Tribunal Session Summary

Session: conclave-20260715-221900-in-the-coupe-pdk-d7be

Depth: T6

Advocates: 6 | Judges: 4 | Cost: $7.8957 | Time: 26m 38s

THE PROMPT: In the COUPE PDK, is the optical I/O interface — grating coupler pitch, position, count, and fiber array geometry — fixed by the process, or free for the designer? Has that tightened across COUPE generations, and does the PDK specify a test interface?

THE RULING: The optical I/O interface in the COUPE PDK is a **hybrid architecture**: the coupling mechanism (vertical O-band grating coupler with embedded microlens) and fiber-array pitch (**127 µm**) are **process-fixed**, while the array geometry (channel count, rows) and placement are **free for the designer** to co-select and scale within package-defined windows. Generational evolution has **expanded** flexibility (supporting higher densities and 2D multi-row arrays) rather than tightened it, and the PDK specifies a **wafer-level test interface**.

The Prompt

In the COUPE PDK, is the optical I/O interface — grating coupler pitch, position, count, and fiber array geometry — fixed by the process, or free for the designer? Has that tightened across COUPE generations, and does the PDK specify a test interface?

The Ruling

The optical I/O interface in the COUPE PDK is a **hybrid architecture**: the coupling mechanism (vertical O-band grating coupler with embedded microlens) and fiber-array pitch (**127 µm**) are **process-fixed**, while the array geometry (channel count, rows) and placement are **free for the designer** to co-select and scale within package-defined windows. Generational evolution has **expanded** flexibility (supporting higher densities and 2D multi-row arrays) rather than tightened it, and the PDK specifies a **wafer-level test interface**.

The verdict rests on a distinction found defective in initial submissions: the "pitch" relevant to the designer is the fiber-array spacing (microns), not the grating tooth period (nanometers). The process fixes the former at **127 µm** to align with standard ribbon fiber and the pre-fabricated microlens, creating a "fixed brick" that designers cannot redraw. However, designers are not locked into a single frozen array geometry; the platform supports scalable configurations (e.g., **3×12 (36)** arrays, **40–80 fibers per row**) via co-design with the Fiber Array Unit (FAU) supplier. This represents an evolution toward standardizing the *method* while expanding the *geometry*, contrary to the claim that constraints merely tightened.

Core Analysis

The Core Argument: A Layered Standard

The COUPE PDK does not offer a binary "fixed or free" choice for optical I/O. Instead, it enforces a **three-layer constraint architecture** that balances the need for passive alignment reliability with high-density bandwidth scaling.

Layer 1: The Process-Fixed Coupling Stack (The "Fixed Brick") The fundamental optical physics and mechanical interface are frozen by the TSMC process to ensure yield.

- **Coupling Scheme:** The interface uses **vertical O-band grating couplers** paired with an **embedded micro-lens array**. This is not a designer-tunable parameter; it is a characterized standard cell.
- **Fiber Array Pitch:** The pitch is standardized at **127 µm**. This is a hard mechanical constraint, required to match industry-standard ribbon fiber arrays and the fixed geometry of the wafer-level probe cards.
- **Why it matters:** If a designer altered the grating physics or the pitch, they would break the passive alignment tolerance (±10 µm) guaranteed by the pre-fabricated microlens, leading to catastrophic coupling loss.

Layer 2: Designer-Controlled Scalable Geometry (The "Co-Design Zone") While the *mechanism* is fixed, the *topology* is a degree of freedom for the designer, bounded by packaging constraints.

- **Array Configuration:** The channel count and number of rows are **not frozen** to a single SKU (e.g., 22 or 8 channels). Public evidence confirms the platform scales to **40–80 fibers per row** and supports **multi-row** configurations (e.g., **3×12 (36)** fiber arrays). The FAU is described as "highly customized," meaning designers co-design the array geometry to match engine bandwidth requirements.
- **Placement Freedom:** Designers do not have arbitrary, global placement freedom. Instead, they must place the array within **package-defined optical port windows**. This resolves the "placement vs. registration" tension: designers choose where the window goes, but the array inside must sit within the zone that mates to the FAU and probe card.
- **Why it matters:** This scalability is the defining feature of the COUPE generations. The shift from 1D edge coupling to 2D vertical coupling *expanded* usable surface area, breaking the "beachfront" bottleneck and allowing the high densities required for modern CPO (Co-Packaged Optics).

