

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

APPLE INC.,

Petitioner,

v.

WECREVENTION, INC.,

Patent Owner.

Patent No. 9,164,942

Filing Date: October 11, 2012

Issue Date: October 20, 2015

Inventors: Weng-Dah Ken and Nicky Lu

Title: HIGH SPEED MEMORY CHIP MODULE AND ELECTRONICS
SYSTEM DEVICE WITH A HIGH SPEED MEMORY CHIP MODULE

PATENT OWNER'S PRELIMINARY RESPONSE

Case No. IPR2026-00239

TABLE OF CONTENTS

	<u>Page(s)</u>
I. INTRODUCTION	1
II. CLAIM CONSTRUCTION	6
III. LEVEL OF ORDINARY SKILL IN THE ART	6
IV. PETITIONER HAS NOT DEMONSTRATED A REASONABLE LIKELIHOOD OF SUCCESS, AND THE PETITION SHOULD BE DENIED	7
A. The Petition Should Be Denied Because Petitioner Takes Inconsistent Claim-Construction Positions Across Forums Without Explanation	9
1. Apple’s Inconsistent Positions Across the Two Forums	9
2. The Unexplained Inconsistency Warrants Denial Under Revvo, Tesla, and Cambridge	11
3. The Inconsistency Has Substantive Consequences for the Petition	12
B. The Petition Independently Warrants Discretionary Denial Under 35 U.S.C. § 325(d) Because the Examiner Specifically Considered Lee and Concluded the Prior Art Did Not Teach the “Converting” Limitation.....	14
C. Grounds 1A and 1I Fail Because Apple’s [19a] Mapping Improperly Fills the “ASIC Processor” Gap in Lee with Expert Testimony and Non-Asserted References, and Because Lee’s “Generating Memory Requests” Does Not Satisfy [19e]’s “Corresponding To” Requirement	15
1. Apple’s [19a] Mapping Improperly Fills the “ASIC Processor” Gap in Lee with Expert Testimony and Non- Asserted References	16

2.	Lee’s “Generating Memory Requests” Does Not Satisfy [19e]’s “Corresponding To” Requirement.....	18
D.	Grounds 1A and 1I Independently Fail Because Apple’s “Converting” Mapping Cross-Cites Nakamura—a Reference Not Asserted in Any Ground—and Because Lee Discloses the Coupling of 128-Bit Parallel Data, Not Conversion	19
1.	Lee Describes Coupling of 128 Bits Through the High-Speed Link, Not Conversion.....	20
2.	Apple Fills the Gap in Lee with Nakamura, a Reference Not Asserted in Any Ground	21
3.	Apple’s Bus-System-68 Alternative Theory Does Not Cure the Defect	22
E.	Grounds 1B–1S (the Lee-Tsern Combinations) Independently Fail Because Inserting Tsern’s Buffer into Lee Destroys Lee’s Principle of Operation	23
F.	The Dependent-Claim Grounds Relying on Lin (Grounds 1F, 1O) and White (Grounds 1H, 1Q, 1S) Independently Fail Because Lin and White Disclose Single-Chip Structures, Not the Per-Memory-Cell-Array-IC Seal Rings and TSVs the Dependent Claims Require.....	25
1.	Lin Discloses a Single-Chip Seal-Ring/TSV Structure, Not Per-Memory-Cell-Array-IC Seal Rings.....	26
2.	White Likewise Discloses TSVs in a Single Die, Not Per-Memory-Cell-Array-IC TSVs.....	27
3.	The Cumulative-Reference Grounds (1P–1S) Compound the Same Defects.....	28
V.	CONCLUSION.....	29

LIST OF EXHIBITS

Exhibit No.	Description of Document
2001	Plaintiff's Disclosure of Asserted Claims and Infringement Contentions in <i>WeCrevention, Inc. v. Apple Inc.</i> , Case No. 7:25-cv-00458-ADA (W.D. Tex.), dated January 20, 2026
2002	Defendant's Invalidity and Subject Matter Eligibility Contentions in <i>WeCrevention Inc. v. Apple Inc.</i> , Case No. 7:25-cv-00458-ADA (W.D. Tex.), dated March 17, 2026
2003	Federal Register, Vol. 90, No. 199, dated Friday, October 17, 2025
2004	Scheduling Order (Dkt. 32) in <i>WeCrevention, Inc. v. Apple Inc.</i> , Case No. 7:25-cv-00458-ADA (W.D. Tex.), dated January 21, 2026
2005	https://marketrealist.com/2019/12/will-china-still-be-second-home-apple/
2006	Appeal Statistics, Patent Trial and Appeal Board, dated October 31, 2025
2007	Patent Trial and Appeals Board (PTAB) data at a glance December 2025
2008	March 24, 2025 guidance memorandum from Chief Administrative Patent Judge Boalick
2009	Director Squires Memorandum, Additional Discretionary Institution Considerations — U.S. Manufacturing and Small Business Use of AIA Proceedings (Mar. 11, 2026)
2010	MacRumors / Financial Times, Apple Plans to Source All U.S.-Bound iPhones from India by End of 2026, Apr. 25, 2025
2011	Korea Economic Daily / MacRumors, Samsung Becomes Apple's Primary LPDDR5X Supplier for iPhone 17, Providing 60–70% of Memory, Dec. 22, 2025
2012	SK hynix Newsroom, <i>Semiconductor 101: SK hynix on Where Chips Are Made & Used</i> (Apr. 3, 2025)
2013	Apple Newsroom, Apple Increases U.S. Commitment to \$600 Billion, Announces American Manufacturing Program, Aug. 6, 2025, https://www.apple.com/newsroom/2025/08/apple-increases-us-commitment-to-600-billion-usd-announces-ambitious-program/

Exhibit No.	Description of Document
2014	Samsung Semiconductor Newsroom, <i>Samsung Begins Mass Production of 16Gb LPDDR5 DRAM at World's Largest Semiconductor Line</i> (Sept. 9, 2020)
2015	SK hynix Inc., <i>SK hynix Signs Investment Agreement of Advanced Chip Packaging with Indiana</i> (Apr. 3, 2024)
2016	Apple's Opening Claim Construction Brief in <i>WeCrevention, Inc. v. Apple Inc.</i> , Case No. 7:25-cv-00458-ADA (W.D. Tex.) Dkt. 40, dated May 6, 2026

I. INTRODUCTION

Apple's Petition for *inter partes* review of the '942 Patent should be denied. It fails for reasons that emerge from Apple's own filings, from the prior art Apple itself selected, and from the prosecution history of the '942 Patent. Each of the independent defects identified below is, on its own, sufficient to defeat institution; taken together, they show a Petition that did not, and on the references Apple chose could not, demonstrate a reasonable likelihood of success on any Challenged Claim.

