment by oscilloscope.

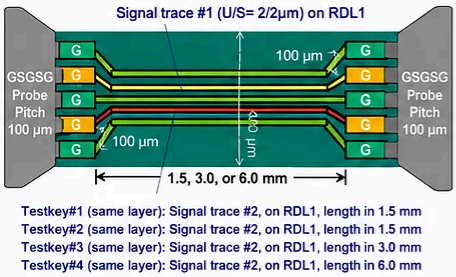




Fig. 6. UCIe-A interconnect test keys design & correlation.

TABLE I. S-PARAMETERS CORRELATION OF COWOS-R INTERCONNECTS TESTKEYS

Design Index	Testkey#1		Testkey#2		Testkey#3		Testkey#4	
	Sim	Mea	Sim	Mea	Sim	Mea	Sim	Mea
L/S (µm)	S:2/2 G:5/2		S:2/2 G:5/2		S:2/2 G:4/2		S:2/2 G:4/2	
Length (mm)	1.50	1.50	1.50	1.50	3.00	3.00	6.00	6.00
IL (dB)	1.75	2.06	2.03	2.05	2.82	2.67	5.84	5.90
RL (dB)	14.9	13.2	11.8	11.5	20.3	20.8	14.9	12.5
Xtalk (dB)	-43.7	-40.9	-41.3	-41.6	-35.0	-37.7	-41.3	-41.6

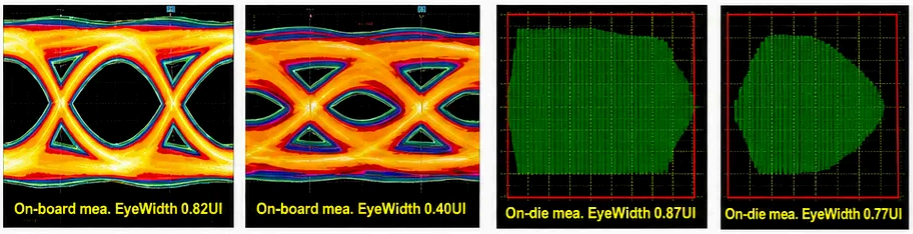


Fig. 7. UCIe-A silicon measurement with 16 and 32 GT/s.

Design flow and STCO for UCIe-A on 8-layer CoWoS-R RDL. Source: GUC, "Design of UCIe-A x64 Chiplets Integration for 16–36 GT/s Using Organic Interposer," ECTC 2026

基于 8 层 CoWoS-R RDL 的 UCIe-A 设计流程与 STCO 方法。来源：GUC，《面向 16–36 GT/s 的 UCIe-A x64 芯粒集成设计（采用有机中介层）》，ECTC 2026

