

**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION**

BiTMICRO LLC,

Plaintiff,

v.

KIOXIA AMERICA, INC.
and KIOXIA CORPORATION,

Defendants.

Civil Action No.: 6:22-cv-00331-ADA

BiTMICRO LLC,

Plaintiff,

v.

INTEL CORPORATION,

Defendant.

Civil Action No.: 6:22-cv-00335-ADA

PLAINTIFF BITMICRO LLC'S SUR-REPLY CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

	Page
I. '939 PATENT	1
A. "predetermined level"	1
B. "means for deactivating the plurality of super capacitors to cut off power to the computing engine based upon the plurality of super capacitors discharging to a predetermined level"	2
C. "reversing the flow of current between the computing engine and the plurality of super capacitors"	4
D. "means for reversing the flow of current between the computing engine and the plurality of super-capacitors"	5
II. '243 PATENT	5
A. "passive port"	5
B. "a plurality of modules each comprising"	7
III. '740 PATENT	8
A. "optimized memory operations"	8
B. "optimally distributed"	10
C. "memory store"	11
IV. '190 PATENT	13
A. "optimal"	13
V. '084 AND '694 PATENTS	14
A. "Direct Memory Access (DMA) descriptors" / "DMA controller descriptor"	14
VI. CONCLUSION	15

TABLE OF AUTHORITIES

	Page(s)
<u>CASES</u>	
<i>Advanced Micro Devices, Inc. v. LG Elecs., Inc.</i> , No. 14-cv-01012-SI, 2017 WL 1383271 (N.D. Cal. Apr. 18, 2017).....	9
<i>Cacace v. Meyer Mktg. (Macau Commercial Offshore) Co.</i> , 812 F. Supp. 2d 547 (S.D.N.Y. 2011).....	13
<i>CCS Fitness, Inc. v. Brunswick Corp.</i> , 288 F.3d 1359 (Fed. Cir. 2002).....	2
<i>Cont’l Circuits LLC v. Intel Corp.</i> , 915 F.3d 788 (Fed. Cir. 2019).....	1
<i>Data Retrieval Technology, LLC v. Sybase, Inc.</i> , No. C 08-5481 VRW, 2010 WL 11586656 (N.D. Cal. Nov. 8, 2010)	13
<i>Info-Hold, Inc. v. Applied Media Techs. Corp.</i> , 783 F.3d 1262 (Fed. Cir. 2015).....	1
<i>Interval Licensing LLC v. AOL, Inc.</i> , 766 F.3d 1364 (Fed. Cir. 2014).....	9
<i>Liebel-Flarsheim Co. v. Medrad, Inc.</i> , 358 F.3d 898 (Fed. Cir. 2004).....	2
<i>Moleculon Research Corp. v. CBS, Inc.</i> , 793 F.2d 1261 (Fed. Cir. 1986).....	9
<i>Nabors Drilling Techs. USA, Inc. v. Helmerich & Payne Int’l Drilling Co.</i> , No. 3:20-CV-03126-M, 2022 WL 1689444 (N.D. Tex. May 26, 2022)	14
<i>Nichia Corp. v. Document Sec. Sys.</i> , No. 2020-2261, 2022 WL 1218036 (Fed. Cir. Apr. 26, 2022)	7
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005).....	1, 5, 15
<i>RF Delaware, Inc. v. Pacific Keystone Techs., Inc.</i> , 326 F.3d 1255 (Fed. Cir. 2003).....	12
<i>Toro Co. v. White Consol. Indus., Inc.</i> , 199 F.3d 1295 (Fed. Cir. 1999).....	2

<i>Versata Software, Inc. v. Zoho Corp.</i> , 213 F. Supp. 3d 829 (W.D. Tex. 2016).....	9
<i>Wenger Mfg., Inc. v. Coating Machinery Sys., Inc.</i> , 238 F.3d 1225 (Fed. Cir. 2001).....	12
<i>WesternGeco L.L.C. v. ION Geophysical Corp.</i> , 876 F. Supp. 2d. 857 (S.D. Tex. 2012)	3
<i>Williamson v. Google Inc.</i> , No. 15-cv-00966-BLF, 2017 WL 3232582 (N.D. Cal. July 27, 2017).....	3

STATUTES

35 U.S.C. § 112(6)	2, 5
--------------------------	------

Plaintiff BiTMICRO LLC (“BiTMICRO”) submits this sur-reply brief in support of its claim constructions.