First, Apple's claim-construction posture is inconsistent across forums. In the parallel district-court litigation,¹ Apple has expressly proposed that the preambles of the '942 Patent's claims are limiting because the preamble phrase "*high speed memory chip module*" supplies "a fundamental characteristic of the claimed invention" that is "necessary to give life, meaning, and vitality" to the claims. Ex. 2016 at 4, 9. In this Petition, however, Apple takes the opposite position by footnote: "Petitioner does not concede that the preambles of the Challenged Claims are limiting." Pet. 9 n.5. Apple offers no explanation for the inconsistency. Under *Revvo Techs., Inc. v. Cerebrum Sensor Techs., Inc.*, IPR2025-00632, Paper 20 at 5 (Director Nov. 3, 2025) (precedential), and *Tesla, Inc. v. Intell. Ventures II LLC*, IPR2025-00340, Paper 18 (Director Nov. 5, 2025), an unexplained inconsistency on whether the preamble is limiting warrants denial of institution.

¹ *WeCrevention, Inc. v. Apple Inc.*, Case No. 7-25-cv-00458 (W.D. Tex.)

Second, the '942 Patent's prosecution history independently warrants denial under § 325(d). Lee (US 2004/0260864) was specifically before the Examiner: Lee was listed in Applicant's Information Disclosure Statement, Ex. 1002, 19–20, and was independently cited by the Examiner in the Notice of References Cited (PTO-892), *id.*, 17 (Row B: "US-2004/0260864 A1 ... Lee et al."). The Examiner expressly confirmed that "ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /H.B./" *Id.*, 19–20. On a record that included Lee, the Examiner allowed the Challenged Claims, *id.*, 14, and stated: "there is no teaching, suggestion, or motivation for combination in the prior art to a high speed memory including where the logic unit is further used for converting the first set of parallel data of the first transmission bus into a second set of parallel data through a second transmission bus." *Id.*, 15. Apple's Petition asks the Board to reach the opposite conclusion on the same reference without acknowledging—much less identifying error in—the Examiner's prior consideration. Thus, denial is warranted. *Advanced Bionics, LLC v. MED-EL Elektromedizinische Geräte GmbH*, IPR2019-01469, Paper 6 at 8 (P.T.A.B. Feb. 13, 2020) (*precedential*).

Third, Apple's mapping for two of Claim 19's central limitations—the "*ASIC processor*" element ([19a]) and the "*executes a predetermined function corresponding to the second set of parallel data*" element ([19e])—does not stand on Lee alone, and so neither Ground 1A nor the Grounds the build off Ground 1A

can be sustained on Apple’s record. For [19a], Apple’s expert testifies that “[a] POSITA would have understood that a North Bridge controller is an ASIC processor,” then cites Reeves (Ex. 1010) and Bruno (Ex. 1011)—neither of which is expressly relied on in Ground 1A—for the proposition that Lee’s controller “is generally considered an . . . application specific integrated circuit (ASIC).” Pet. 11 (quoting Ex. 1010, 5:30–33; citing Ex. 1011, [0032]; Ex. 1003, ¶40). Framing Ground 1A in that manner runs afoul of 37 CFR § 104(b)(4) in at least two ways. First, the expert-testimony characterization that supplies the missing “ASIC processor” limitation is the precise use of general knowledge the Stewart Memorandum prohibits: “expert testimony . . . may not be used to supply a missing claim limitation.” Stewart Memorandum, July 31, 2025, at 2. Second, the Reeves/Bruno citations Apple’s expert offers to support that characterization are prior-art references the Petition should have included as an obviousness combination for Ground 1A but did not—turning what Apple labels a single-reference §103 ground into, in substance, a multi-reference §103 combination disguised as a single-reference ground. Separately, for [19e], Apple equates Lee’s controller “generat[ing] memory requests” with the claim’s requirement that the ASIC processor execute a function “*corresponding to*” the second set of parallel data. But Lee’s controller generates a request *before* the read; the second set of parallel data is the *result* of that request, not the input to it. Pet. 28. Lee’s controller therefore does nothing

“*corresponding to*” the second set of parallel data within any natural reading of Lee and Claim 19. The same two defects carry into Ground 1I (for Claims 5, and 9–10) and the eleven Grounds that build off it (Grounds 1J–1S), because Apple’s Ground 1I mapping at [5b] expressly directs the Board to “*See [19a] and [19e] of Ground 1A.*” Pet. 90.

Fourth, Apple’s mapping for the “*converting*” limitation is defective on the same two axes, and the defect reaches every Ground in the Petition because the “*converting*” language appears identically in both independent Claim 19 ([19d]) and independent Claim 1 ([1c]). Lee, on its face, does not describe “*converting*” one set of parallel data into another set of parallel data. Lee describes coupling 128 bits of data through the high-speed link 58 either by “increasing the speed of the link 58” or by varying “the size of the data word coupled through the link 58.” Ex. 1004, [0017]; Pet. 7–8. That language describes how the same 128-bit payload travels onto the link, not a conversion of one set of data into a different set. Faced with that gap in Lee, Apple does precisely what the Stewart Memorandum forbids: Apple’s expert simply characterizes Lee’s coupling as a “*convers[ion]*” by citing Nakamura (Ex. 1019)—a reference asserted nowhere in the Petition—for the proposition that “a parallel-to-parallel converter for converting an ‘m’-bit parallel signal into an ‘n’-bit parallel signal” was known in the art. Pet. 25–27 (citing Ex. 1003, ¶¶ 66–68 (citing Ex. 1019, Abstract)). Apple’s Ground 1I mapping for Claim 1 element [1c] expressly

directs the Board to “*See* [19d] of Grounds 1A-1B,” Pet. 90, so the [19d] defect carries directly into Ground 1I and the eleven Grounds that build off it. Every Ground in the Petition therefore depends, by Apple’s own cross-references, on a [19d]/[1c] mapping that violates 37 CFR § 42.104(b)(4).