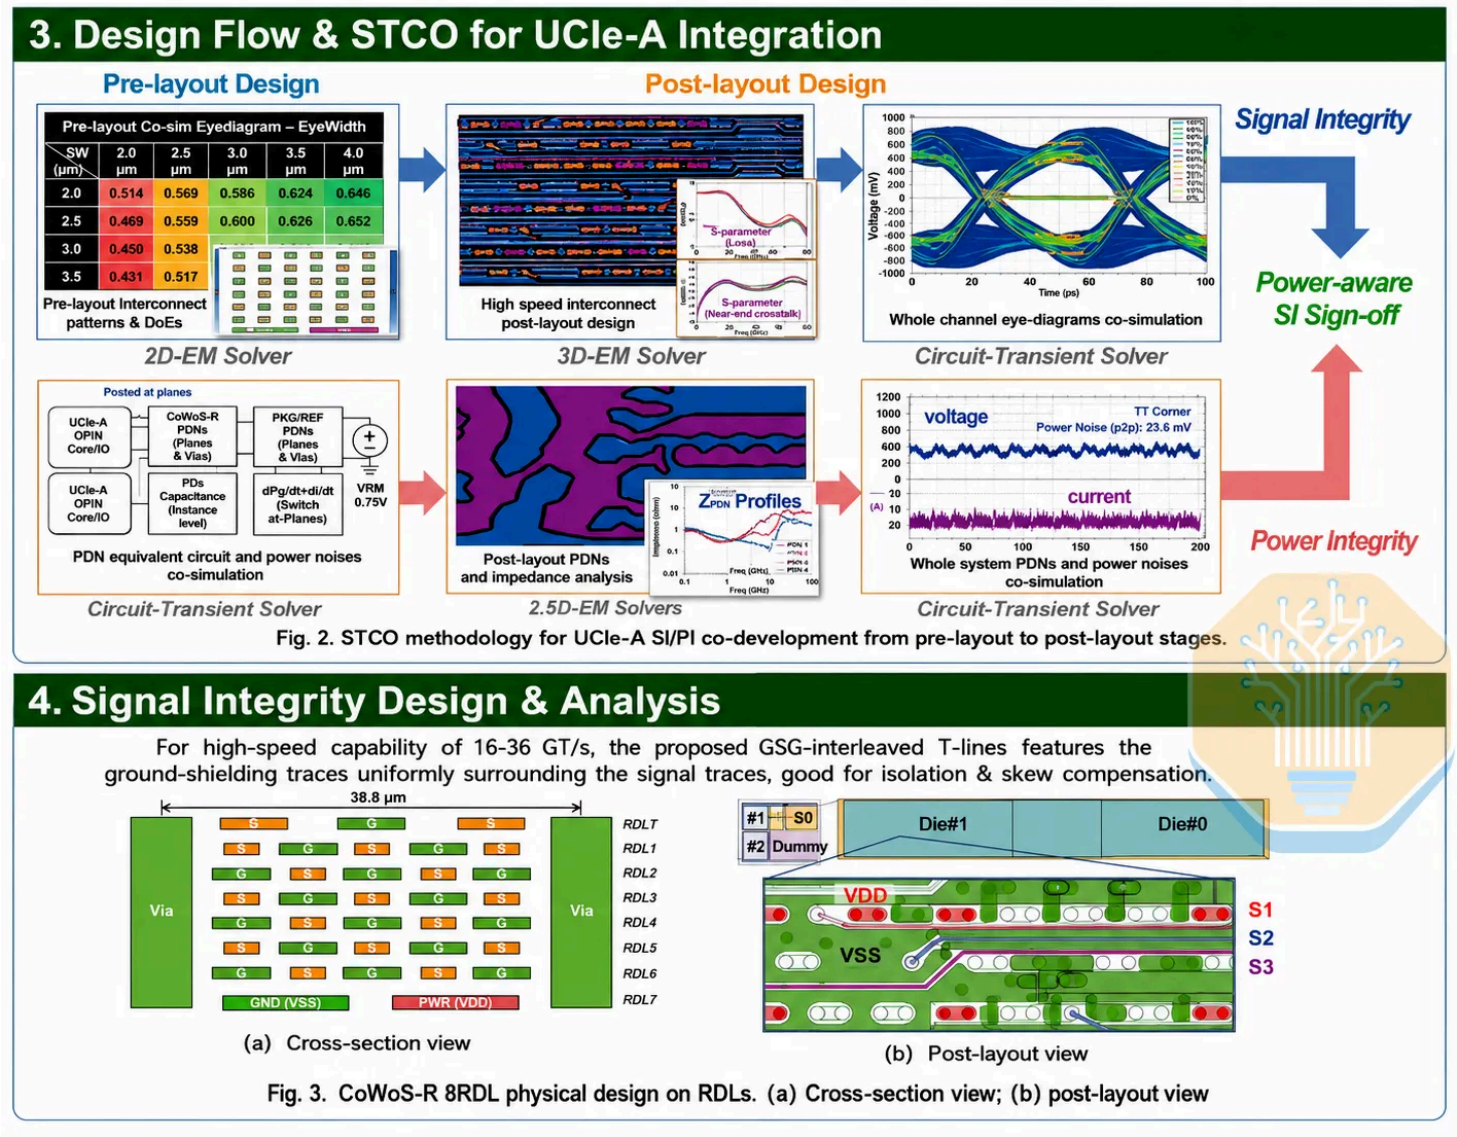
As the most advanced RDL manufacturer, TSMC also collaborated with GUC to present their work on 8-layer RDL scaling, which is believed to be the near-term limit of the CoWoS-R platform. GUC demonstrated an STCO-based design and validation flow for integrating a 64-bit UCIe-A interface fabricated on TSMC N3 and integrated

on an 8-layer CoWoS-R RDL.