I. '939 PATENT

A. “predetermined level”

Patent	BiTMICRO’s Proposal	Defendants’ Proposal
'939 Patent Claims 1, 10	Plain and ordinary meaning.	“a preset minimum operating voltage necessary for proper operation, in which the super capacitors are not fully discharged”

“In construing claims, district courts give claims their ordinary and customary meaning, which is ‘the meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention.’” *Cont’l Circuits LLC v. Intel Corp.*, 915 F.3d 788, 796-797 (Fed. Cir. 2019) (quoting *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312-13 (Fed. Cir. 2005)). Defendants have provided no reason to diverge from that principle with respect to “a predetermined level.” They have not identified anything in the claims that compels the import of the two limitations in their construction. Nor have Defendants identified anything in the specification or prosecution history wherein the patentee “clearly set forth a definition of the disputed claim term other than its plain and ordinary meaning” or “expressions of manifest exclusion or restriction, representing a clear disavowal of claim scope.” *Cont’l Circuits*, 915 F.3d at 1796-797 (internal quotations and citations omitted). At best, they have identified a description of an embodiment or even a “best mode,” but that is not enough. *Id.*

Defendants did not make any “salient claim arguments” in their opening brief, nor does BiTMICRO agree with Defendants’ mischaracterizations of the '939 Patent therein or on reply. Sure, the '939 Patent specification discloses one embodiment wherein “a pre-determined level” is such that “the down-converter has insufficient voltage differential to continue proper operation and shuts down” ('939 Patent at 5:39-43), but Defendants have not shown that the term must be construed as limited to that embodiment, nor should it be. *Cont’l Circuits*, 915 F.3d at 797; *Info-Hold, Inc. v. Applied Media Techs. Corp.*, 783 F.3d 1262, 1267 (Fed. Cir. 2015); *Liebel-Flarsheim*

Co. v. Medrad, Inc., 358 F.3d 898, 906 (Fed. Cir. 2004). And Defendants oversell the cited portions of the specification at 5:44-48 and 6:45-49, which describe “two positive side effects” of that embodiment. *Compare* Reply at 1 with ’939 Patent at 5:44-57. Nor is it true, as Defendants contend, that “a threshold set below a minimum operating voltage necessary for proper operation would fatally undermine the ’939 Patent’s stated goals of ensuring that ‘a user is able to correctly store large amounts of newly written and modified data’ and ‘rapidly and irretrievably erase data.’” Reply at 1. Indeed, Defendants mischaracterize the relationship between those goals and the “minimum operating voltage necessary for proper operation.” Provided that the claimed system can supply enough power to the computing engine to store or erase the data in question *before* the capacitors are discharged to a predetermined level (*i.e.*, the capacitors have sufficient capacity), the system operates properly. That is perhaps why those portions of the ’939 Patent that discuss “goals” of the claimed invention do not mention the “minimum operating voltage necessary for proper operation.” See *e.g.*, ’939 Patent at 1:8-2:21, 2:25-43, 3:3-27. For these reasons, Defendants’ cited cases are inapposite. See *CCS Fitness, Inc. v. Brunswick Corp.*, 288 F.3d 1359, 1366-67 (Fed. Cir. 2002); *Toro Co. v. White Consol. Indus., Inc.*, 199 F.3d 1295, 1301 (Fed. Cir. 1999) (statements admitting the “important to the invention”).

Finally, Defendants try to run away from the fact that they used “predetermined level” in their own patents directed to similar subject matter. But BiTMICRO is not, as in the cases cited by Defendants, seeking to import a definition or limitations from another patent. Rather, it is simply relying on Defendants’ use in its other patents to confirm what is already obvious to the Court—that the term “a predetermined level” has a plain and ordinary meaning to those skilled in the art.

B. “means for deactivating the plurality of super capacitors to cut off power to the computing engine based upon the plurality of super capacitors discharging to a predetermined level”

Patent	BiTMICRO’s Proposal	Defendants’ Proposal
’939 Patent Claim 10	<u>Function</u> : deactivating the plurality of super capacitors to cut off power to the computing engine based upon the plurality of super capacitors discharging to a predetermined level.	Indefinite under 35 U.S.C. § 112(6): Lacks corresponding structure.