Fifth, Apple’s Lee-Tsern combinations (Grounds 1B–1S) inherit Ground 1A’s defects and independently fail because inserting Tsern’s serialization/deserialization buffer into Lee’s architecture destroys Lee’s principle of operation. Lee’s defining feature is its high-bandwidth mode in which “the memory hub 60 addresses all of the memory devices 70 simultaneously so that 128 bits of data are written to or read from the memory devices 70 each memory access.” Ex. 1004, [0017]. Tsern’s buffer device, by contrast, “performs serialization and deserialization functions” depending on “memory device access width” to bridge mismatched interface widths between memory devices and a single buffer interface. Ex. 1005, 31, 38. Substituting Tsern’s serialization-based bandwidth-concentration scheme for Lee’s simultaneous parallel access either eliminates Lee’s parallel-128-bit operation (defeating Lee’s purpose) or requires re-architecting Tsern’s buffer in a way Tsern does not contemplate. *See In re Gordon*, 733 F.2d 900, 902 (Fed. Cir. 1984). The motivation that Apple does articulate is conclusory *KSR* boilerplate (“known technique applied to known system to yield predictable results”) that does not grapple with the architectural conflict. Pet. 41–43.

Any one of the foregoing—the unexplained claim-construction inconsistency, the disguised multi-reference reliance on Reeves, Bruno, and Nakamura in violation of § 104(b)(4) and the Stewart Memorandum, the textual mismatch between Lee’s generating of memory requests and the claimed “*executing a predetermined function corresponding to the second set of parallel data*,” the § 325(d) basis for denial under *Advanced Bionics*, and the destruction of Lee’s principle of operation in the Lee-Tsern Grounds—is independently a sufficient basis to deny institution. Together, they leave Apple far short of the reasonable-likelihood threshold. The Petition should be denied.

II. CLAIM CONSTRUCTION

Patent Owner agrees with Apple’s position in the Petition that all terms should be given their plain and ordinary meaning. Pet. at 4-5.

III. LEVEL OF ORDINARY SKILL IN THE ART

For purposes of this Preliminary Response only, Patent Owner applies Petitioner’s definition of a person of ordinary skill in the art (POSITA). Pet. 4. Patent Owner reserves the right to advance an alternative definition in any subsequent filing should the Board institute review.

**IV. PETITIONER HAS NOT DEMONSTRATED A REASONABLE
LIKELIHOOD OF SUCCESS, AND THE PETITION SHOULD
BE DENIED**

Obviousness rests on factual determinations including the scope and content of the prior art, the differences between the claimed subject matter and the prior art, and the level of ordinary skill. *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1, 17–18 (1966); *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question is whether the claimed invention as a whole would have been obvious. *Stratoflex, Inc. v. Aeroquip Corp.*, 713 F.2d 1530, 1537 (Fed. Cir. 1983). Failure to identify the differences between the claimed subject matter and the prior art is fatal to an obviousness challenge. *See Apple, Inc. v. Contentguard Holdings, Inc.*, IPR2015-00355, Paper 9 at 9–10 (P.T.A.B. June 26, 2015).

The Board must support the conclusion that a POSITA not only could have made, but *would have* been motivated to combine the prior-art references in the way the claim recites and reasonably expected success. *Pers. Web Techs., LLC v. Apple, Inc.*, 848 F.3d 987, 994 (Fed. Cir. 2017); *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1368 (Fed. Cir. 2016). The combinability determination requires a plausible articulated rationale. *Broadcom Corp. v. Emulex Corp.*, 732 F.3d 1325, 1335 (Fed. Cir. 2013). “Knowledge of a problem and motivation to solve it ***are entirely different from*** motivation to combine particular references.” *Innogenetics, N.V. v. Abbott Lab’ys*, 512 F.3d 1363, 1373 (Fed. Cir.

2008) (emphasis added). Where the proposed combination would render the primary reference unsatisfactory for its intended purpose, or change its principle of operation, there is no motivation to combine. *In re Gordon*, 733 F.2d at 902.

35 U.S.C. § 312(a)(3) and 37 C.F.R. § 42.104(b) require the petition to identify “*with particularity*” each claim challenged, the grounds for the challenge, and how each claim element is found in the prior art **relied upon for that ground**. The Board may deny petitions that fail to comply. *See* July 31, 2025 Memorandum from Then-Acting Director Stewart Regarding Enforcement and Non-Waiver of 37 C.F.R. § 42.104(b)(4).

Independently, the Board may exercise its discretion to deny institution under 35 U.S.C. § 314(a). *Cuozzo Speed Techs., LLC v. Lee*, 579 U.S. 261, 273 (2016) (“the agency’s decision to deny a petition is a matter committed to the Patent Office’s discretion”). The Director has held in a precedential decision that a petitioner who advances different claim-construction positions in district court and at the Board “**should explain sufficiently why the different positions are warranted, even in instances that do not implicate Section 112(f),**” and that failure to do so warrants denial. *Revvo Techs.*, IPR2025-00632, Paper 20 at 5 (emphasis added).

A. The Petition Should Be Denied Because Petitioner Takes Inconsistent Claim-Construction Positions Across Forums Without Explanation

After filing this Petition, Apple filed an Opening Claim Construction Brief in the parallel district-court litigation, *WeCrevention, Inc. v. Apple Inc.*, Case No. 7:25-cv-00458-ADA (W.D. Tex.) (Ex. 2016), taking the position that the preambles of the independent claims of the '942 Patent are limiting. Here, however, Apple expressly does not concede that the preambles of the Challenged Claims are limiting. The disparity is unexplained, and under the Director's guidance, that defect alone warrants denial of institution.

