Intel Crescent island GPU-PCB

PCB Analysis

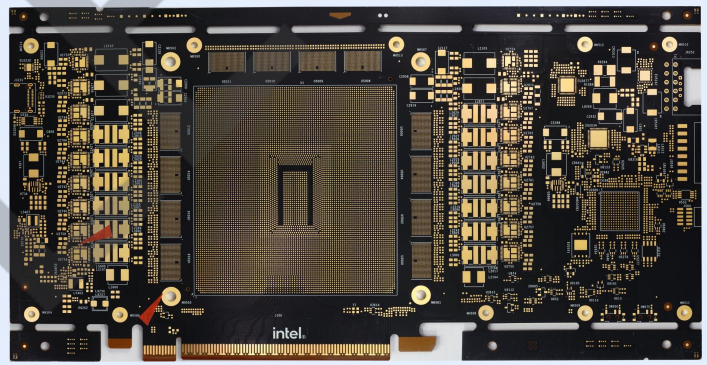

Kurnal insights

Version 1.0 | 2026 Q2

Kurnal Insights

Kurnal-insights.com

Version of this Report

Version	Date	Updates	Author
V0	2026/06/07 22:32	Analyze PCB data	Kurnal
V1	2026/06/19	Finish Full PCB Analysis	Kurnal
V1.1	2026/06/21	Debug with Flyrway	Kurnal

About Authors

This Report is made From @Kurnal

Kurnal
Insights

BiliBili: @Kurnal
X: @Kurnalsalts
WeChat: KurnalWeChat
Website: Kurnal-insights.com

VoidTech Lab | OOO实验室

Website: http://void-tech.cn/

Overview

Methodology

Executive Summary

•This report is a physical-layer teardown: PCB, power delivery, GPU package & die estimate, and the LPDDR5X memory subsystem.
•Direct measurement of PCB, package and BGA pad pitch (0.7 mm)
•Cross-section inspection of layer stack-up, copper weight and surface finish
•Die size estimated from BGA sparse-region × pad pitch; DPW from die area on 300 mm wafer
•Memory capacity & bandwidth derived from package count, BGA563 type and LPDDR5X data rate
•Benchmarked against NVIDIA / AMD reference parts

•Form factor: Shortest board in its class — 235 × 110 mm, full-height air-cooled PCIe Gen5 x16, single 12V-2x6 power.
•Power: 18 VRM phases (13 core / 3 memory / 2 I/O), estimated ~350 W board power.
•PCB: 10-layer stack-up, heavy 30 µm outer copper, ENIG surface finish.
•GPU die: ~24.5 × 31.5 mm ≈ 772 mm² (reticle-class, ~75 dies/wafer) — between AD102 (616) and H100 (814).
•Memory: up to 20× LPDDR5X BGA563 → 480 GB / 1280-bit / 1.36–1.54 TB/s — below HBM but huge capacity at low cost/power.
•Takeaway: Not built to win bandwidth benchmarks — built to host very large models cheaply, in standard servers, at scale. A capacity-and-TCO play for the inference era.

Company

Profile & Road-map

Company Profile

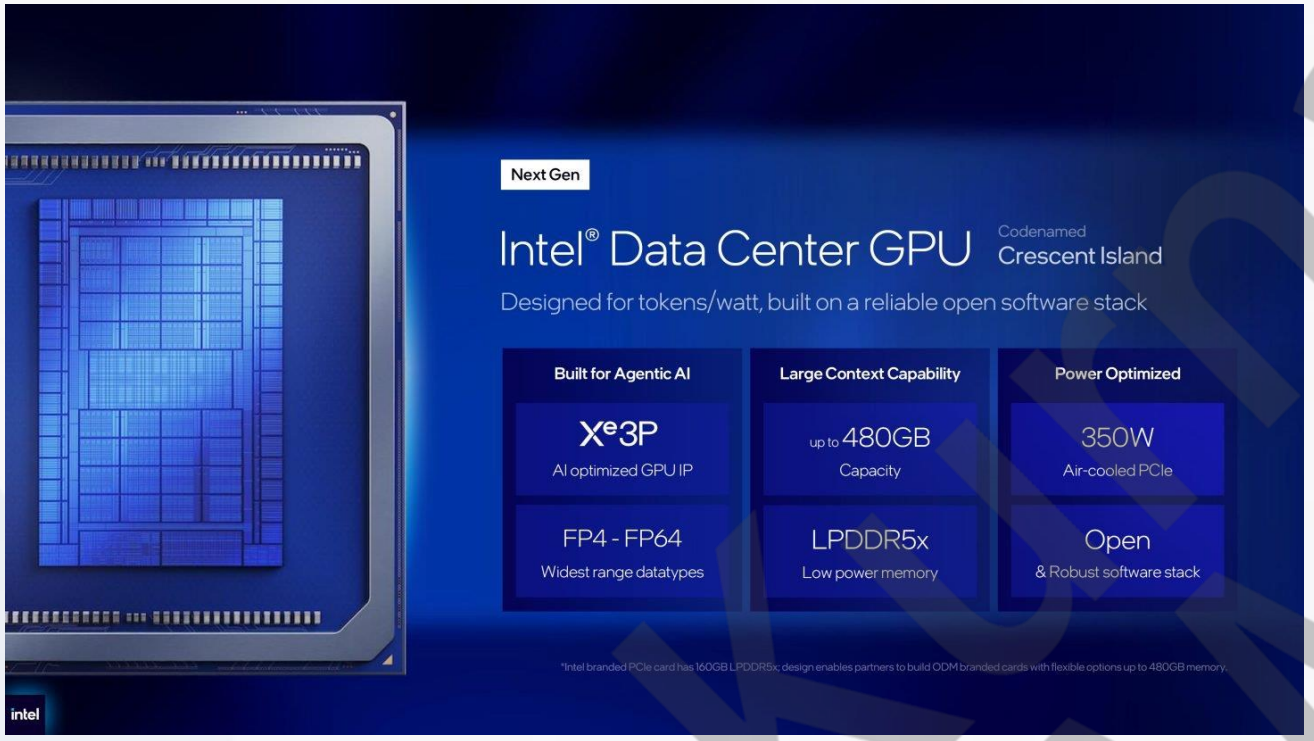

Intel Crescent Island is Intel's next-generation AI inference accelerator targeting cloud and enterprise datacenter deployments.

Built on the Xe3P architecture, Crescent Island is optimized for performance-per-watt, high memory capacity, and cost-efficient AI inference workloads rather than large-scale AI training.

Unlike competing accelerators that rely on HBM memory, Crescent Island adopts LPDDR5X memory technology, enabling significantly higher memory capacity while reducing board complexity, power consumption, and overall system cost.

The accelerator is designed for standard air-cooled PCIe server platforms, making it suitable for large-scale inference deployment and agentic AI applications.

Current disclosures indicate a reference configuration featuring 160GB LPDDR5X memory, with future partner implementations supporting capacities up to 480GB

Company Road-map

Data Center GPU Max 1100
(Ponte Vecchio)