作为最先进的 RDL 制造商，台积电还与 GUC 合作展示了其在 8 层 RDL 缩放方面的成果，这被认为是 CoWoS-R 平台的近期极限。GUC 演示了一种基于 STCO 的设计与验证流程，用于集成在台积电 N3 工艺上制造、并集成于 8 层 CoWoS-R RDL 的 64 位 UCIe-A 接口。

Its STCO framework uses ground-signal-ground interleaved transmission lines to control crosstalk and skew, while simulations show that C4-side IPDs provide localized decoupling and reduce voltage fluctuation at the chiplet microbumps.

其 STCO 框架采用地-信号-地交错传输线来控制串扰和时滞，而仿真表明，C4 侧 IPD 可提供局部去耦并减少小芯片微凸点处的电压波动。



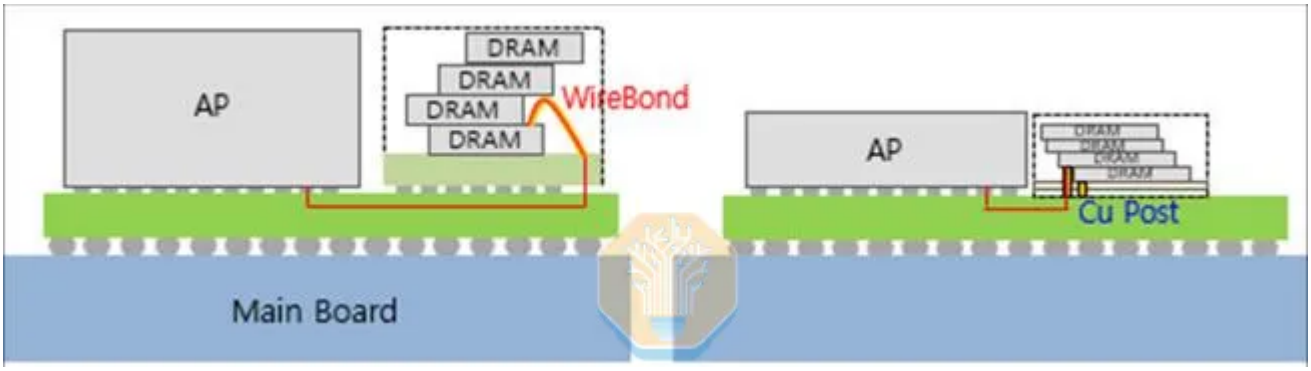
Measured UCIe-A eye diagrams at 16 and 32 GT/s. Source: GUC, “Design of UCIe-A x64 Chiplets Integration for 16–36 GT/s Using Organic Interposer,” ECTC 2026

测量得到的 UCIe-A 眼图（16 GT/s 和 32 GT/s）。图片来源：GUC，《基于有机中介层的 16–36 GT/s UCIe-A x64 芯粒集成设计》，ECTC 2026

The design targets 16-36 GT/s with a 64-bit, 10-column UCIe-A interface at 45 µm bump pitch. Signal traces were routed at 2/2 µm across six layers, with the seventh reserved for power delivery. The test chip achieved a measured on-die eye width of 0.77 UI at 32 GT/s, while simulations showed an eye width of 0.74 UI at 36 GT/s. The results show that organic interposers can meet the signal and power integrity requirements for heterogeneous chiplet systems.