1. Apple's Inconsistent Positions Across the Two Forums

Apple's district-court position. Apple's Opening Claim Construction Brief expressly proposes that "*high speed memory chip module*" term of the preambles of the '942 Patent is limiting. Ex. 2016 at 6. Apple maintains in district court that the preamble "contains the core of what the inventors actually invented," that the preamble "supplies fundamental structure required by the body of the claims," and that "the preamble here therefore breathes 'life, meaning, and vitality' into the claim because no other limitations in the claims capture the modular structure imparted by '[reconfigurable] high speed memory chip module.'" *Id.* at 7, 9, 12. Apple goes so far as to argue in district court that **without** treating the preamble as limiting, the claims "could map to disparate and remote components: a 'memory cell array group'

in a data center, a ‘logic unit’ in a remote device, and the ‘buses’ being network links.” *Id.* at 12.

Apple’s IPR position. Apple’s Petition takes the opposite position by footnote: **“Petitioner does not concede that the preambles of the Challenged Claims are limiting.”** Pet. 9 n.5 (emphasis added). The reservation is not incidental. For example, Apple’s mapping of the preamble of Claim 19 in Ground 1A—“*A electronics system device with a high speed memory chip module, the device comprising*”—directs the Board to Lee’s **computer system 100** and Lee’s **memory system 50** (Ex. 1004, [0023]; Pet. 9–11) without ever showing that Lee specifically discloses a “*high speed memory chip module*” in the sense Apple urges is required in district court. Apple’s mapping equates Lee’s memory module 54 with the claimed module simply because “it includes memory chips” (Pet. 10), and characterizes the module as “high speed” because it “is coupled to the controller ‘through a high-speed link.’” Pet. 11 (citing Ex. 1004, [0015], [0017]). That is precisely the kind of “disparate and remote components” mapping that Apple’s district-court brief argues is inappropriate, should the preambles be deemed limiting in district court.

Apple’s positions are irreconcilable, as Apple cannot simultaneously be right in both forums. Either the preamble carries the structural weight Apple ascribes to it in district court (in which case Apple’s IPR mapping ignores limitations the claims

contain), or the preamble does not carry that weight (in which case Apple’s district-court construction is wrong). Apple offers no explanation for the inconsistency. The hedge “Petitioner does not concede” is itself an admission that Apple knew, when it filed the Petition, that it was taking a position inconsistent with one Apple intended to advance in district court.

2. The Unexplained Inconsistency Warrants Denial Under *Revvo*, *Tesla*, and *Cambridge*

An unexplained inconsistency between a petitioner’s claim-construction positions in district court and at the Board warrants denial of institution. In *Revvo Techs., Inc.*, for example, the Director held: **“the petitioner should explain sufficiently why the different [claim construction] positions are warranted.”** See IPR2025-00632, Paper 20 at 4–5 (emphasis added). On remand, the Board again granted institution, and the Director subsequently *sua sponte* initiated review and again denied institution. *Id.*, Papers 30 (Jan. 14, 2026) and 36 (Director Jan. 26, 2026).

The Director applied *Revvo* in *Tesla, Inc.*, IPR2025-00340. The *Tesla* petitioner argued indefiniteness in district court but said in the IPR that no claim term required express construction. The Director rejected the petitioner’s explanation: “Petitioner’s statement that it cannot raise indefiniteness challenges in an *inter partes* review is not a sufficient explanation. Indeed, the statement amounts

to an assertion that a petitioner should be permitted to raise inconsistent invalidity challenges in the two forums.” *Id.*, Paper 18 at 3. Applying *Revvo*, the Director was express in denying institution in *Tesla*: “[a]llowing a petitioner to advance a claim construction before the Board when that petitioner has made inconsistent indefiniteness arguments in district court fails to further, but instead detracts from, the Office’s goal of ‘**providing greater predictability and certainty in the patent system.**’” *Id.* at 4 (emphasis added) (citing *Revvo*, IPR2025-00632, Paper 20 at 4–5; see also *Cambridge Mobile Telematics, Inc. v. Sfara, Inc.*, IPR2024-00952, Paper 12 at 8 (P.T.A.B. Dec. 13, 2024) (“[a]t a minimum, the Petition should have explained why the plain and ordinary meaning . . . should apply in this proceeding (why the inconsistent positions are warranted) or set forth an alternative . . . claim construction.”)).

Apple’s footnote 5 is not an explanation; it is a flag. It is a candid admission that Apple is taking inconsistent positions and intends to maintain both. *Revvo* requires that Apple explain *why* the different positions are warranted—but Apple has merely reserved the right to maintain them. That is insufficient. *Tesla*, IPR2025-00340, Paper 18 at 3.

3. The Inconsistency Has Substantive Consequences for the Petition

The inconsistency is not a technicality. Apple’s district-court position—that

the preamble is limiting because it conveys “the modular structure imparted by ‘*[reconfigurable] high speed memory chip module*’” (Ex. 2016 at 12)—is incompatible with several of Apple’s IPR mappings. Most prominently, Apple’s Ground 1I–1S mappings (Claim 1 and its dependents over Lee-Tsern-Matsubara) depend on importing teachings from Matsubara, which is a **system LSI with embedded DRAM** (Ex. 1006, 3:65–4:7)—precisely the kind of “SOC with embedded memory in one single chip” that Apple’s own district-court brief argues falls **outside** the preamble-limited claim scope. Ex. 2016 at 11 (citing extrinsic evidence that “memory chip module” connotes “a self-contained and usually standardized[] unit . . . [that] can be disconnected and removed as a single unit”). Apple cannot use its IPR mapping to import an embedded-DRAM teaching for a limitation in claims that Apple has just told the district court are limited to a self-contained module structure. Thus, for Apple to argue in district court that the preambles of the ’942 Patent are limiting (to further Apple’s non-infringement positions) only to then argue here that the preambles are not limiting (to further Apple’s invalidity defenses) certainly does not provide greater predictability and certainty in the patent system—and is a plainly unfair tactic.