	Structure: Col. 3:3-12, 3:28-4:8, 5:35-42, 6:37-44, Figs. 1, 3, 5, of the '939 Patent, and equivalents thereof.	
--	---	--

After being called out for putting forward nonsensical arguments that rely on nonsensical meanings for “deactivating the plurality of supercapacitors,” Defendants seek to foist the liability for their arguments on BiTMICRO and the '939 Patent. *Compare e.g.*, Resp. at 9 (questions posed to Defendants) *with* Reply at 3 (Defendants deflecting and refusing to answer). But that is absurd.

The claim refers to “deactivating the plurality of super capacitors to cut off power to the computing engine.” The plain and ordinary meaning of “deactivating the plurality of super capacitors” is to make inactive, *e.g.*, by disconnecting from other elements of the system as in “the switch deactivates the alarm.” Hence, the plain and ordinary meaning of the claim language is to make the plurality of super capacitors inactive, *e.g.*, by disconnecting them from the rest of the system, “to cut off power to the computing engine.” The claim language also conveys when this is done: *i.e.*, “based upon the plurality of super capacitors discharging to a predetermined level.”

This language is so straightforward that neither a POSITA nor a jury would require more explanation. Consistent therewith, Defendants did not challenge this same language in claim 1 for indefiniteness, even though it contains the exact same language save for the phrase “means for.”

Claim 1 (not challenged)	Claim 10 (challenged)
(c) deactivating the plurality of super capacitors to cut off power to the computing engine based upon the plurality of super capacitors discharging to a predetermined level.	<i>means for</i> deactivating the plurality of super capacitors to cut off power to the computing engine based upon the plurality of super capacitors discharging to a predetermined level.

In other words, Defendants do not believe this limitation is indefinite, but nevertheless seek to take advantage of the case law built up around means-plus-function terms to try to get it rendered indefinite anyway. But to do so, they had to take nonsensical positions to which no expert would sign on. *See Williamson v. Google Inc.*, No. 15-cv-00966-BLF, 2017 WL 3232582, at *10 (N.D. Cal. July 27, 2017); *see also WesternGeco L.L.C. v. ION Geophysical Corp.*, 876 F. Supp. 2d. 857, 875 (S.D. Tex. 2012).

The simple fact is that the '939 Patent specification discloses structure that deactivates the plurality of super capacitors to cut off power to the computing engine based on the plurality of

super capacitors discharging to a predetermined level. As noted in BiTMICRO’s response, that structure consists of the “charge level sensors 30,” which informs the computer system whether the plurality of capacitors have discharged to a predetermined level, the “down converter DC 42,” which turns off or shuts down thereby disconnecting the plurality of super capacitors from the rest of the system and accordingly “cut[ting] off power to the computing engine,” and other related circuitry (like the power isolator and up converter that revert to their original states). *See* Resp. at 6-9. Defendants continue to push the absurd argument that “deactivating” super capacitors requires a different structure than that disclosed, but have never been able to identify what such a structure might be. And Defendants’ argument that the structures disclosed are not “clearly linked” to the claim language is patently absurd, since a 1:1 mapping of structure and claim is clearly evident.

C. “reversing the flow of current between the computing engine and the plurality of super capacitors”

Patent	BiTMICRO’s Proposal	Defendants’ Proposal
’939 Patent Claim 6	Plain and ordinary meaning; not indefinite.	Indefinite.

Here, Defendants again push arguments that it could not get a single expert to support. But BiTMICRO’s explanation of the current flow in its response brief is exactly what the specification describes. For example, consider the dashed line through the center of Figure 1; the specification explains that “the SC 34 and associated circuitry 30, 32, 36, 38, 42 [may be] configured as a removable option to the system 20.” ’939 Patent at 3:63-4:2. Such an embodiment clearly illustrates the “shared power plane of the EPS, Power Isolator and CE” that seems to confuse the Defendants. *See* Reply at 5, Resp. at 11. Moreover, Defendants’ admission that current flows in one direction when external power is supplied, and another direction when power is supplied by the capacitors, is entirely consistent with how a POSITA and jury would understand the claim. *Compare e.g.*, Reply at 5 with Resp. at 11 (“These disclosures in the specification make clear that ‘reverse’ in the claim language is referring to the change in the flow of current in the system from the normal state (when power is supplied by the EPS to the computing engine and to the capacitors to charge them) and backup state (when power is supplied from the capacitors to the computer

engine)”).