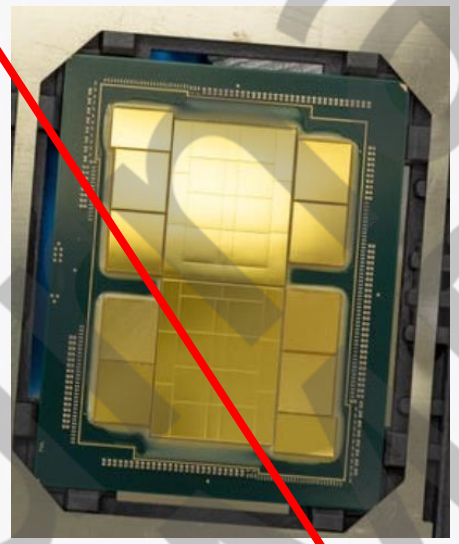

Rialto Bridge

Physical analysis

PCB-analysis

Physical analysis

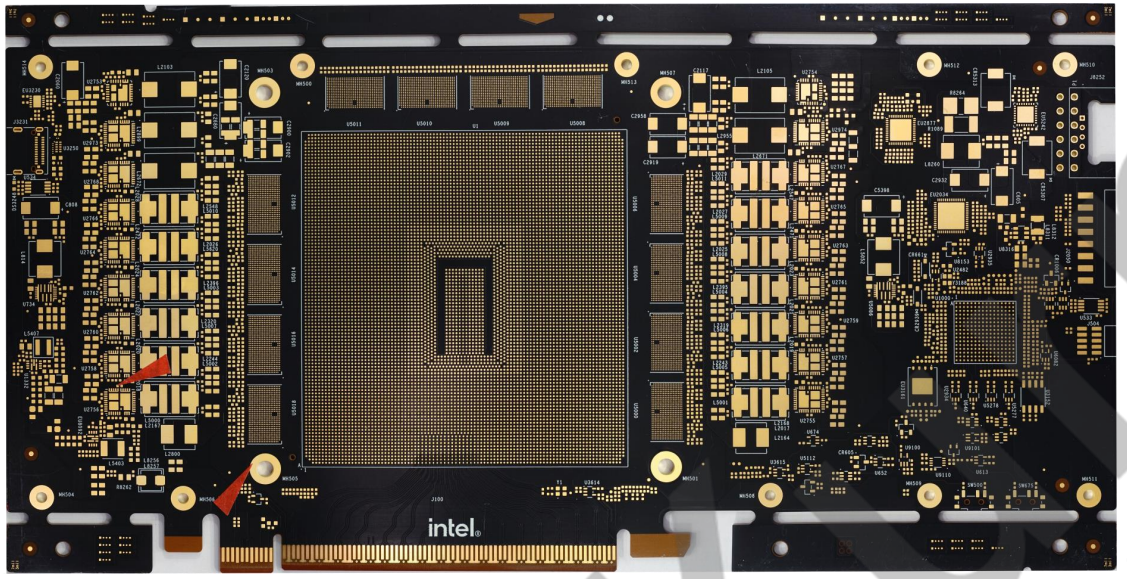

GPU-PCB Topview

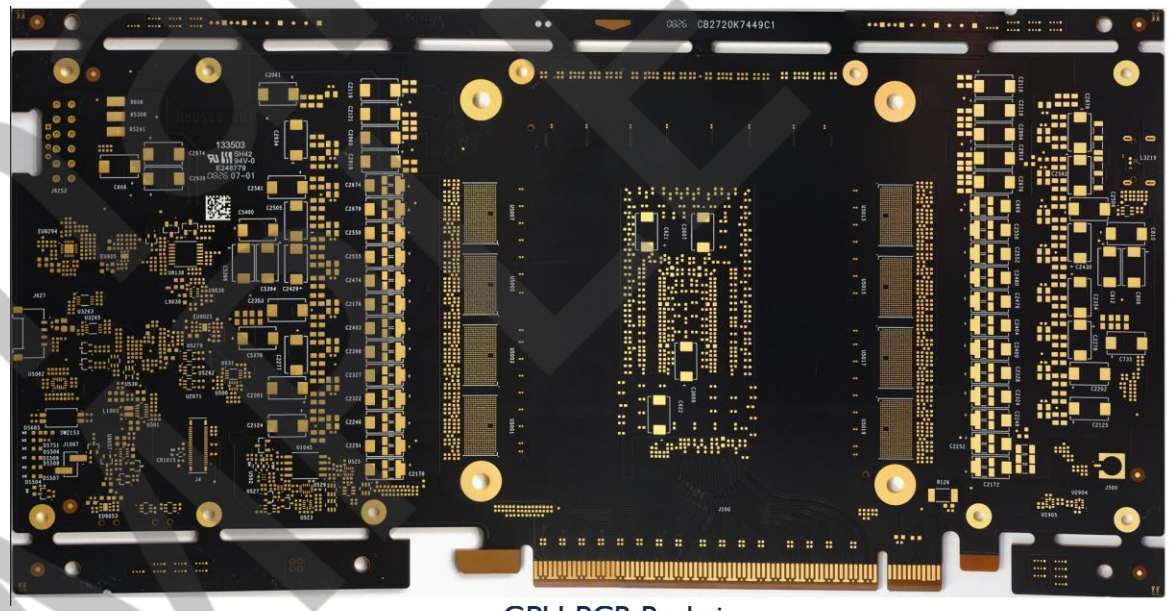

GPU-PCB Backview

我们只购买到了这张PCB板,所以只能制作一些基于这个PCB的分析报告

Physical analysis-PCB size

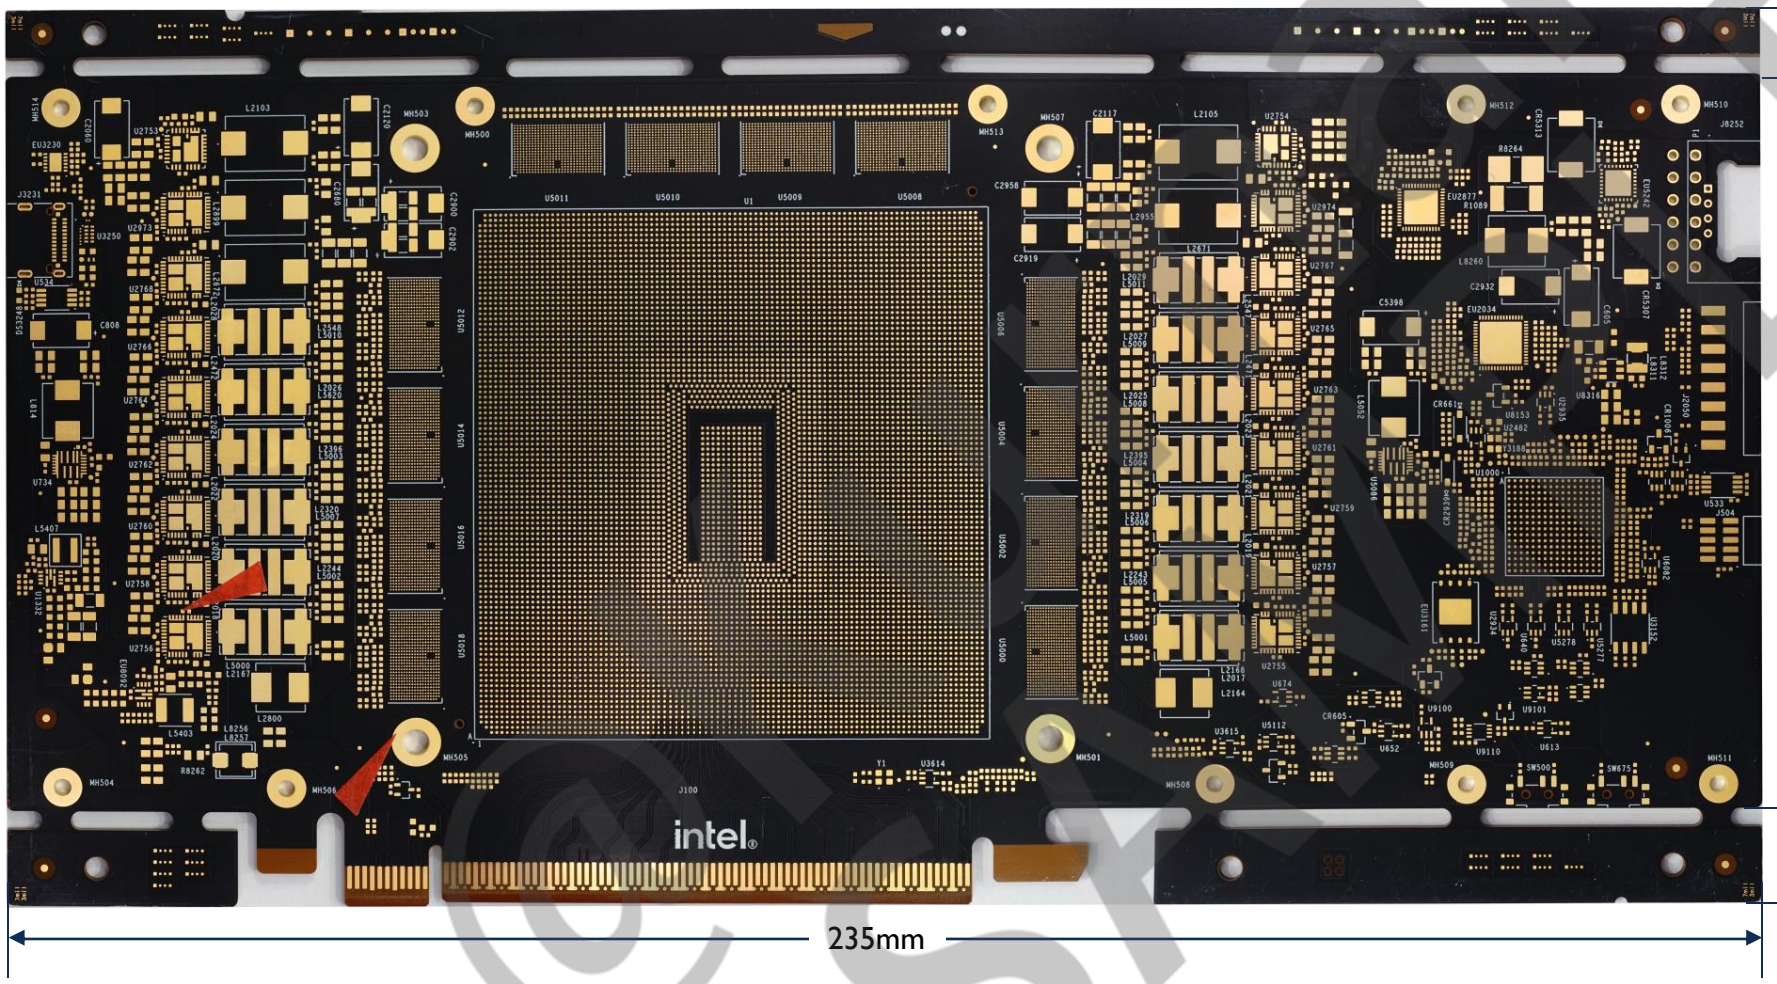

Physical analysis-PCB size

Metric	Intel Crescent island	Intel Ponte Vecchio	Nvidia H200 PCIE	Nvidia RTX 6000 Ada	Nvidia RTX 4090 FE	AMD MI210 PCIE
Length	235mm	266.7mm	267mm	267mm	304mm	290mm
Height	110mm	111.15mm	112mm	112mm	137mm	125mm