B. The Petition Independently Warrants Discretionary Denial Under 35 U.S.C. § 325(d) Because the Examiner Specifically Considered Lee and Concluded the Prior Art Did Not Teach the “*Converting*” Limitation

Lee (US 2004/0260864) was specifically considered by the Examiner during prosecution of the '942 Patent. Lee was listed in Applicant's Information Disclosure Statement, Ex. 1002, 19–20, and was independently cited by the Examiner in the Notice of References Cited (PTO-892), *id.*, 17 (Row B: “US-2004/0260864 A1 | 12-2004 | Lee et al.”). The Examiner expressly confirmed that “ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /H.B./” *Id.*, 19–20. Lee was not lined through.

On a record that included Lee, the Examiner allowed the Challenged Claims. *Id.*, 14. In the Reasons for Allowance, the Examiner stated: “there is no teaching, suggestion, or motivation for combination in the prior art to a high speed memory including where the logic unit is further used for **converting the first set of parallel data of the first transmission bus into a second set of parallel data through a second transmission bus.**” *Id.*, 15 (emphasis added).

Under *Advanced Bionics*, and per § 325(d), the Board must consider whether “the same or substantially the same prior art or arguments previously were presented to the Office,” and if so, whether “the petitioner has demonstrated that the Office erred in a manner material to the patentability of challenged claims.” IPR2019-

01469, Paper 6 at 7-8. Both prongs favor denial here. Lee **was** presented to the Office and **was** specifically considered. The Petition does not acknowledge that Lee was before the Examiner, much less identify any material error in the Examiner’s evaluation of Lee. Apple’s affirmative position—that Lee discloses precisely the “*converting*” limitation the Examiner found absent from the prior art (including Lee)—is a request that the Board reach a result opposite to the Examiner’s express finding on the same reference.

Apple’s renewed challenge on Lee, on the same record the Examiner already considered and squarely addressed, is the exact situation § 325(d) and the *Advanced Bionics* framework are intended to prevent. Thus, denial under § 325(d) is warranted on this record.

C. Grounds 1A and 1I Fail Because Apple’s [19a] Mapping Improperly Fills the “ASIC Processor” Gap in Lee with Expert Testimony and Non-Asserted References, and Because Lee’s “Generating Memory Requests” Does Not Satisfy [19e]’s “Corresponding To” Requirement

Ground 1A is presented as a single-reference § 103 ground over Lee alone. Pet. 1, 5. Ground 1I (Claims 1, 5, and 9–10) is presented as a § 103 combination over Lee-Tsern-Matsubara. Pet. 83. But Apple’s Ground 1I mapping at Claim 5 element [5b] expressly directs the Board to “*See [19a] and [19e] of Ground 1A.*” Pet. 90. Both Grounds therefore stand or fall on the Ground 1A [19a] and [19e]

mappings—which fail on textual grounds and on Rule 104(b)(4) grounds enforced under the Stewart Memorandum.

1. Apple’s [19a] Mapping Improperly Fills the “ASIC Processor” Gap in Lee with Expert Testimony and Non-Asserted References

Claim 19 recites “*an ASIC processor.*” ’942 Patent, Cl. 19. Lee describes its controller in non-ASIC terms—“a controller 52” coupled to memory modules through a “high-speed link 58.” Ex. 1004, [0015]. Elsewhere Lee describes “a system controller 110, which is also sometimes referred to as a ‘North Bridge’ or ‘memory controller.’” *Id.*, [0023]. Lee says nothing about whether the controller is implemented as an ASIC, especially an ASIC **processor**.

To bridge that gap, Apple offers expert testimony that “[a] POSITA would have understood that a North Bridge controller is an ASIC processor.” Pet. 11 (citing Ex. 1003, ¶ 40). That is the precise use of “expert testimony . . . to supply a missing claim limitation” that the Stewart Memorandum prohibits. Stewart Memorandum, July 31, 2025, at 2. The “general knowledge” of a POSITA that Lee’s North Bridge purportedly “is” an ASIC processor is not a teaching of Lee; it is a characterization drawn from outside Lee.

Apple’s expert, in turn, supports that characterization by citing Reeves (Ex. 1010) and Bruno (Ex. 1011: “*See* EX1010, 5:30–33 (‘Northbridge 14 is generally considered an . . . application specific integrated circuit (ASIC)’); EX1011, [0032]

(‘an ASIC die . . . such as a north bridge’); EX1003, ¶ 40.” Pet. 11. The “ASIC” characterization Apple’s expert offers comes from those references, not from Lee. The result is a multi-reference § 103 combination in substance—Lee + Reeves + Bruno—presented as a single-reference § 103 ground. That framing of an obviousness ground is expressly prohibited by 37 CFR § 42.104(b)(4), which requires that the Petition “specify where each element of the claim is found in the prior art patents or printed publications relied upon” within a given ground. 37 C.F.R. § 42.104(b)(4). The Stewart Memorandum makes the prohibition explicit: a petitioner cannot use expert testimony or general knowledge to import additional prior art into a ground that does not assert that art. Stewart Memorandum, July 31, 2025, at 2.

Apple offers an alternative theory that “a POSITA would have found it obvious to implement Lee’s controller as an ASIC processor.” Pet. 12. But the support Apple offers for that alternative—that ASIC implementation “was standard” and provides “performance advantages”—is itself sourced to the same out-of-ground references (Reeves (Ex. 1010) at 5:30–33, Bruno (Ex. 1011) at [0032], and Ex. 1031). Pet. 12–13. Without Reeves, Bruno, and Ex. 1031, the alternative theory is unsupported attorney argument and conclusory *KSR* boilerplate. *See Pers. Web Techs., LLC v. Apple, Inc.*, 848 F.3d 987, 994 (Fed. Cir. 2017) (obviousness requires “articulated reasoning with rational underpinning,” not conclusory rationales).