Defendants’ fallback argument—that the claim language requires the “computing engine” to provide power to super capacitors when external power is being supplied—is technical to a fault. *Phillips* makes clear that claims are to be construed from the perspective of a POSITA, and that the claims “do not stand alone” but are guided by the specification. A POSITA would have little trouble understanding the claim language at issue in view of those sources. Resp. at 9-11.

D. “means for reversing the flow of current between the computing engine and the plurality of super-capacitors”

Patent	BiTMICRO’s Proposal	Defendants’ Proposal
’939 Patent Claim 15	<u>Function</u> : reversing the flow of current between the computing engine and the plurality of super capacitors. <u>Structure</u> : Col. 3:3-12, 3:28-4:8, 6:22-28, Figs. 1, 3, 5, of the ’939 Patent, and equivalents thereof.	Indefinite under 35 U.S.C. § 112(6): Lacks corresponding structure.

This term is not indefinite for the same reasons as discussed with respect to the previous term. And the disclosures recited with respect to that term and the “deactivating” term above also identify the corresponding structure, i.e., the “circuitry” in the cited passages of the specification.

II. ’243 PATENT

A. “passive port”

Patent	BiTMICRO’s Proposal	Intel’s Proposal
’243 Patent Claims 1, 2	“connection between a passive ball and a passive pad”	“connection between a passive ball on one surface of a SDRAM module and a passive pad on another surface of the same SDRAM module”

Intel incorrectly argues that the ’243 applicants affirmatively stated that they were acting as their own lexicographers to define “passive port.” The applicants merely stated to the examiner the general principle that “[a]pplicants may be their own lexicographers,” along with the legal principle that “[d]efiniteness of claim language must be analyzed, not in a vacuum, but in light of Applicants written disclosure, the prior art, and the claim interpretation that would be given by one of ordinary skill in the art” Ta Decl., Ex. F at 23. Merely reciting the state of

the law to the examiner is not the same as explicitly defining a term.

The applicants did not intend to limit “passive port” to just the Figure 19 embodiment, as Intel contends. To the contrary, the specification repeatedly notes that Figure 19 is merely an “example.” Ta Decl., Ex. F at 2 (“The figure illustrates the controller and SDRAM modules as an example.”); *id.* at 3 (“The example embodiment illustrated in Figure 19 comprises . . .”). In addition, as BiTMICRO previously noted (Resp. at 15-16), the applicants’ discussion with the examiner emphasized that the defining characteristic of “passive port” is a connection between a ball and pad, rather than the specific connection within a single SDRAM module proposed by Intel. Such a narrow construction would even be inconsistent with the Figure 19 embodiment, which shows that a passive port can connect a ball and pad on different modules (1905) as well as on the same module (1909). *See id.* at 15-16. Intel responds that there is no inconsistency because the applicants said that “[t]he above description should be read in conjunction with amended Figure 19,” which is nonsensical. Reply at 9. The specification’s description of Figure 19 specifically points out an example of a passive port as a connection between a ball and pad on different modules: ball X1 on SDRAM Module 1 and pad X2 on SDRAM Module 0. *See* Ex. 5 at 9:8; Ta Decl., Ex. F at 3 (“passive ports 1905 connecting . . . passive ball X1 to passive pad X2”).¹ Indeed, the principle cited by Intel that a patentee may draft claims to cover different embodiments has little weight here, where the embodiment that would be excluded from Intel’s proposed construction (Figure 19) is the very embodiment that the applicants cited to the examiner in their discussion of “passive port.”

Finally, Intel cannot credibly argue that claim differentiation does not apply because dependent claim 2 adds multiple limitations, given that it takes the exact opposite position with respect to the ’740 claim terms. *See* Reply at 14-15 (“BiTMICRO is wrong that the presumption of claim differentiation does not apply simply because some of the dependent claims have

¹ For the Court’s convenience, attached as Appendix A to this reply brief is a chart showing the language in the ’243 specification regarding Figure 19 after applying the changes in the Certificate of Correction.

additional limitations.”).