PCB Mark Analysis

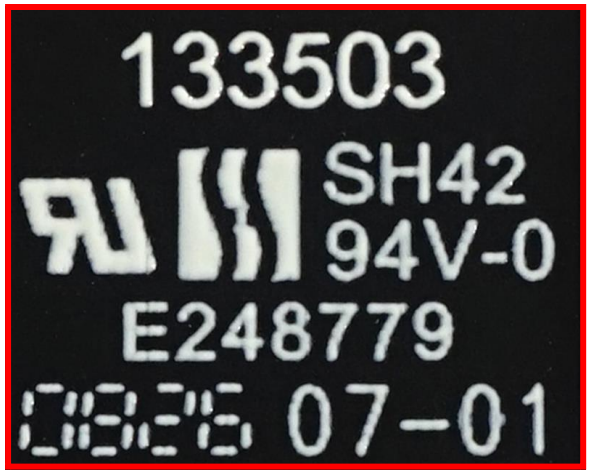

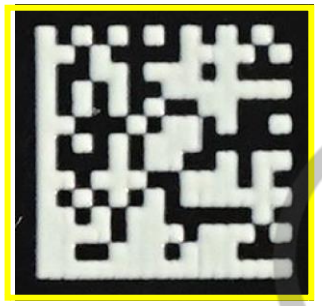

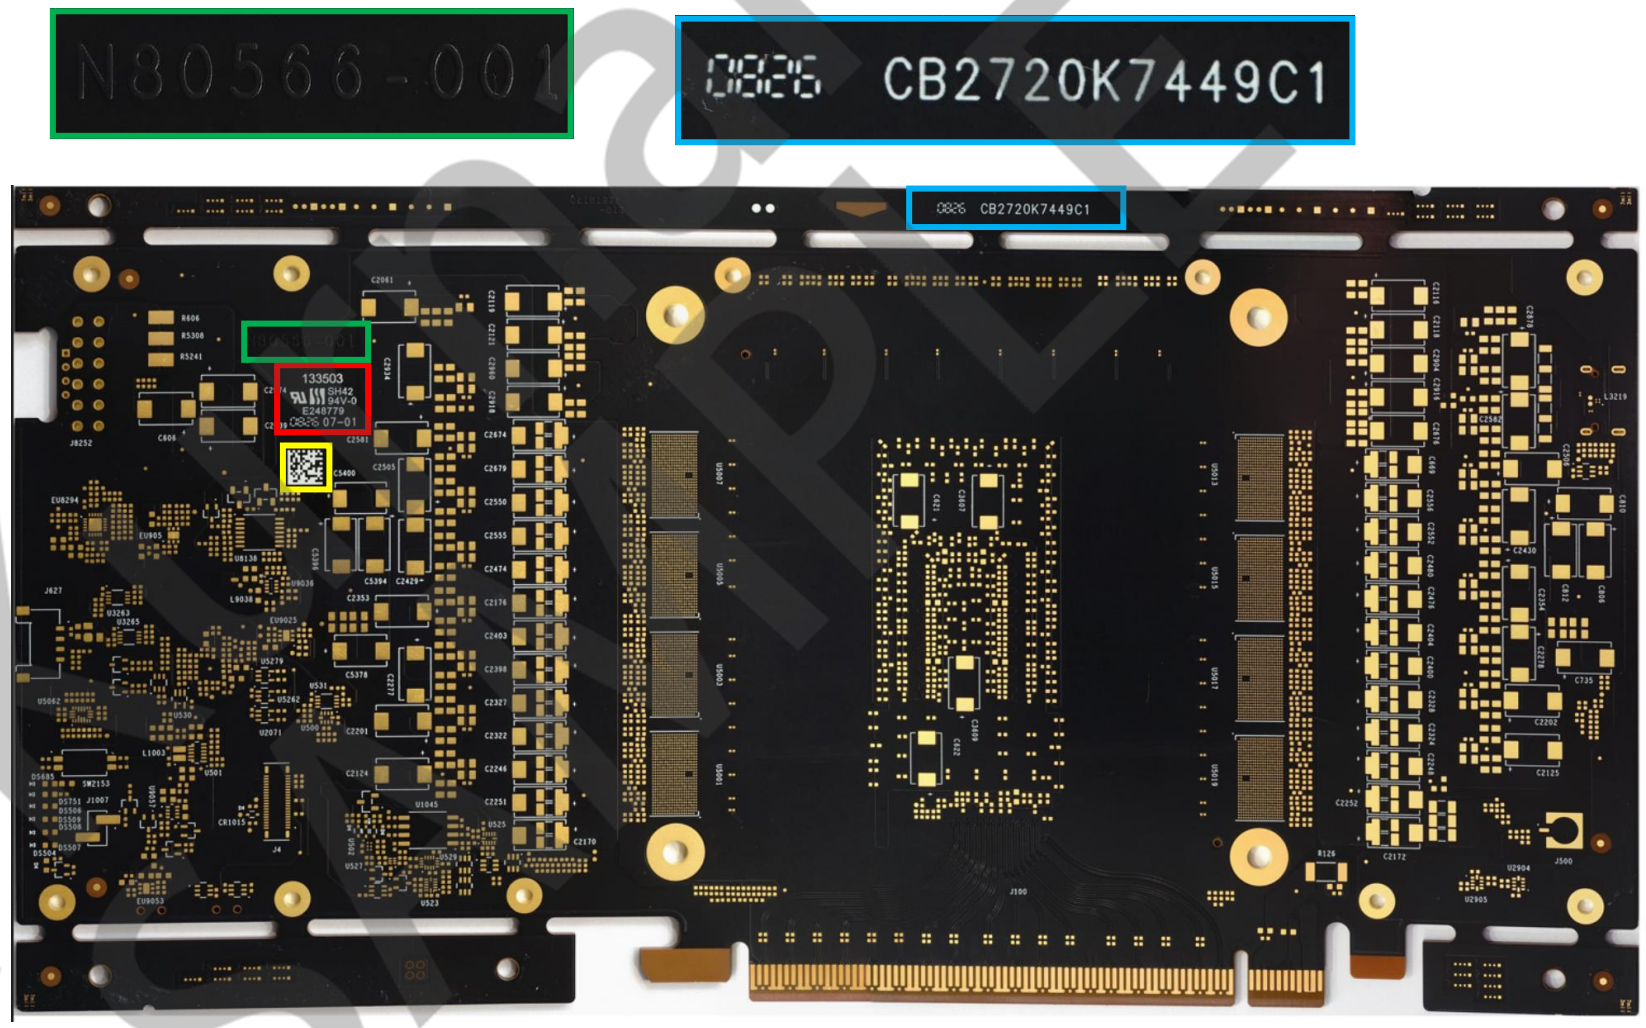

GPU-PCB Backview

PCB Mark Analysis

PCB Mark

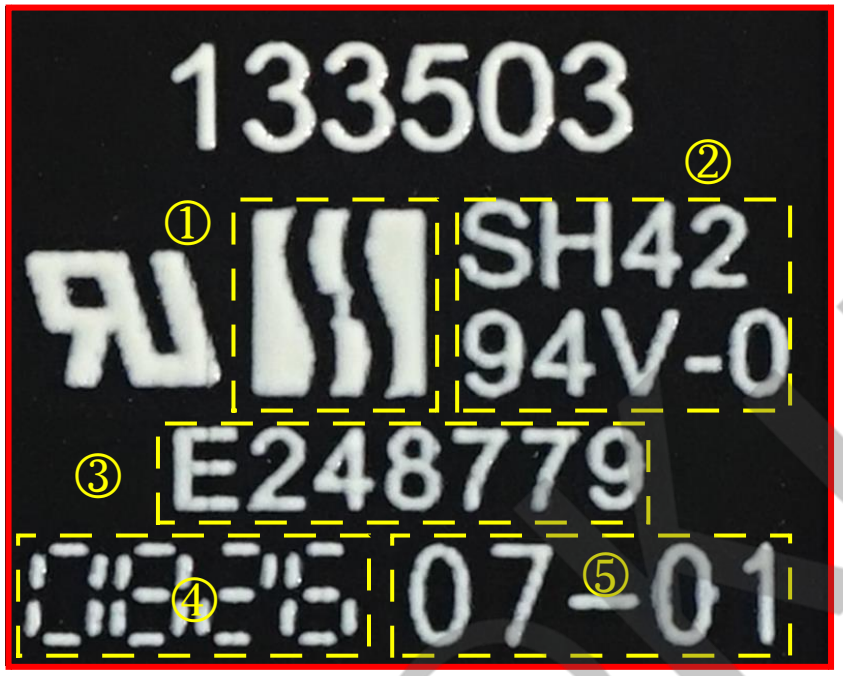

① Manufacturer Logo ② Material/Flammability Rating

③ UL-File Number ④ Data Code(WWYY)

⑤ Revision Code

PCB Information (Decoded)

Manufacturer	Victory Giant Technology (Huizhou) 胜宏科技(惠州)
PCB Material / Process	SH42 (ASP1) Assembly Solder Process I
Flammability Rating	94V-0
UL File Number	E248779
Production Date	2026 Year Week 08 (2026W08)

PCB Mark Analysis

PCB Mark

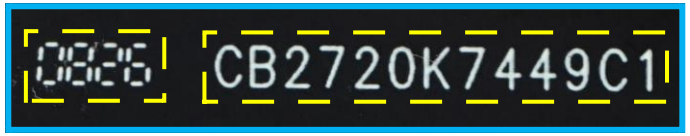

①　②

③

④

① Data Code(WWYY)