The Ground 1A [19a] defect is incorporated by Apple into Ground 1I (beginning at Claim 5) and the Grounds that build off Ground 1 (Grounds 1M, 1N), and into Grounds 1B–1H through Ground 1B’s express incorporation of Ground 1A’s [19e]–[19f] mapping (Pet. 49). Apple’s cross-references mean the [19a] defect propagates through every Ground that touches Claims 5, 19, and claims that depend from those claims.

2. Lee’s “Generating Memory Requests” Does Not Satisfy [19e]’s “Corresponding To” Requirement

Limitation [19e] requires that “*the ASIC processor executes a predetermined function corresponding to the second set of parallel data while the ASIC processor is under operating condition.*” ’942 Patent, Cl. 19 (emphasis added). The ’942 specification confirms what the limitation requires: “the ASIC processor 112 can execute a predetermined function *according to* the second set of parallel data.” *Id.*, 4:15-17 (emphasis added).

Apple maps [19e] to Lee’s controller 52 ““generat[ing] memory requests that are coupled to the memory modules.”” Pet. 28 (quoting Ex. 1004, [0026]). Apple contends that “[g]enerating a memory request by the controller (predetermined function) *results in* the transmission of the second set of data along Lee’s high-speed bus to/from the memory module.” Pet. 28 (emphasis added). The “results in” framing is fatal to Apple’s mapping. The second set of parallel data is the *result* of

Lee’s controller generating a memory request—it is what comes back from the read. The function (generating the request) precedes the data; the data does not exist when the function is performed. There is nothing in Lee that describes Lee’s controller doing anything responsive to, or operating on, the contents of the second set of parallel data—which is what [19e]’s “corresponding to” language requires.

Apple expressly incorporates the Ground 1A [19e] mapping into the Grounds that build off Ground 1A. Pet. 49 (Ground 1B: “[19e]–[19f] *See* [19e]–[19f] of Ground 1A. EX1003, ¶ 128.”); Pet. 90 (Ground 1I, Claim 5: “[5b] *See* [19a] and [19e] of Ground 1A.”). The [19e] defect therefore carries through every Claim 19 Ground (1A–1H) and into Ground 1I at [5b], which carries to every ground that builds off 1I and reaches Claim 5 or its dependents.

D. Grounds 1A and 1I Independently Fail Because Apple’s “*Converting*” Mapping Cross-Cites Nakamura—a Reference Not Asserted in Any Ground—and Because Lee Discloses the Coupling of 128-Bit Parallel Data, Not Conversion

Limitations [19d] and [1c] both recite the same “*converting*” limitation: the logic unit “*is further used for converting the first set of parallel data of the first transmission bus into a second set of parallel data through a second transmission bus.*” ’942 Patent, Cls. 1 and 19 (emphasis added). The “*converting*” limitation appears, in identical terms, in both independent claims, and Apple’s mapping for [19d] is the same mapping it asks the Board to apply to [1c]: Apple’s Ground 1I

mapping at [1c] expressly directs the Board to “*See* [19d] of Grounds 1A-1B.” Pet. 90. The “*converting*” defect therefore reaches Ground 1A ([19d] directly), Grounds 1B–1H (which build off Ground 1A, Pet. 49), Ground 1I (via [1c] cross-reference, Pet. 90), and Grounds 1J–1S (which build off Ground 1I).

1. Lee Describes Coupling of 128 Bits Through the High-Speed Link, Not Conversion

The Lee paragraph on which Apple’s “conversion” theory rests is [0017]. Pet. 7–8. Lee says: “The 128 bits of data can be coupled through the high-speed link 58 by either increasing the speed of the link 58 or the size of the data word **coupled** through the link 58. For example, in the high-speed mode, a 128-bit data word may be **coupled** through the link 58. Therefore, for every memory access, 128 data bits will be **coupled** through the link 58. Alternatively, the link 58 may transfer only a 32-bit data word, but it may operate at four times the speed of the memory devices 70.” Ex. 1004, [0017] (emphasis added).

Lee uses the language of coupling, not “*converting*.” Lee describes two formats through which the same 128-bit payload may travel onto the link—either as a single 128-bit word or as four serially-sent 32-bit words at four times the speed—and confirms that “[t]herefore, for every memory access, 128 data bits will be **coupled** through the link 58.” *Id.* The data quantity is preserved (128 bits), the data content is preserved (the same bits read from memory), and the difference between

the two formats is a **transmission-format** difference rather than a transformation of one set of data into a different set. That is bandwidth equivalence; it is not conversion. The '942 specification, by contrast, describes “*converting*” as a substantive operation performed by a “*parallel-to-serial controller*” within the logic unit, where the first bus operates at one bit width and the second bus operates at a different bit width: “the logic unit 106 utilizes a parallel-to-serial controller 1062 included thereof to convert the first set of parallel data ... to the [ASIC processor 112] through a second transmission bus.” '942 Patent, 4:4-9. Lee’s mode-switching language describes addressing reconfiguration (the high-bandwidth mode addresses all memory devices simultaneously; the high-memory-depth mode addresses one rank at a time), not conversion of data between bus widths. Ex. 1004, [0017]–[0018].

2. Apple Fills the Gap in Lee with Nakamura, a Reference Not Asserted in Any Ground

In view of Lee’s coupling language, Apple simply has its expert characterize Lee’s coupling as “*converting*” by importing Nakamura: “A POSITA would have found obvious that receiving 128 bits of data simultaneously through multiple bus systems and then transmitting that data in a 64- or 32-bit word through a single bus system involves the memory hub converting the received data (including the first set of parallel data) to the transmitted data (second set of parallel data), e.g., by multiplexing the received data. EX1003, ¶¶66–68 (citing EX1019 (Nakamura),

Abstract (‘a parallel-to-parallel converter for converting an ‘m’-bit parallel signal into an ‘n’-bit parallel signal’)).” Pet. 25–26. Per Apple, Nakamura supplies the substantive proposition that Lee’s asserted multiplexing is “*converting*” within the meaning of the claim. Apple repeats the same citation chain at Pet. 26 and Pet. 27.