B. “a plurality of modules each comprising”

Patent	BiTMICRO’s Proposal	Intel’s Proposal
’243 Patent Claim 1	Part of the preamble and not a limitation of the claim.	“each module in the stacked module comprises”

The preamble of claim 1 recites “[a] stacked module comprising a plurality of modules each comprising:” Intel’s proposal to break off the latter part of this phrase and construe it as a separate claim limitation would render claim 1 nonsensical, as the actual limitations in claim 1 recite components that are not within “each module in the stacked module,” but are rather a part of the stacked module as a whole. *See Resp.* at 18-19.

Reliance on *Nichia Corp. v. Document Sec. Sys.*, No. 2020-2261, 2022 WL 1218036 (Fed. Cir. Apr. 26, 2022), is misplaced. In that case, the claim in dispute began with the language “[a] display comprising a plurality of plastic leaded chip carrier LEDs, **the plastic leaded chip carrier LEDs each comprising.**” *Id.* at *4 (emphasis added). There, the phrase “a display comprising” was separated by a comma from the subsequent phrase “the plastic leaded chip carrier LEDs each comprising.” Thus, the natural reading was that “each comprising” referred to the immediately preceding phrase “the plastic leaded chip carrier LEDs,” rather than to the term “display.” In ’243 claim 1, however, the claim language recites “[a] stacked module comprising a plurality of modules each comprising,” with no intervening comma. Here, “each comprising” does not naturally refer to “plurality of modules.” Rather, it more sensibly refers to “a stacked module” when read in context with the rest of the limitations in claim 1. As discussed in BiTMICRO’s responsive brief, claim 1 includes the limitations of “a first serial chain route” and “a second serial chain route,” both of which are routes that run through multiple modules in the stack as described in the specification. *Resp.* at 18-19. But under Intel’s construction where “each module in the stacked module” has all of the limitations of the claims, each individual module in the stack would have both “a first serial chain route” and “a second serial chain route,” which does not make sense. Tellingly, Intel does not address this argument in its reply brief. It does not do so because it cannot

reconcile this inconsistency between its proposed construction and the specification’s description of the “first serial chain route” and “second serial chain route.”

Intel’s focus on the portion of the ’243 specification discussing that all modules have a control circuit also misses the mark. The issue is not whether each module in the stack can have a control circuit. Rather, the problem with Intel’s construction is that claim 1 recites a specific requirement for the control circuit – namely, that the control circuit must be “for enabling a routing path that connects the first serial chain route with the second serial chain route ***within an end module***.” The plain language of claim 1 therefore requires the control circuit of the claim to be in the end module. Indeed, Intel offers no explanation for how a control circuit in any other module could connect the first serial chain route with the second serial chain route within the end module. Intel’s construction requiring each module in the stack to have all the limitations of claim 1 therefore makes no sense and should be rejected.

III. ’740 PATENT

A. “optimized memory operations”

Patent	BiTMICRO’s Proposal	Defendants’ Proposal
’740 Patent Claim 1	Plain and ordinary meaning; not indefinite.	Indefinite.

Defendants cannot show by clear and convincing evidence that the term “optimized memory operations” is indefinite. As discussed in BiTMICRO’s responsive brief, during prosecution of the ’740 patent the applicants explained to the examiner how the specification teaches a mapping table that causes a storage device to perform “optimized memory operations,” thereby providing a POSITA with sufficient guidance as to what that term means. Resp. at 21. Defendants incorrectly suggest that that portion of the file history is irrelevant because the discussion involved a response to a written description rejection, not an indefiniteness rejection. But it does not matter what type of rejection the applicants were responding to. Rather, what matters is that the applicants made specific statements during prosecution regarding the meaning of “optimized memory operations” from which a POSITA would understand the scope of that term

with reasonable certainty. Defendants fail to present any evidence to the contrary. Having offered no expert testimony to support their assertions, Defendants cannot prove that “optimized memory operations” is indefinite. *See Advanced Micro Devices, Inc. v. LG Elecs., Inc.*, No. 14-cv-01012-SI, 2017 WL 1383271, at *4 (N.D. Cal. Apr. 18, 2017) (defendant failed to prove indefiniteness because it “provides no expert testimony to support . . . its broader indefiniteness contentions, despite needing to prove its position by clear and convincing evidence.”).