② Card-Part Number

③ Customer Part Number / Intel P/N

④ Data Matrix Code:V08261335030016

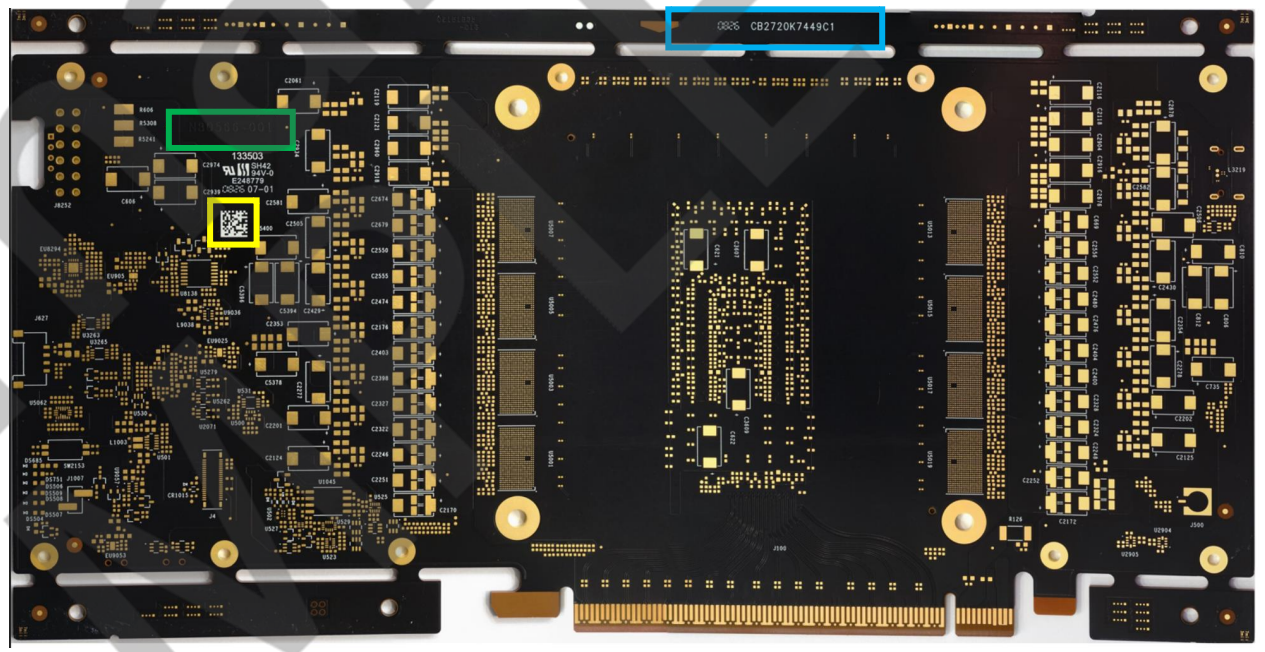

GPU-PCB Backview

Socket Analysis

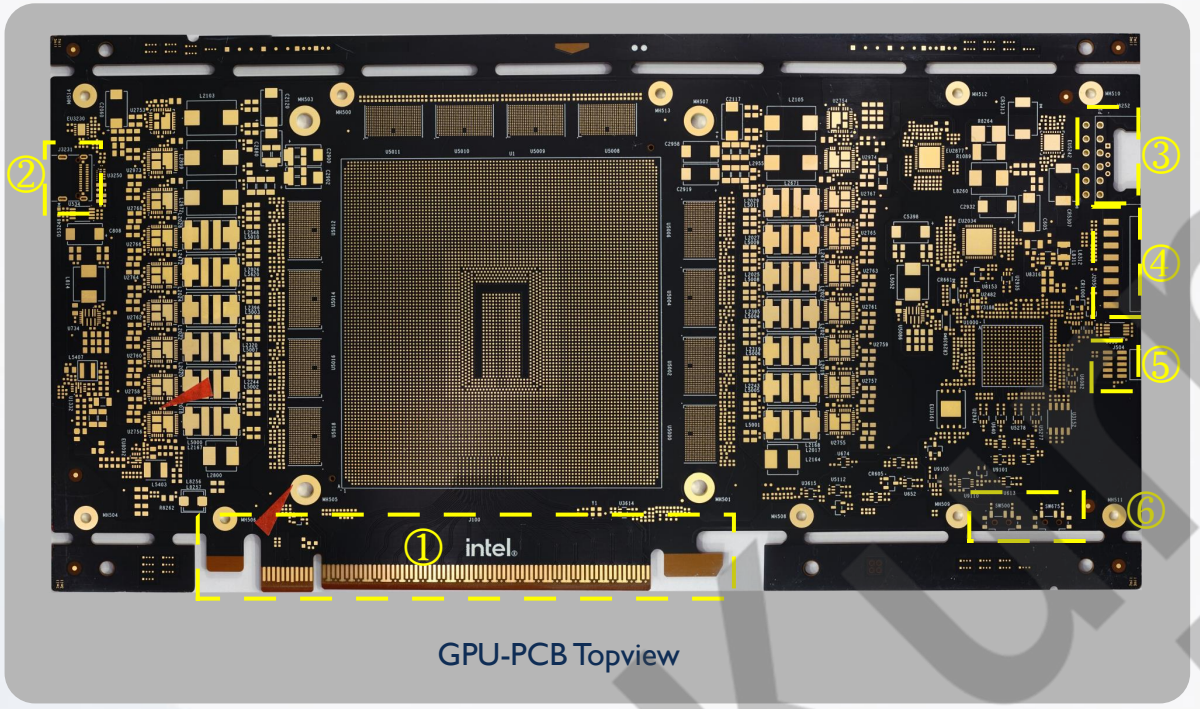

GPU-PCB Topview

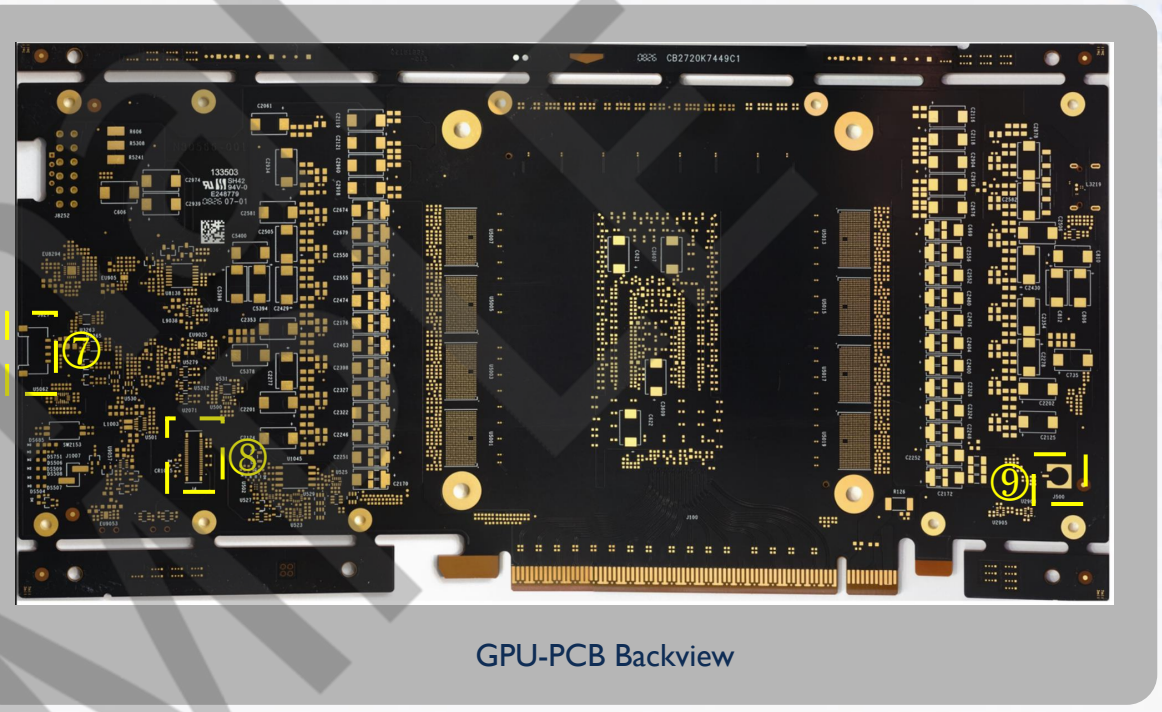

GPU-PCB Backview

① PCIE5.0 X16
② USB-C
③ 12VHPWR12V-2x6
④ Fan control?
⑤ Jumper set
⑥ SMD switch
⑦ PCIE5.0 X16
⑧ GPIO
⑨ Smp Socket

PCB Cut Layers Analysis

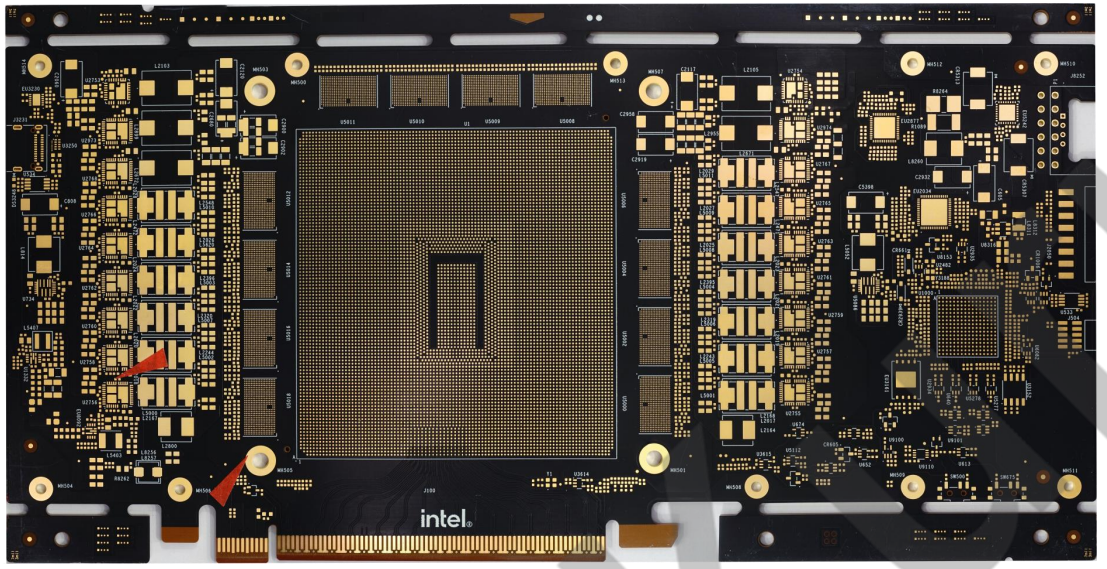

GPU-PCB Topview

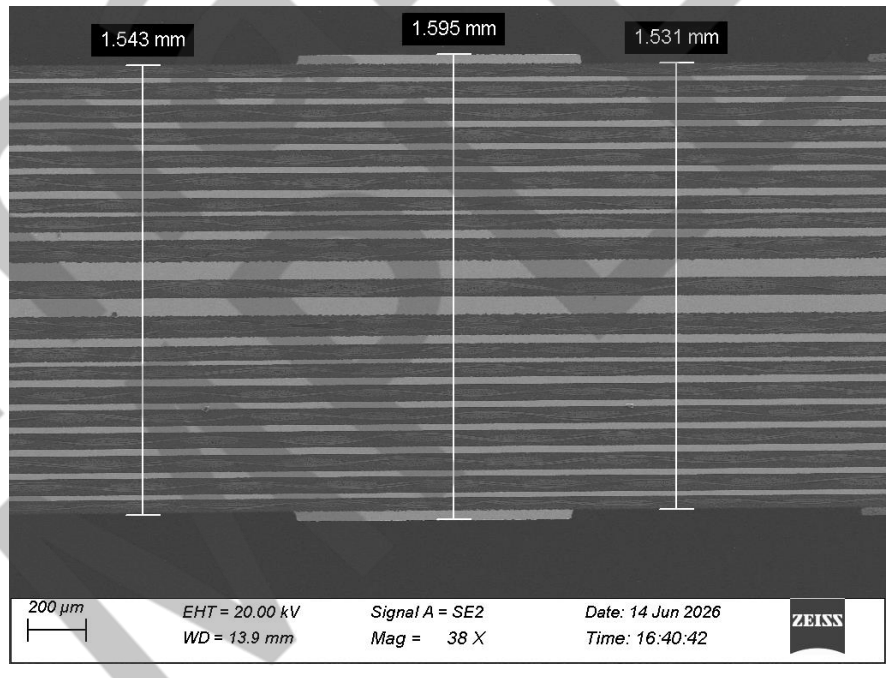

Layer Stack Summary

Total Layers
20 Layers
(10 + 10) Symmetrical)

PCB Thickness
1.6 mm
(±10%)

Via Type
PTH
(Plated Through Hole)

PCB Type	High Layer Count PCB
Board Size	120mm x 235mm
PCB Thickness	1.6 mm
Material	FR-4