Nakamura (Ex. 1019) is not in Ground 1A’s asserted basis (“§ 103: Lee”). Pet. 1. It is not asserted in any other Ground. Pet. 1–2. Yet Nakamura supplies the dispositive characterization that purportedly transforms Lee’s function into the claimed “*converting*.” That is again the precise use of expert testimony and non-asserted prior art that the Stewart Memorandum forbids: “expert testimony, common sense, and other evidence that is not ‘prior art consisting of patents or printed publications’ . . . may not be used to supply a missing claim limitation.” Stewart Memorandum, July 31, 2025, at 2. The Petition must “specify where each element of the claim is found in the prior art patents or printed publications relied upon” within the ground. 37 C.F.R. § 42.104(b)(4).

3. Apple’s Bus-System-68 Alternative Theory Does Not Cure the Defect

Apple offers an “alternative/additional” theory in which Lee’s collective bus system 68 (composed of the four individual bus systems 76, 82, 90, 94) is the “*first transmission bus*.” Pet. 19–22, 27–28. On that theory, Apple says, the 128 bits read simultaneously across bus systems 76, 82, 90, 94 are the first set of parallel data; the

64-bit word transmitted on link 58 is the second set. *Id.* 27.

The alternative theory inherits the same defects. Lee still does not describe “*converting*”; Lee describes how the same 128 bits are coupled through link 58 in different formats. Apple’s alternative mapping cross-cites Nakamura to bridge the gap: “receiving 128 bits of data simultaneously through multiple bus systems and then transmitting that data as a 64- or 32-bit word through a single bus system involves converting the received data . . . , e.g., by multiplexing the received data. EX1003, ¶72 (citing EX1019, Abstract).” Pet. 27.

The Petition thus leaves Patent Owner and the Board to guess which alternative theory operates in any of the eleven downstream Grounds 1B–1S that incorporate Ground 1A’s [19d] mapping by reference. Apple does not specify which of the two competing first-transmission-bus theories is operative for any of those Grounds. That ambiguity is itself incompatible with the “particularity” § 104(b)(4) requires. Ground 1A’s [19d] mapping, Ground 1I’s [1c] mapping (which incorporates [19d] by reference, Pet. 90), and the Grounds that build on either of them, thus all fail.

E. Grounds 1B–1S (the Lee-Tsern Combinations) Independently Fail Because Inserting Tsern’s Buffer into Lee Destroys Lee’s Principle of Operation

Apple’s Lee-Tsern combinations all depend on substituting Tsern’s serialization/deserialization buffer for Lee’s memory hub. That substitution destroys

Lee's principle of operation. Lee's defining feature is its high-bandwidth mode in which "the memory hub 60 addresses all of the memory devices 70 simultaneously so that 128 bits of data are written to or read from the memory devices 70 each memory access." Ex. 1004, [0017]. The 128-bit simultaneous parallel access is the very feature Lee uses to deliver high bandwidth, and Lee's configurability is achieved by reprogramming "a non-volatile register 98 in the memory module 54" to switch modes—not by inserting a serialization layer between the memory devices and the hub. *Id.* Tsern's buffer interface, by contrast, includes "a [c]onfigurable serialization/deserialization circuit 1891" that "performs serialization and deserialization functions" depending on "memory device access width" to bridge mismatched interface widths between memory devices and a single buffer interface. Ex. 1005, 31, 38. Substituting Tsern's serialization-based bandwidth-concentration architecture for Lee's simultaneous-parallel-access architecture either eliminates Lee's 128-bit parallel operation (defeating Lee's stated purpose) or re-architects Tsern's buffer in a manner Tsern itself does not contemplate. *See In re Gordon*, 733 F.2d at 902 (combination unpatentable where it would render the primary reference "inoperable for its intended purpose" or "require[] a substantial reconstruction").

Apple's motivation rationale—"known technique applied to known system to yield predictable results"—is conclusory *KSR* boilerplate that does not engage with the architectural conflict between Lee and Tsern. Pet. 41–43. *KSR* requires an

“articulated reasoning with some rational underpinning to support the legal conclusion of obviousness.” *KSR Int’l Co.*, 550 U.S. at 418 (quoting *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006)).

F. The Dependent-Claim Grounds Relying on Lin (Grounds 1F, 1O) and White (Grounds 1H, 1Q, 1S) Independently Fail Because Lin and White Disclose Single-Chip Structures, Not the Per-Memory-Cell-Array-IC Seal Rings and TSVs the Dependent Claims Require

Several of the ’942 Patent’s dependent claims recite multi-chip-module-specific structural features—per-IC seal rings, TSVs outside those seal rings, and a stacked arrangement—that the references Apple cites do not disclose. Specifically Claim 11 recites that “*each of the memory cell array ICs of the type of memory cell array group corresponds to a first seal ring **and** the logic unit corresponds to a second seal ring, a circuit of each of the memory cell array ICs is surrounded by a first seal ring and a circuit of the logic unit is surrounded by the second seal ring.*” ’942 Patent, Cl. 11 (emphasis added). Claim 12 adds that “*at least one through silicon via (TSV) exists outside a first seal ring **of each of the memory cell array ICs.***” *Id.*, Cl. 12 (emphasis added). Claims 27, 28 recite the same per-IC seal-ring and per-IC TSV arrangement in the electronics-system-device claims. *Id.*, Cls. 27, 28. Claim 14 adds that the per-IC TSVs “are used for connecting to a power or a ground supply.” *Id.*, cl. 14.

**1. Lin Discloses a Single-Chip Seal-Ring/TSV Structure,
Not Per-Memory-Cell-Array-IC Seal Rings**

Apple’s Ground 1F (Claims 27–28 and 31) and Ground 1O (Claims 11–12 and 15) rely on Lin (Ex. 1013) to supply the seal-ring and TSV limitations. Pet. 1–2, 65–73, 96–97. Lin describes **a single wafer** with “an active circuit region 140” encompassed by a “seal ring 120” and TSVs in adjacent “TSV regions.” Ex. 1013, [0016], [0024]; Pet. 65–67. Lin’s Fig. 3b and Fig. 8 each depict the wafer-level structure: one active circuit region, one seal ring, TSV regions outside the seal ring. Ex. 1013, Figs. 3b, 8.