Defendants argue that expert testimony is not required prove indefiniteness, but the cases cited by Defendants do not support their position. In *Interval Licensing LLC v. AOL, Inc.*, 766 F.3d 1364, 1370 n.6 (Fed. Cir. 2014), the Federal Circuit affirmed an indefiniteness finding without supporting expert testimony only because in that case “there are no disputes about underlying questions of fact.” But in the present case, there is a factual dispute as to what a POSITA would understand based on the disclosures in the ’740 specification. And *Moleculon Research Corp. v. CBS, Inc.*, 793 F.2d 1261, 1270 (Fed. Cir. 1986), does not even discuss indefiniteness, and is therefore completely inapposite.

In addition, Defendants’ reliance on *Versata Software, Inc. v. Zoho Corp.*, 213 F. Supp. 3d 829 (W.D. Tex. 2016), is misplaced. In that case, the court held that the term “space-constrained display” was indefinite because the specification at best provided only a few examples of such displays, and did not define what types of displays would not be “space-constrained.” *See id.* at 836-37. There is no such problem here. The ’740 Patent specifically defines “optimized memory operations” as “increasing the likelihood that, in response to an I/O transaction initiated by a host, the operational load imposed on the storage device by these memory operations will be optimally distributed across different storage device resources, such as by interleaving or parallel memory operations, reducing memory operation latency, increasing operational device efficiency, or both.” ’740 Patent, 2:14-21. The patent also describes multiple ways of achieving this optimal distribution of memory operations. *See id.* at 6:46-7:18. The patent therefore provides sufficient disclosure for a POSITA to understand with reasonable certainty what falls within as well as outside the scope of “optimized memory operations.”

Defendants’ additional arguments regarding claim differentiation also lack merit. Defendants assert that the mapping table in Figure 2 (as well as any similar mapping table with more than two LBA sets) will necessarily have adjacent LBA sets, and so the scope of claim 1 must be the same as claim 2. This makes no sense. Whether a mapping table has adjacent LBA sets or not is not the issue. Rather, the issue is **which** of the LBA sets are involved in the “optimized memory operations.” Claim 1 recites that the claimed mapping table causes “optimized memory operations on memory locations respectively associated with said first PBA and said second PBA, if the I/O transaction request is associated with said first and second LBA sets.” ’740 Patent, 10:19-24. Thus, claim 1 covers optimized memory operations where the I/O transaction request is associated with **any two LBA sets** in the mapping table, regardless of whether they are adjacent or non-adjacent. Claim 1 is therefore broader than claim 2, which is limited to optimized memory operations only if the I/O transaction request involves adjacent LBA sets. *Id.* at 10:25-26 (“2. The mapping table of claim 1, wherein: said first and second LBA sets are adjacent LBA sets . . .”). Thus, there is no claim differentiation problem.

B. “optimally distributed”

Patent	BiTMICRO’s Proposal	Defendants’ Proposal
’740 Patent Claim 32	Plain and ordinary meaning; not indefinite.	Indefinite.

Defendants’ arguments for indefiniteness of “optimally distributed” fail for the same reasons that their arguments for indefiniteness of “optimized memory operations” fail. In addition, Defendants incorrectly argue that “claim 32 does not limit, and therefore extends, its reach to ‘I/O transaction request[s]’ with memory operations that use just one PBA.” Reply at 13. Not so. Defendants ignore the language in claim 32 stating that “said mapping table increases the likelihood that the operational load imposed on the storage apparatus during the processing of the I/O transaction request will be optimally distributed **across storage device resources**.” ’740 Patent, 14:1-5 (emphasis added). The claim language here of distributing “across storage device resources” necessarily means that multiple PBAs are involved in “optimally distributed.” *See id.*

at 9:19-25. Claim 32 is therefore not unduly broad in scope as Defendants suggest. The scope of claim 32 still differs from claim 34, however, in that claim 32 does not specify which PBAs are used for the memory operations, whereas claim 34 specifically identifies memory operations using the recited “said first PBA” and “said second PBA.”

C. “memory store”


Patent	BiTMICRO’s Proposal	Defendants’ Proposal
’740 Patent Claims 1, 15, 32	“non-volatile, solid-state storage, e.g., flash”	Plain and ordinary meaning.