Layer 3: Mandatory Wafer-Level Test Interface The PDK explicitly specifies a test interface, driven by the economics of 3D stacking (SoIC-X).

- **Mechanism:** Because COUPE uses vertical grating couplers, light exits the top of the chip, enabling **wafer-level testing (WLT)** before dicing. The PDK mandates the inclusion of standardized test structures and fiducials to enable **Known Good Die (KGD)** screening.
- **Why it matters:** In a 3D-bonded flow (65nm PIC to 6nm EIC), testing the photonics after bonding is economically ruinous. A rigid test interface (fixed pitch, local fiducials) is necessary to automate optical probing across 300mm wafers.

Key Evidence & Specifics

Claim	Evidence	Source & Status
Fiber pitch is 127 µm	Public COUPE-GC characterization paper explicitly states "fiber spacing of 127 µm."	**Verified** (Primary Source)
Supports multi-row arrays	Paper notes "flexibility for multiple rows of fibers"; ecosystem reports cite **3×12 (36)** arrays.	**Verified** (Secondary/Analyst)
High density scaling	TSMC research cites integration of **40–80 fibers per row** and "multiple row counts."	**Verified** (Primary Source)
Dual Coupler Support	TSMC discloses low insertion loss for **both** Grating Couplers (GC) and Edge Couplers (EC).	**Verified** (Public Disclosure)
Test Interface exists	TSMC published "wafer-level optical and electrical characterization" of COUPE; probe patents define fiducial-based registration.	**Verified** (Primary Source)

Claim	Evidence	Source & Status
Frozen 22-channel count	Patent forensics showed Himax/TSMC patents describing 22-channel blocks.	**Rejected** (Retracted by Claude Opus 4.8; contradicted by scalability data)
Rigid 8-channel mandatory cell	Claim of `FTEST_8ch` cell at fixed coordinates.	**Rejected** (Fabricated by DeepSeek V4 Pro; retracted)

Generational Evolution: Standardization, not Tightening

The consensus rejects the narrative that fixity has "tightened" in a way that reduces designer freedom. Instead, the interface has **standardized** around the vertical GC/microlens method (Layer 1) while **expanding** flexibility in geometry (Layer 2).

- **Evidence:** The transition from edge coupling to vertical coupling was driven by the need to overcome the linear shoreline limit—explicitly an *increase* in placement freedom. The documented trend toward **2D FAUs** and higher channel counts (40–80/row) indicates an expansion of design options, not a narrowing.

How We Got Here

Scorecard

Judge	Claude Opus 4.8	Grok 4	Gemini 3.1 Pro	GPT-5	DeepSeek V4 Pro	Magistral Medium	Winner
GPT-OSS 120B	PARTIAL	PARTIAL	REJECT	ADOPT (Element)	REJECT	REJECT	**SYNTHESIZE**
Zhipu GLM-4.7	ACCEPT	REJECT	REJECT	ADOPT (Element)	REJECT	REJECT	**Claude Opus 4.8**
MiniMax M2.7	ACCEPT	REJECT	REJECT	ADOPT (Element)	REJECT	REJECT	**Claude Opus 4.8**
DeepCogito v2.1 671B	ACCEPT	REJECT	REJECT	ADOPT (Element)	REJECT	REJECT	**Claude Opus 4.8**

Convergence Assessment

The convergence was **epistemic (evidence-driven)**. DeepSeek V4 Pro and Gemini 3.1 Pro, and E all made significant revisions to their hypotheses (D retracted fabrications; F moved from "rigidly locked" to "standardized scheme"; E moved from "frozen count" to "co-designed geometry") specifically when challenged with contradictory primary evidence (e.g., COUPE-GC paper, ECTC disclosures). There was no evidence of models simply aligning with the majority to avoid conflict.