Lin’s single-chip structure does not disclose the multi-die arrangement the ’942 Patent’s dependent claims require. Claims 11 and 12 recite that **each** of the multiple memory cell array ICs has its own first seal ring with TSVs outside, and that the logic unit has a **separate, second** seal ring. ’942 Patent, Cls. 11, 12, 27, 28. Lin discloses one seal ring on one chip; the Petition’s mapping does not (and cannot) point to a Lin disclosure of multiple ICs each with its own seal ring.

Apple’s mapping does not identify any disclosure in Lin showing multiple memory cell array ICs each having its own seal ring. Apple instead simply asserts, without supporting citation to any of the asserted references, that “[i]n the combination, the memory hub and the memory devices are ICs each including a seal ring” Pet. 67. That is attorney argument—not a teaching of any of the asserted

prior art. The Stewart Memorandum and particularity requirement of § 42.104(b)(4) obligates Apple to identify **where in Lin** the multi-IC arrangement is shown, or to articulate a sufficient motivation for a POSITA to replicate Lin’s single-chip structure across multiple discrete ICs in a combination. Apple does neither beyond conclusory *KSR* formulations: “application of known techniques . . . to a known structure . . . to yield predictable results.” Pet. 68. The motivation Apple does articulate (TSVs allow chip stacking; seal rings block moisture) (Pet. 68) is generic and does not address why a POSITA would replicate Lin’s single-chip structure across each memory IC **and** the separate logic IC. 37 C.F.R. § 42.104(b)(4).

2. White Likewise Discloses TSVs in a Single Die, Not Per-Memory-Cell-Array-IC TSVs

Grounds 1H, 1Q, and 1S rely on White (Ex. 1014) to supply the TSV-power-or-ground-supply limitations of Claims 14, 18, and 30. Pet. 1–2, 79–82, 97–98, 101. White describes TSVs in a single die. Apple’s mapping (Pet. 80–82) does not identify where in White the multi-die per-IC TSV arrangement of the ’942 dependent claims is disclosed. As with Lin, Apple supplies the multi-IC arrangement by stipulation—“In the combination . . . based on White” (Pet. 82)—without identifying a White disclosure of the multi-IC arrangement or a sufficient motivation to replicate White’s single-die structure across multiple ICs.

3. The Cumulative-Reference Grounds (1P–1S) Compound the Same Defects

Grounds 1P through 1S layer additional references (Ali, Liao) onto Lee-Tsern-Matsubara-Lin or Lee-Tsern-Matsubara-Lin-White to address particular dependent claim features (Claim 13, 16, 17). Pet. 1–2, 97–101. Each of those Grounds inherits the foundational defects of Grounds 1A, 1B, 1I, and 1O (or 1H) it builds upon. None of the additional references (Ali for the topmost-memory-cell-array metal layer (Cl. 29); Liao for the stacked arrangement (Cl. 16)) cures the foundational defects in the “*ASIC processor / predetermined function*” mapping (§ IV.B), the “*converting*” mapping (§ IV.C), the Lee-Tsern destruction-of-purpose problem (§ IV.E), or the per-IC seal-ring/TSV mismatch addressed in this section.

Apple’s strategy of stacking up to six references in a single ground (Ground 1S: Lee-Tsern-Matsubara-Lin-Liao-White, Pet. 1–2, 101) does not strengthen the obviousness showing; it compounds the motivation-to-combine problem. The more references that must be combined, the more rigorously Apple must articulate **why** a POSITA would have combined them. Apple’s rationales remain conclusory *KSR* formulations across the cumulative-reference grounds. *In re Van Os*, 844 F.3d 1359, 1361 (Fed. Cir. 2017); *Pers. Web Techs.*, 848 F.3d at 994.

Grounds 1F, 1H, 1O, 1P, 1Q, 1R, and 1S should be denied for these independent reasons in addition to the foundational defects in Grounds 1A, 1B, and

1I that they inherit.

V. CONCLUSION

For each of the foregoing reasons—any one of which is independently sufficient—Patent Owner respectfully requests that the Board deny institution.

Respectfully submitted,

Dated: May 20, 2026

By: /Vincent J. Rubino, III/
Vincent J. Rubino, III (Reg. No. 68,594)
Lead Counsel for Patent Owner
Fabricant, Rubino & Lambrianakos LLP
411 Theodore Fremd Avenue,
Suite 206 South
Rye, New York 10580
Tel. 212-257-5797
Fax. 212-257-5796
Email: vrubino@frlip.com

CERTIFICATE OF WORD COUNT

The undersigned hereby certifies that the portions of the above-captioned PATENT OWNER'S PRELIMINARY RESPONSE has 6,488 words in compliance with the 14,000 word limit set forth in 37 C.F.R. § 42.24. This word count was prepared using Microsoft Word for Office 365.

Respectfully submitted,

May 20, 2026

By: /Vincent J. Rubino, III/
Vincent J. Rubino, III (Reg. No. 68,594)
Lead Counsel for Patent Owner
Fabricant, Rubino & Lambrianakos LLP
411 Theodore Fremd Avenue,
Suite 206 South
Rye, New York 10580
Tel. 212-257-5797
Fax. 212-257-5796
Email: vrubino@frlip.com

CERTIFICATE OF SERVICE

A copy of the foregoing Patent Owner's Preliminary Response and Exhibit
2016 has been served on Petitioner's counsel of record as follows:

Email: IPR50095-0275IP@fr.com
W. Karl Renner
Gretchen DeVries
Hyun Jin In
Andrew B. Patrick
FISH & RICHARDSON P.C.
60 South Sixth Street, Suite 3200
Minneapolis, MN 55402

Attorneys for Apple Inc.

May 20, 2026

By: /Vincent J. Rubino, III/
Vincent J. Rubino, III (Reg. No. 68,594)
Lead Counsel for Patent Owner