As noted in BiTMICRO’s responsive brief, the following passage in the ’740 specification equates “memory store” with solid-state storage:

Memory store 6 may be configured to include a set of solid state memory devices that are coupled to a set of buses and that are controlled by a set of FDEs. In FIG. 1, these memory devices, buses, and FDEs are implemented in the form of flash memory devices 12-01 and 12-02; buses 8-0 and 8-1; and FDEs 22-01, 22-02, 22-11, and 22-12, respectively. The term “flash memory device” is intended to include any form of non-volatile memory, including those that use blocks of nonvolatile memory cells, named flash blocks.

’740 Patent, 4:54-62. The “may be configured” language in the first sentence does not leave open the possibility of using memory devices other than solid state memory devices, as Defendants contend, but rather leaves open the possibility of other ways in which the solid state memory devices can be connected in the memory store. Defendants argue that this language is not focused on the manner of coupling between the memory devices, but the plain language of the sentence shows otherwise, referring to “solid state memory devices that are *coupled to a set of buses and that are controlled by a set of FDEs.*” *Id.* (emphasis added).

Also, this passage refers to Figure 1, which shows that the memory store is composed of flash storage, as can be seen from the presence of the “Flash Buffer Controller 16” in the figure (highlighted in red):



Id., Fig. 1. In addition, the entire Background section of the '740 Patent is focused on describing the state of the art in “solid state storage devices” and “flash drives,” and problems with existing “flash memory related techniques.” *Id.* at 1:28-2:21. From this context, a POSITA would understand that the '740 invention is directed to an improved system for managing non-volatile, solid-state storage, rather than other types of memory.

Finally, on Defendants’ claim differentiation argument, Defendants incorrectly argue that it does not matter that dependent claims 7, 8, 15, 17 and 19 each add multiple limitations beyond independent claim 1. It is well-established Federal Circuit law that claim differentiation is not a “hard and fast rule of construction,” but rather is most applicable where “there is a dispute over whether a limitation found in a dependent claim should be read into an independent claim, and that limitation is the **only meaningful difference** between the two claims.” *RF Delaware, Inc. v. Pacific Keystone Techs., Inc.*, 326 F.3d 1255, 1263 (Fed. Cir. 2003) (quoting *Wenger Mfg., Inc. v. Coating Machinery Sys., Inc.*, 238 F.3d 1225, 1233 (Fed. Cir. 2001)). None of the dependent claims relied

on by Defendants fit this description, and thus claim differentiation principles do not control here.

IV. '190 PATENT

A. “optimal”

Patent	BiTMICRO’s Proposal	Defendants’ Proposal
'190 Patent Claims 1, 29, 45, 59, 80, 98	Plain and ordinary meaning; not indefinite.	Indefinite.

The '190 Patent makes clear in context what is optimal. Defendants incorrectly and misleadingly suggest that the term “optimal” is indefinite because of the “innumerable situations the device might encounter” and the possible ways in which the device profile can be “optimal.” Reply at 16. Defendants ignore that the term “optimal” in the '190 claims occurs in the specific context of only **two** possible data types – random and sequential. Specifically, the claims recite “wherein said first device profile is optimal for a data type subject to a memory transaction, wherein said data type comprises one of **a random data type or a sequential data type.**” Thus, “optimal” simply means that the first device profile is better for either random or sequential data. And the '190 specification provides specific guidance in the form of at least three different examples of device profiles that are better for handling random or sequential data depending on how the attributes in the device profile are set. See Resp. at 28; '190 Patent, 9:16-43.

Defendants’ response is pure attorney argument—claiming that three examples is not enough guidance to a person of ordinary skill in the art. But this is unsupported by any actual evidence, and as such is not enough to meet Defendants’ burden of proving indefiniteness by clear and convincing evidence. See *Cacace v. Meyer Mktg. (Macau Commercial Offshore) Co.*, 812 F. Supp. 2d 547, 561 (S.D.N.Y. 2011) (“[M]ere attorney argument is insufficient to establish invalidity based on indefiniteness.”).

Defendants’ reliance on *Data Retrieval Technology, LLC v. Sybase, Inc.*, No. C 08-5481 VRW, 2010 WL 11586656 (N.D. Cal. Nov. 8, 2010), is misplaced. Notably, in that case the court found indefiniteness based on expert testimony that the “optimized” database recited in the claim could be achieved by any one of hundreds of possible improvements to the database. *Id.* at *4-5.