PCB Stack-up Architecture Analysis

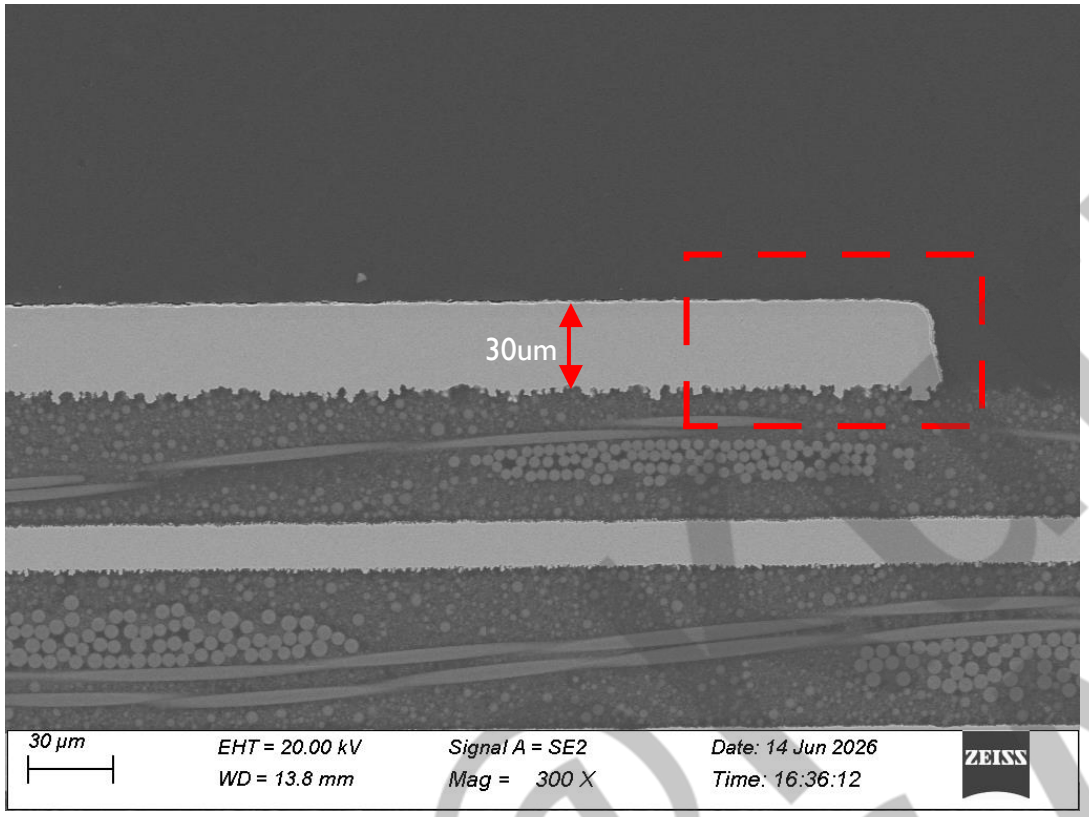

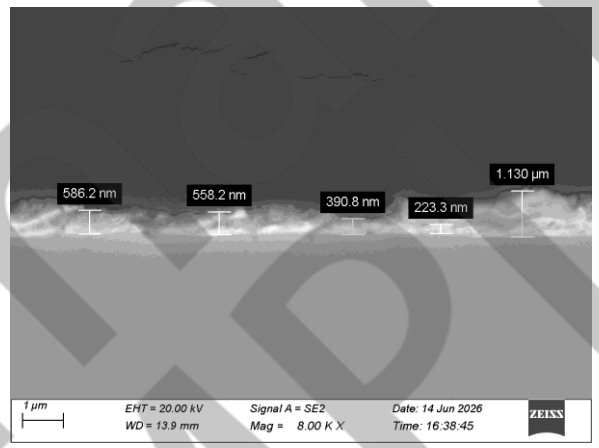

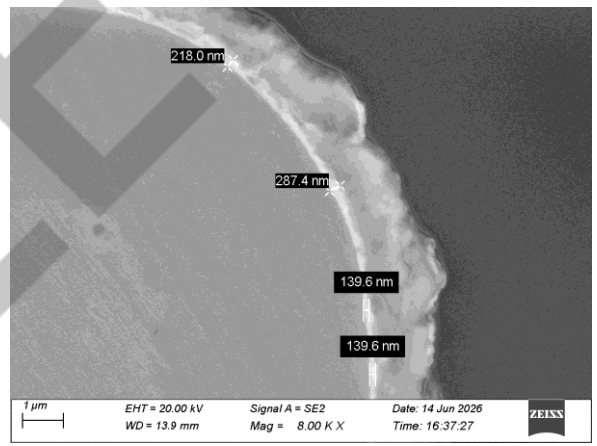

Core PCB Stack-up Findings

Copper Power Layer

- Thickness: **30 µm**
- Designed for GPU high transient current
- Lower IR drop under load spikes

ⓘ The top copper layer thickness is approximately **30 µm**, designed for high-current transient performance.

PCB Inter-layer Dielectric & Spacing Analysis

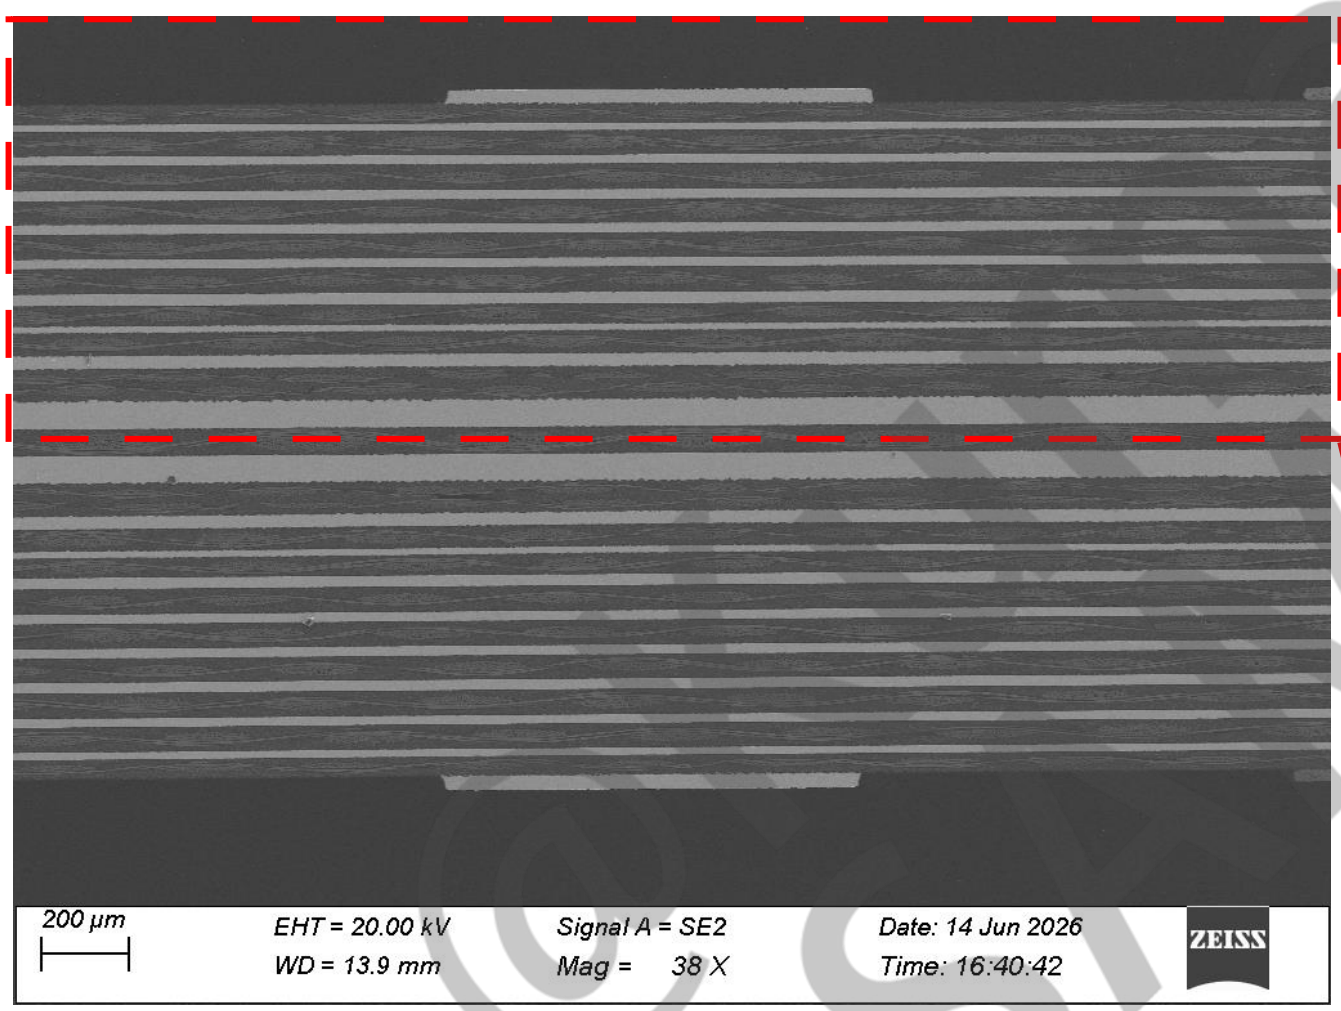

Layers	Cu Thickness	Layers Gap
Layer1	30um	47um
Layer2	16um	57um
Layer3	20um	60um
Layer4	20um	60um
Layer5	20um	60um
Layer6	20um	60um
Layer7	26um	47um
Layer8	16um	53um
Layer9	30um	76um
Layer10	60um	60um