该设计目标速率为 16-36 GT/s，采用 64 位、10 列 UCIe-A 接口，凸点间距为 45 µm。信号走线在六层中以 2/2 µm 规格布线，第七层预留用于电源传输。测试芯片在 32 GT/s 速率下测得的片上眼宽为 0.77 UI，而仿真结果显示在 36 GT/s 时眼宽为 0.74 UI。结果表明，有机中介层能够满足异构芯粒系统的信号与电源完整性要求。

Stacked Memory 堆叠式存储器



Conventional wire-bond stack versus Vertical Cu post Stack. Source: Samsung, "Multi Stacked FOWLP utilizing Extreme Aspect Ratio Cu Post for Mobile on-Device AI Memory Solution," ECTC 2026

传统引线键合堆叠与垂直铜柱堆叠对比。图片来源：三星，《利用超高深宽比铜柱的多层 FOWLP 用于移动端设备端 AI 内存解决方案》，ECTC 2026

Samsung showed an approach to DRAM stacking that avoids TSVs entirely. Vertical Cu post Stack (VCS) connects four stacked memory dies through <56 µm-pitch, <30 µm-wide high-aspect-ratio copper posts embedded in molding compound, rather than etching vias through silicon, while also replacing the package substrate with an RDL. Shortening the interconnect cut power consumption by 41% versus a conventional wire-bonded stack, from 0.646 W to 0.384 W at iso-speed. Samsung also reported better signaling, with maximum data rate rising from 8.6 Gb/s to 11.8 Gb/s with only

8% higher power consumption.

三星展示了一种完全避开 TSV 的 DRAM 堆叠方法。垂直铜柱堆叠（VCS）技术通过嵌入模塑化合物中的<56 微米间距、<30 微米宽的高深宽比铜柱连接四层堆叠存储芯片，而非在硅中蚀刻通孔，同时用 RDL 替代封装基板。相较于传统引线键合堆叠，这种互连缩短方式在同等速度下将功耗从 0.646W 降至 0.384W，降幅达 41%。三星还报告了更优的信号传输性能，最高数据速率从 8.6 Gb/s 提升至 11.8 Gb/s，功耗仅增加 8%。



Cross-section of the four-high VCS DRAM stack connected by extreme-aspect-ratio copper posts. Source: Samsung, "Multi Stacked FOWLP utilizing Extreme Aspect Ratio Cu Post for Mobile on-Device AI Memory Solution," ECTC 2026

四层高深宽比铜柱堆叠的 VCS DRAM 截面图。来源：三星，《采用超高深宽比铜柱的多层 FOWLP 技术用于移动端设备 AI 内存解决方案》，ECTC 2026

The form factor also improves substantially, with package height and footprint each reduced by 40% while bandwidth increases 2.6× and I/O count 6×. Although Samsung's focus is on mobile platforms like smartphones, this approach is also promising for higher-power workloads. We believe VCS and similar approaches could help future AI accelerators reach higher bandwidth at lower power and a smaller footprint, while also enabling higher-density memory modules like SOCAMMs for server CPUs.

外形尺寸也大幅改善，封装高度和面积各减少 40%，同时带宽提升 2.6 倍，I/O 数量增加 6 倍。尽管三星的重点是智能手机等移动平台，但这种方法对更高功耗的工作负载也颇具前景。我们认为 VCS 及类似方案有助于未来 AI 加速器在更小面积、更低功耗下实现更高带宽，同时还能支持服务器 CPU 使用 SOCAMM 等更高密度的内存模组。

Thanks for reading! 感谢阅读！

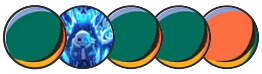
Recommend SemiAnalysis to your readers

向您的读者推荐 SemiAnalysis

Bridging the gap between the world's most important industry, semiconductors, and business.

连接全球最重要的半导体产业与商业之间的桥梁。

Recommend 推荐



46 Likes 46 个赞 · 3 Restacks 3 次转发

← Previous 上一篇

Next 接下来 →

Discussion about this post

关于此帖的讨论

Comments 评论 Restacks



Write a comment...

© 2026 Dylan Patel · Privacy · Terms · Collection notice
Substack is the home for great culture