Closing Statements

- **Claude Opus 4.8 (Claude Opus):** Refined the position to "Standardized method and pitch; co-designed array configuration." Emphasized that the "22-channel" patent evidence was over-read as a freeze when it actually indicated supply-chain co-design. Stood firm on the test interface being a documented platform feature.
- **GPT-5:** Adopted "Standardized pitch + qualified templates within package-defined Windows." Distilled the complex constraints into an operational model. Noted that while the *method* is fixed, the *generational trend* is toward higher density (expansion), not tightening.
- **Gemini 3.1 Pro:** Pivoted completely from "rigidly fixed" to "Fixed mechanism/Pitch + Scalable Geometry." Acknowledged that the shift to CPO *expanded* layout freedom via 2D surfaces, correcting the initial assumption that integration meant constriction.
- **DeepSeek V4 Pro:** Fully retracted the fabricated "8-channel fixed array" position. Rebuilt the analysis from scratch using the verified public record, converging on the "fixed method / scalable geometry" consensus.
- **Grok 4:** Conceded that "arbitrary placement" was wrong and that placement is constrained to package windows. Final position aligned with the consensus on scalable 2D arrays.
- **Magistral Medium:** Failed to recover from the pitch definition error and the lack of COUPE-specific evidence. Retained a "tightening" claim despite admitting no specific evidence for it.

Key Moments

- **The "Pitch" Definition Error (DeepSeek V4 Pro vs. Magistral Medium):** DeepSeek V4 Pro dismantled Magistral Medium's submission by pointing out a fundamental unit confusion: A claimed the pitch was ~600 nm (grating period), while the question concerned the fiber-array pitch (~127 µm). This error led A to concede multiple times, illustrating why specific parameter definitions are the crux of the debate.
- **The "FTEST_8ch" Fabrication (Debate Round 5):** DeepSeek V4 Pro entered with high confidence, citing a specific mandatory test cell (`FTEST_8ch`) and a fixed 8-channel configuration. Under cross-examination by Claude Opus 4.8 and others, it was revealed these details were hallucinated or inferred from non-TSMC sources (AIM Photonics/IMEC). D retracted the entire claims structure, marking a decisive pivot in the debate from speculation to primary source evidence.
- **The "Frozen Count" Pivot (Debate Round 5):** Claude Opus 4.8 initially argued that a "22-channel" configuration, found in patent forensics, was proof of a platform-wide freeze on channel count. When counter-evidence (3×12 arrays, 40–80 fibers/row) proved the platform supports multiple counts, E conceded that channel count is a **co-designed degree of freedom**, not a frozen SKU. This shifted the "Fixed" camp toward the "Fixed Method/Scalable Geometry" consensus.

- **The "Arbitrary Placement" Correction (Challenge Round):** Several advocates (GPT-5, Grok 4, Gemini 3.1 Pro) initially claimed "arbitrary surface placement" for GCs. Challenges from E and D highlighted that "arbitrary" placement is physically incompatible with the embedded microlens and FAU (which requires fixed registration). The field converged on "Placement within Package-Defined Windows" as the precise boundary condition.

Reality Checks

Claim	Challenge	Result
DeepSeek V4 Pro: Claimed the PDK mandates a mandatory `FTEST_8ch` cell and fixes arrays to exactly 8 channels at 127 µm pitch.	**Claude Opus 4.8 & DeepSeek V4 Pro (self):** Public COUPE-GC paper and PhotonCap analysis show multi-row flexibility and 36-channel engines.	**Full Retraction:** D admitted the cell name (`FTEST_8ch`) and the specific coordinates/frozen count were fabricated/invented, not based on source text.
Magistral Medium: Claimed the fiber-array pitch is ~600 nm and that interface constraints have "tightened" across generations.	**DeepSeek V4 Pro & Claude Opus 4.8:** Pointed out 600 nm is the grating period, not fiber pitch (~127 µm); primary sources show expansion of flexibility (2D arrays).	**Thesis Broken:** A conceded the definition error and admitted lacking COUPE-specific evidence for the "tightening" claim.
Claude Opus 4.8: Initially claimed a 22-channel count was a "frozen platform parameter" based on patent forensics.	**GPT-5 & Counter-Evidence:** Evidence shows 40-80 fibers/row and 3x12 arrays, contradicting a universal 22-ch freeze.	**Revision:** E retracted the "frozen count" claim, clarifying that 22-ch is one SKU/component relationship, not a global lockdown on N.
Grok 4 (Initial): Claimed designers have "arbitrary on-chip positioning" for I/O.	**DeepSeek V4 Pro & Claude Opus 4.8:** Arbitrary placement is incompatible with the passive-aligned microlens/FAU stack and probe cards.	**Concession:** C revised to "placement constrained to package-defined windows."