Power supply design

PCB-analysis

Power supply design

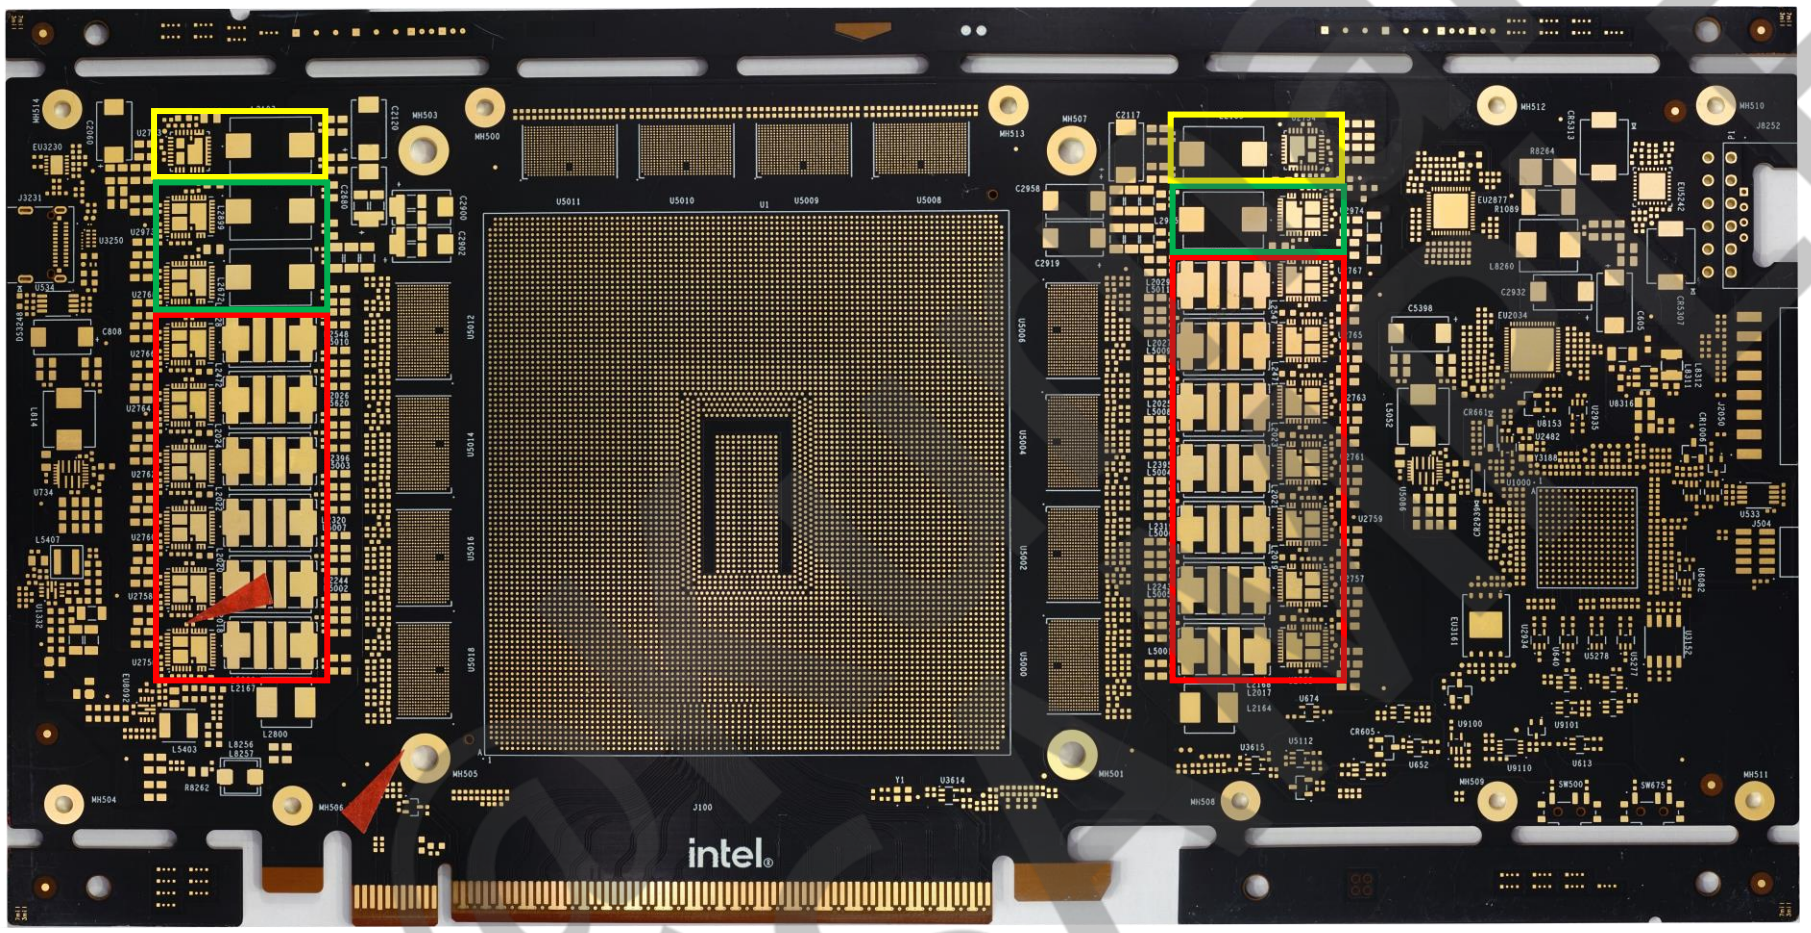

VRM Power Summary

Power Rail	Phase Count	Description
VDD_Core	13	GPU Core Power
VDD_Mem	3	LPDDR5X Memory Power
VDD_IOM	2	I/O & System Power
Total	**18**	Total Phases

Key Observations

- 13-phase VRM for GPU core power delivery
- 3-phase VRM for LPDDR5X memory
- 2-phase VRM for I/O and system power

Estimated total board power: **350W+**

GPU Analysis

Crescent island

GPU package size

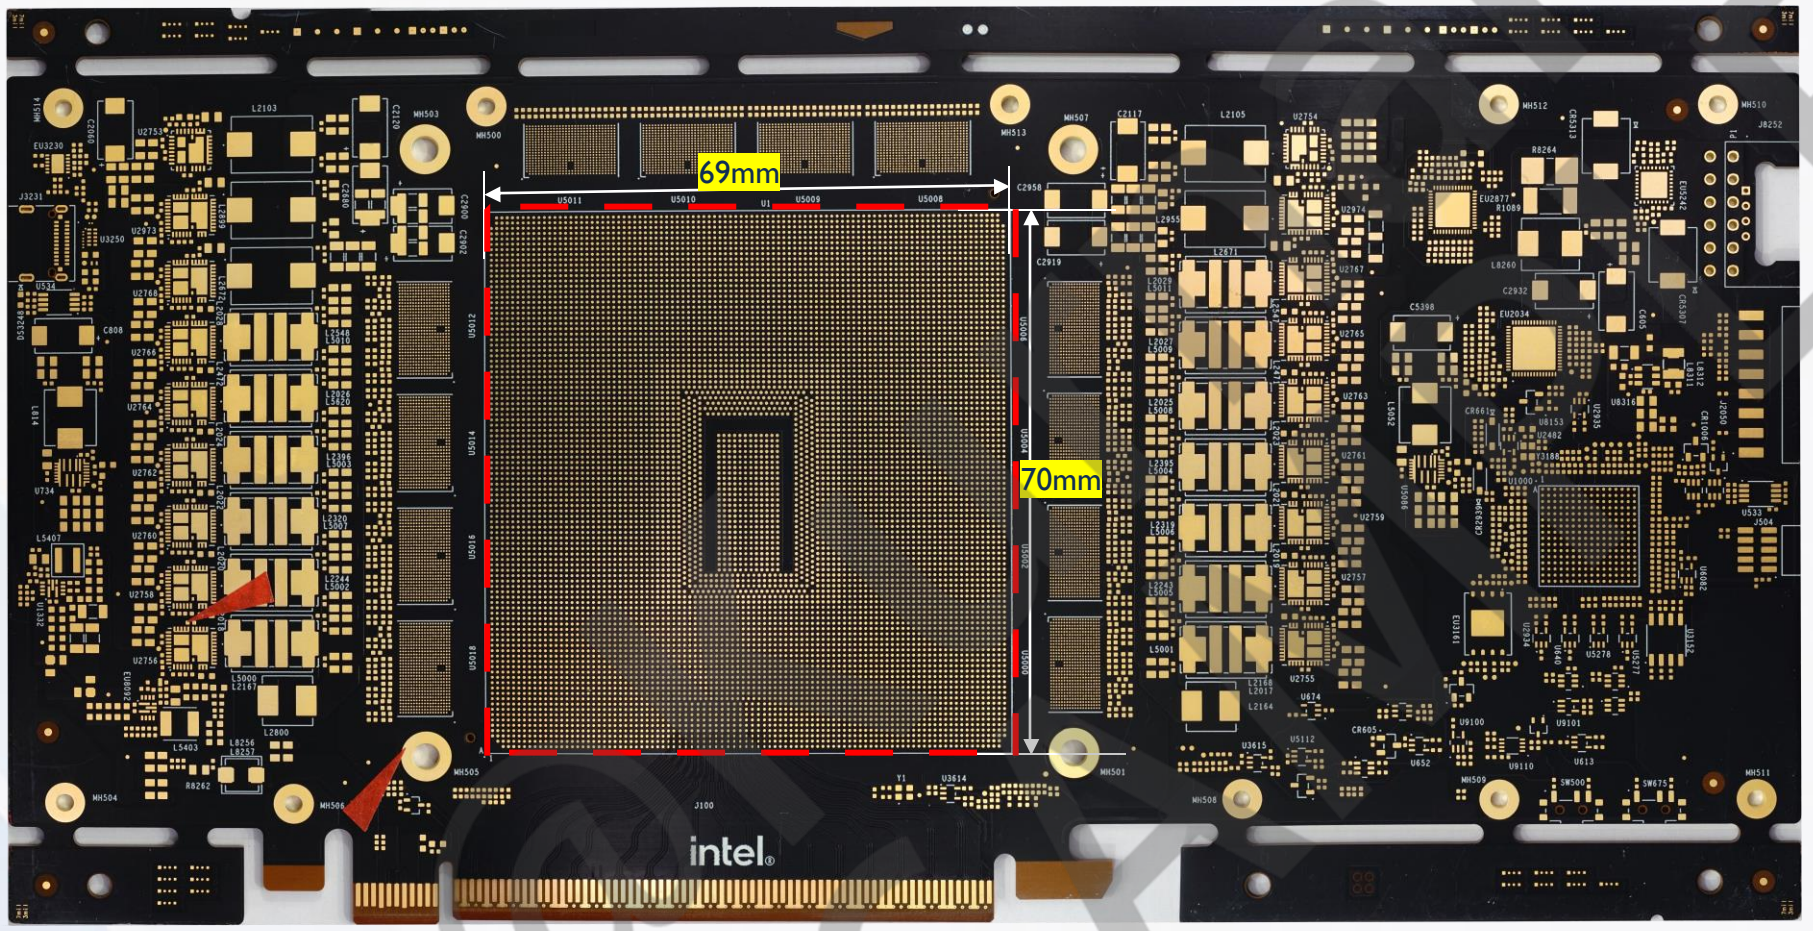

Package Size
69 mm × 70 mm
Length × Width

Pad Pitch
0.7 mm
Center-to-center

Pad Count
9,575
Total BGA Pads

Package Type
BGA
Flip Chip Package

BGA Pad Pitch Measurement

Measured pad pitch between centers

Key Specifications

Package Size	69 mm × 70 mm Length × Width
Pad Pitch	0.7 mm Center-to-center
Pad Count	9,575 Total BGA Pads
Package Type	BGA Flip Chip Package

Key Findings

- Large 69mm × 70mm BGA package design
- Fine-pitch 0.7mm ball array
- Total of 9,575 solder balls
- High-density interconnect substrate
- Comparable to modern high-end AI / GPU packages

SoC Die Size Estimation

Based on BGA Sparse Region Analysis

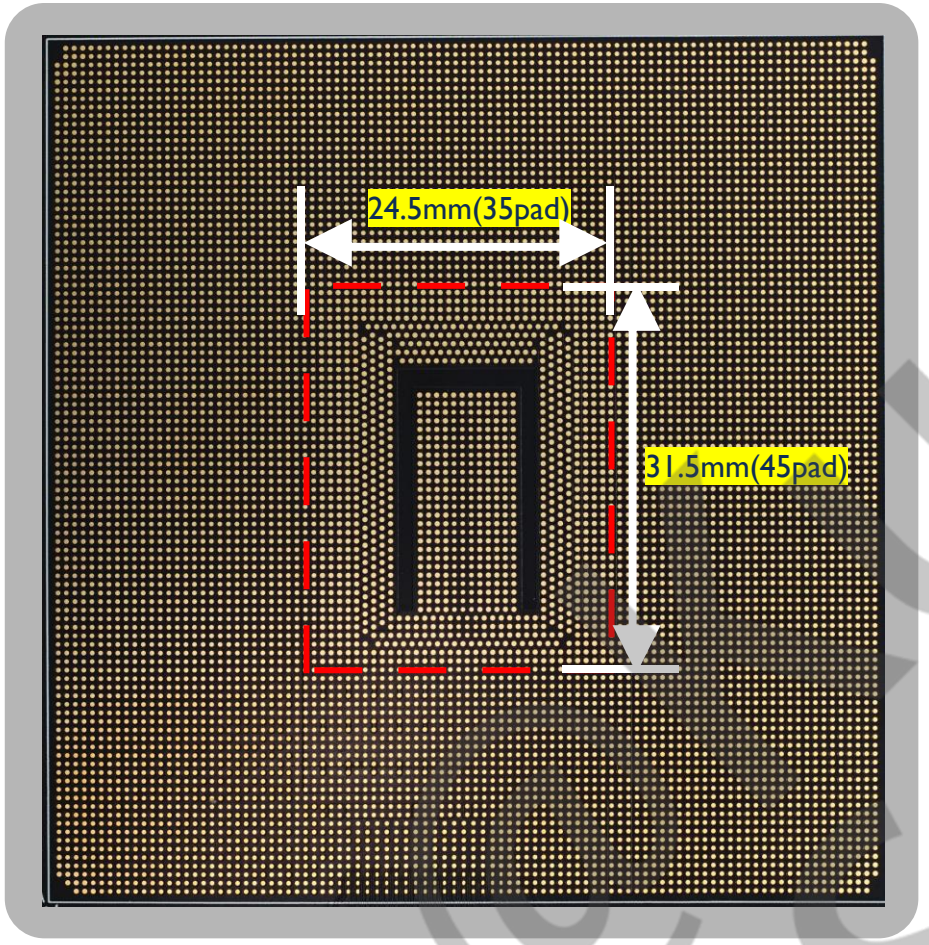

Sparse Pad Region Analysis

The central BGA array exhibits a low-density pad region. Assuming this sparse region corresponds to the GPU die placement area, the die footprint can be estimated from the measured dimensions.

Measured Dimensions

Parameter	Value
Width (Horizontal)	24.5 mm (35 pad)
Height (Vertical)	31.5 mm (45 pad)
Pad Pitch	0.7 mm
Total Pad Count	9,575
Package Size	69 mm (W) × 70 mm (H)

Estimated Die Size

24.5 mm × 31.5 mm

Estimated Die Area

≈ 772 mm²

Calculation: 24.5 × 31.5 ≈ 771.75 mm²

Die area is approximately 772 mm².

Key Findings

- Central sparse region likely corresponds to GPU die location
- Estimated die footprint: **24.5 mm × 31.5 mm**
- Estimated silicon area: **~772 mm²**
- Represents a large reticle-class GPU die
- Calculation based on BGA pad sparse region and pad pitch (0.7 mm)

Die Area Comparison

Device	Die Area (mm²)	Notes
NVIDIA AD102 (Ada)	616.03	Monolithic GPU
NVIDIA GB202 (Blackwell)	761.56	Monolithic GPU
NVIDIA H100 (Hopper)	824.686	Monolithic GPU
NVIDIA GB200 (Blackwell)	804 × 2	Two GPU dies (per package)
Intel Crescent Island (Est.)	**~772**	**Estimated**

The estimated die area is comparable to modern high-end GPU/AI processors.

DPW(Die-Per-Wafer) Analysis

Based on Estimated Die Size

1. Die Size (From Previous Analysis)

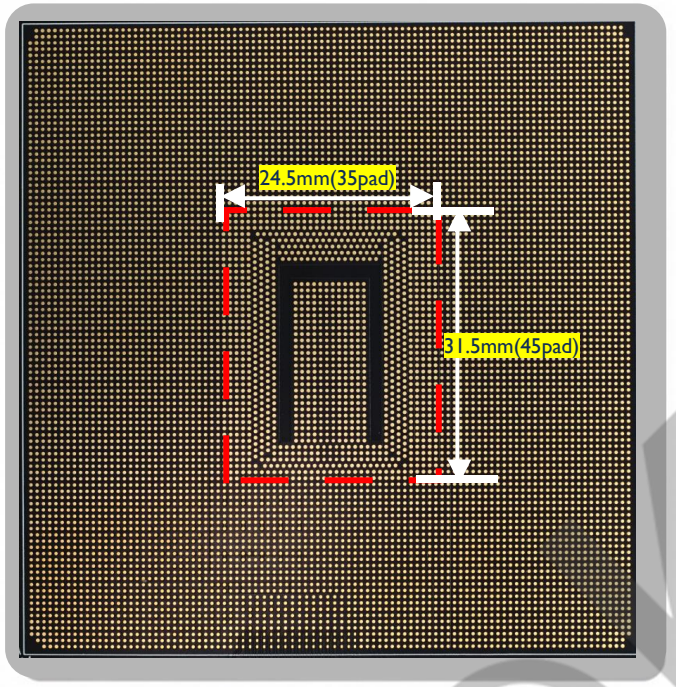

Estimated Die Size	
Width (Horizontal)	24.5 mm
Height (Vertical)	31.5 mm
Die Area	772 mm²

2. DPW Calculation Result

DPW
(Dies Per Wafer)

75
dies / wafer

3. Mask Information

Mask Type
Full Mask 1:1

Scaling
1:1 (No Scaling)

4. Utilization Metrics

Mask Utilization (MFU)	**89.947%**
Wafer Utilization (OWE)	**81.885%**
Mask on Wafer Efficiency (MOWE) MFU × OWE	**73.65%**

5. Wafer Map (Full Mask 1:1)

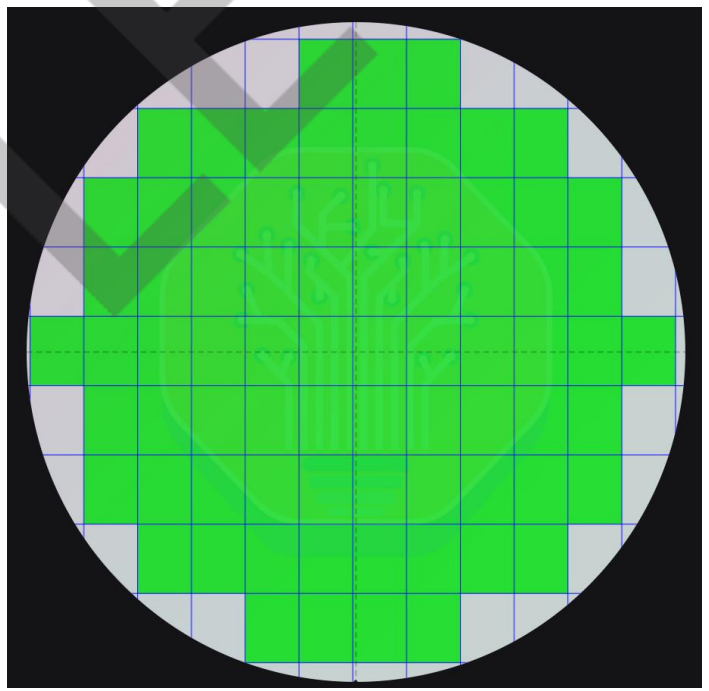

Summary

- Based on estimated die size of 24.5 mm × 31.5 mm (772 mm²)
- Full mask 1:1 is assumed for DPW calculation
- The resulting DPW is 75 dies per wafer
- Overall wafer efficiency (MOWE) is 73.65%