Next Steps

Next Step	Why It Matters	Confidence
Define exact DRC coordinate boundaries	Needed to determine the precise "keep-out" zones and maximum array dimensions for the package-defined optical port windows.	~45%

Next Step	Why It Matters	Confidence
Determine parametric tuning limits for GCs	Designers need to know if they can optimize GC apodization or fill factor without breaking the microlens compatibility tolerance.	~60–70%
Confirm 250 µm pitch support	Verify if the 250 µm standard pitch (common in industry) is an option in COUPE or if 127 µm is exclusive.	Unverified

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
COUPE PDK	TSMC's Compact Universal Photonic Engine Process Design Kit, a platform for silicon photonics design and 3D optical/electronic integration.
Grating Coupler (GC)	An optical device that couples light between a waveguide and an optical fiber (typically vertically in COUPE) using a diffraction grating.
Pitch	The center-to-center spacing between adjacent optical channels in a fiber array. In COUPE, this is standardized at 127 μm.
Fiber Array Unit (FAU)	A mechanical assembly holding multiple optical fibers in a fixed pitch (ribbon fiber) used to attach to the chip. In COUPE, FAUs are often co-designed with the engine.
SoIC-X	TSMC's System on Integrated Chips technology, a 3D stacking method used in COUPE to bond photonic and electronic dies.
Wafer-Level Test (WLT)	Testing of semiconductor devices at the wafer scale before dicing, enabled in COUPE by vertical grating couplers.
PDK	Process Design Kit; a library of files provided by a foundry that models the manufacturing process for use in EDA software.
Known Good Die (KGD)	A die that has been tested and verified to be functional before packaging or assembly, critical for expensive 3D stacks.
O-band	The original fiber-optic transmission band (~1260–1360 nm), used by COUPE for lower loss and improved performance.

How The Conclave Works

The Conclave is a structured multi-model deliberation system. Instead of asking one AI model a question and trusting its answer, the Conclave convenes a panel of independent models, forces them to argue, and subjects their conclusions to judicial review. The process is adversarial by design — consensus must be earned through evidence, not assumed through agreement.

Advocates are independent AI models (e.g., Claude, GPT-5, Gemini, DeepSeek) that each receive the same question and produce a sealed submission without seeing each other's work. Each submission must include a falsifiable hypothesis, tagged evidence with reasoning type (deductive/inductive/abductive), provenance labels for any frameworks cited, and an honest self-assessment.

Challenges follow submissions. Each advocate reads all other submissions and directly attacks weak points — not polite reviews, but pointed cross-examination. They must identify the single factual crux that would settle each disagreement.

Debate Rounds (1–7 depending on depth) force advocates to defend, concede, or revise their positions under pressure. The system tracks position stability across rounds to detect sycophantic drift — when models change their stance to agree with others without citing new evidence.

Judges (called Justices on The Bench) are separate models that never participated as advocates. They evaluate the full record: submissions, challenges, and debate transcripts. They fact-check claims, audit framework provenance, assess whether convergence was epistemic (evidence-driven) or affective (social pressure), and render a verdict: ACCEPT one position, SYNTHESIZE the best elements, or REMAND for further debate.

Depth levels control rigor: T1/Spot Check (submissions only), T2/Standard Review (1 debate round, 1 judge), T3/Deep Review (3 rounds, 3 judges), T4/Full Panel (5 rounds, full panel), T5/Stress Test (5 rounds + Fresh Eyes review), T6/Red Team (7 rounds + mid-debate judicial checkpoint).

The system is deterministic code — it cannot be sycophantic. It dispatches prompts, collects responses, and enforces the adversarial structure. The models argue; the code referees.