Memory Analysis

Crescent island

Lpddr5x Package placement analysis

GPU-PCB Topview

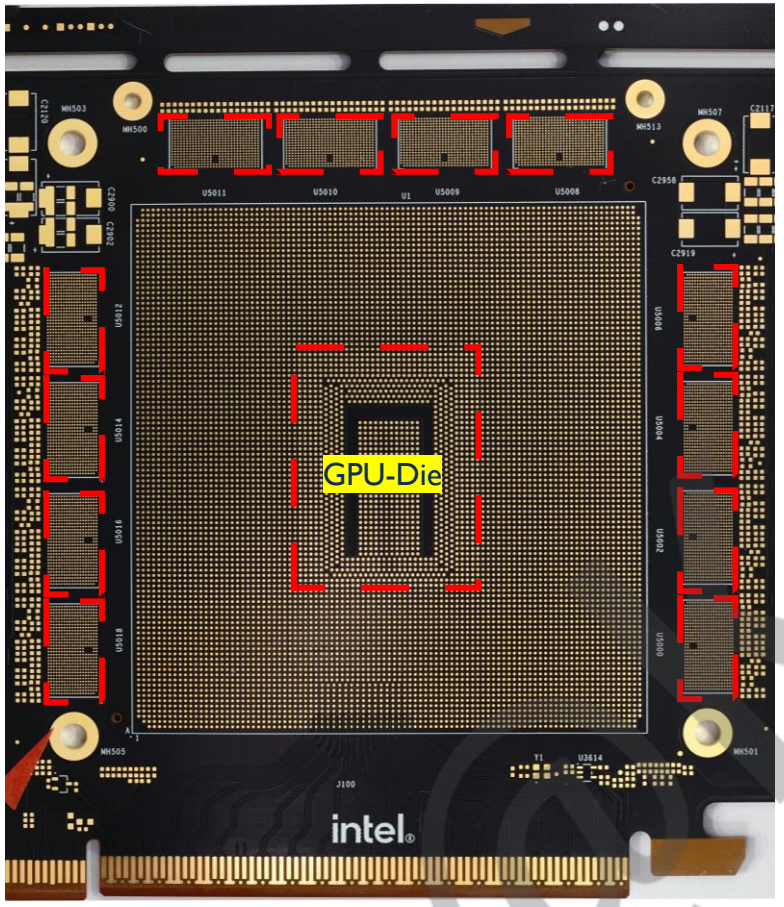

12 lpddr5x packages placed on Front side surrounding the GPU

GPU-PCB Backview

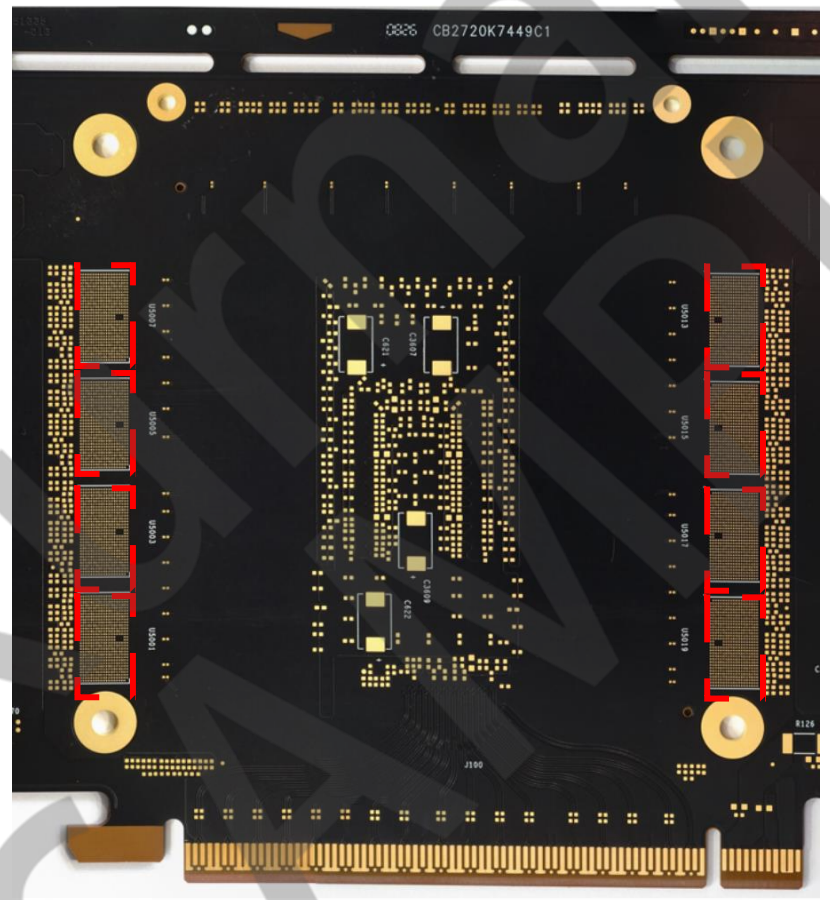

8 lpddr5x packages placed on Back side of the PCB

SUMMARY

MEMORY CONFIGURATION

Front Side	12 Packages
Back Side	8 Packages
Total Packages	**20 Packages**

KEY FINDINGS

- Dual-side LPDDR5X population
- Symmetrical package placement
- Short routing distance to GPU
- Optimized signal integrity

LPDDR5X BGA563 Package analysis

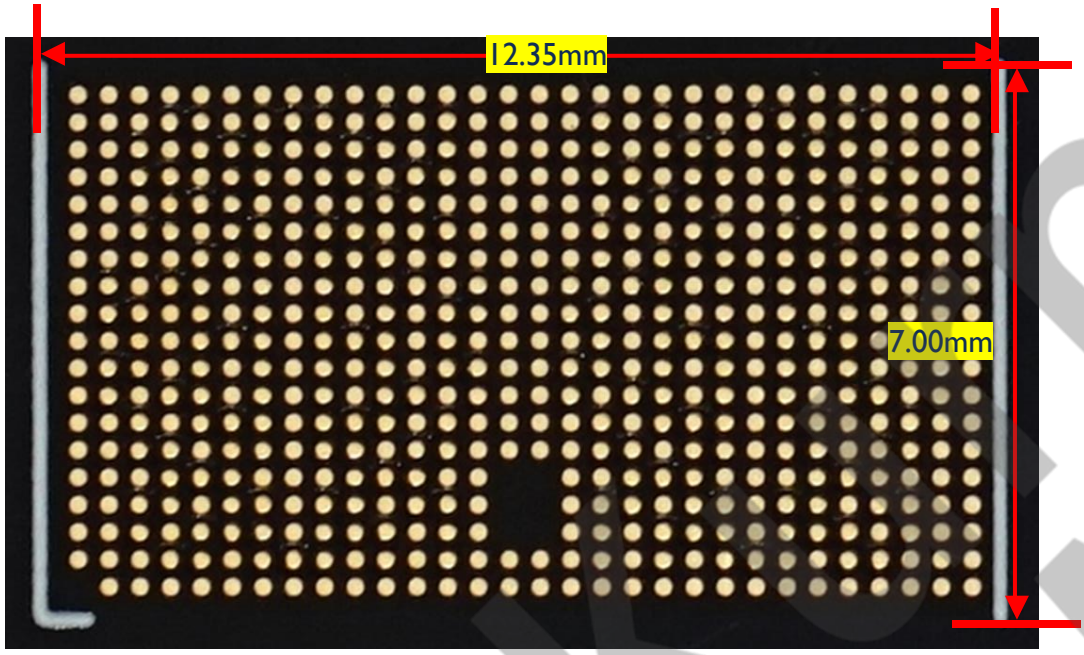

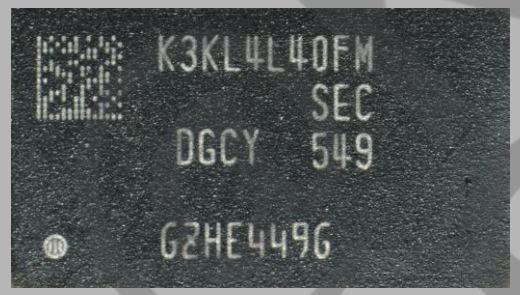

Common BGA563 Lpddr5x chips

PACKAGE SPECIFICATIONS	
Package Type	LPDDR5X
Package Code	BGA563
Ball Count	563
Ball Array	Matrix Layout
Pad Pitch	0.40 mm
Package Size (L x W)	7.00 mm x 12.35 mm
Package Thickness	0.90 mm (typ.)
Substrate Type	Organic
Ball Material	SAC (Sn-Ag-Cu)
RoHS Compliance	Yes

KEY FINDINGS

- BGA563 package identified
- Fine-pitch memory package with 0.40 mm pad pitch
- High-density package design with 563 balls
- Optimized for high bandwidth and power efficiency
- Suitable for AI accelerator and high-performance applications

Memory Capacity Estimation

480 GB
Total Capacity

1280 -bit
Bus Width

20
Memory Packages

563
Balls per Package

Technical Summary

Memory Type	LPDDR5X
Package Type	BGA563
Packages	20
Capacity	480 GB
Bus Width	1280-bit

Architecture Highlights

Ultra-wide memory subsystem
1280-bit bus width enables massive bandwidth for AI workloads

High-capacity memory design
480 GB total capacity supports large-scale model inference and training

Dual-side package placement
20 packages distributed on both sides of the module for optimal density

AI and HPC optimized architecture
Low-power LPDDR5X memory optimized for AI and HPC applications

Memory Bandwidth Estimation

BANDWIDTH CALCULATION

FORMULA Bandwidth = Bus Width × Data Rate ÷ 8

BUS WIDTH		DATA RATE (Effective)		BUS CONVERSION
1280-bit (20 Channels)	×	8533 MT/s (Base) UP to 9600 MT/s	÷	÷ 8

CALCULATED BANDWIDTH

At 8533 MT/s (Base)	At 9600 MT/s (Maximum)
1.36 TB/s (Base)	1.54 TB/s (Maximum)

MEMORY BANDWIDTH COMPARISON

Peak Theoretical Memory Bandwidth

Higher is Better

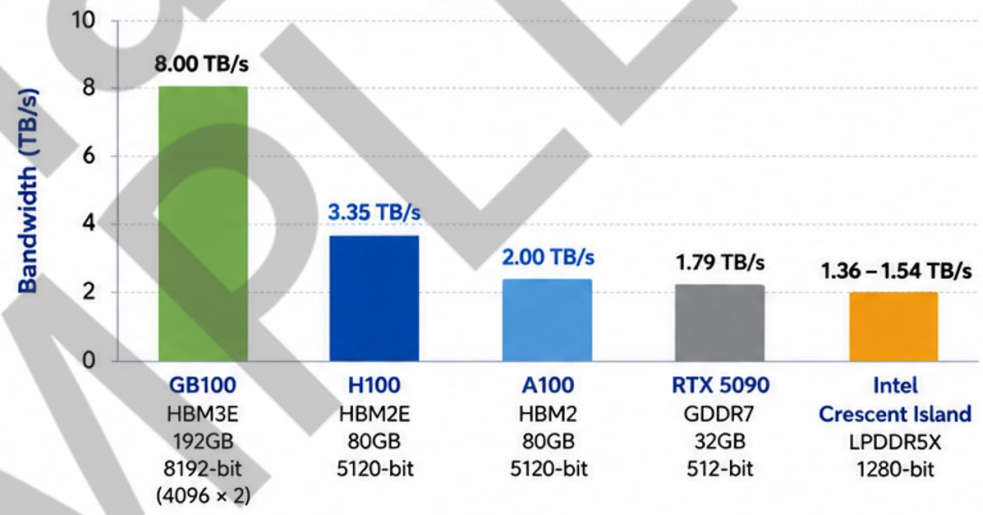

	GB100	H100	A100	RTX 5090	Intel Crescent Island
Memory Type	HBM3E	HBM2E	HBM2	GDDR7	LPDDR5X
Capacity	192GB	80GB	80GB	32GB	480GB
Interface Width	8192-bit (4096 × 2)	5120-bit	5120-bit	512-bit	1280-bit
Bandwidth (TB/s)	8.00	3.35	2.00	1.79	1.36 – 1.54

KEY FINDINGS

- High-bandwidth LPDDR5X subsystem delivers up to 1.54 TB/s peak theoretical bandwidth
- Competitive bandwidth for power-efficient AI inference
- Optimized for edge and data center inference workloads