Similarly, in *Nabors Drilling Techs. USA, Inc. v. Helmerich & Payne Int’l Drilling Co.*, No. 3:20-CV-03126-M, 2022 WL 1689444, at *14-15 (N.D. Tex. May 26, 2022), the court also found indefiniteness based on expert testimony that the patent failed to provide sufficient guidance as to what would be an “optimized path” toward a drill plan. Both of these cases therefore support BiTMICRO’s position by highlighting how essential expert testimony is to prove indefiniteness. Defendants have no such expert evidence here, and indeed, they offer no evidentiary support whatsoever for their claim that there are “innumerable” ways in which a device profile could be “optimal” such that a person of skill in the art would not understand the scope of that term. Without such supporting evidence, Defendants’ indefiniteness argument necessarily fails.

V. ’084 AND ’694 PATENTS

A. “Direct Memory Access (DMA) descriptors” / “DMA controller descriptor”

Patent	BiTMICRO’s Proposal	Intel’s Proposal
’084 Patent: Claims 1, 19 ’694 Patent: Claims 1, 6, 12	Plain and ordinary meaning.	“data structure including at least a source and destination address and ~~byte count~~ size”

Intel’s Reply Brief discusses Intel’s view on what “a DMA Controller needs . . . to perform a DMA transfer.” This is a separate question from the properly understood plain meaning of DMA Descriptors. And Intel is wrong, as BiTMICRO showed, because there are applications known to a POSITA where, for instance, a second address is not needed in the Descriptor because the source or destination is fixed—or the second address is found elsewhere. Resp. at 32-33; Foty Decl. ¶¶23-24. But even if Intel were correct about the requirements, it does not follow that each and every DMA Descriptor must be limited to containing each and every thing on Intel’s list. There is no support in the patent or intrinsic evidence for this. And as BiTMICRO also showed, there are applications known to a POSITA where multiple Descriptors with different subsets of the items may exist. Resp. at 32 (citing from Dr. Alpert’s own evidence: “we specify *either* the source or the destination, depending on the direction of the transfer”).

Compelled by its expert’s concessions, Intel has switched to “size” from the overly narrow

“byte count.” Reply at 17. This still leaves out the possibility of a DMA system that operates on fixed-size blocks by design, such that size needn’t be spelled out—or with multiple linked descriptors in which only some contain the size. There is no reason to arbitrarily limit the claims to particular applications that Intel puts forward for its own purposes.

Intel’s position here flips the *Phillips* claim interpretation process on its head, proposing arbitrary and improper limitations based on selective citation to extrinsic evidence. BiTMICRO pointed to contrary extrinsic evidence from Intel’s own evidence and documents. Resp. at 32-33. So now Intel tries to fault the patents for not disproving the improper limitation’s premise. Reply at 17 (“the patents do not describe”). This is wrong and should not be credited. Claim construction starts from what a POSITA would understand to be the ordinary meaning in the context of the patent and intrinsic evidence. The patent claims and specification describe what the DMA Descriptors are used for. Resp. at 31. And the patent claims and specification contain no limits on the particular format of those descriptors. *Id.*; Ex. 16 (Alpert Dep. Tr.) at 32:21-33:25:

Q. Now, will you agree with me that there’s nowhere in this patent that actually describes the sort of bitwise format of a descriptor, right?

THE WITNESS: Yes, that is correct. I mean, it’s correct that it doesn’t describe it.

...

Q. Right. So the patent is explicitly saying that there’s not one given format; that a number of possible structures are contemplated, right?

A. Yes.

Intel’s self-serving attempt to add limitations via claim construction to the ordinary meaning of these terms should be rejected.

VI. CONCLUSION

BiTMICRO respectfully requests that the Court adopt BiTMICRO’s proposed constructions and decline Defendants’ invitation to improperly limit the claims or invalidate them for indefiniteness.

Dated: December 28, 2022

Respectfully submitted,

/s/ Michael Flynn-O'Brien

Denise M. De Mory (*Pro Hac Vice*)
California State Bar No. 168076
Corey Johanningmeier (*Pro Hac Vice*)
California State Bar No. 251297
Richard C. Lin (*Pro Hac Vice*)
California State Bar No. 209233
Michael Flynn-O'Brien (*Pro Hac Vice*)
California State Bar No. 291301
Bunsow De Mory LLP
701 El Camino Real
Redwood City, CA 94063
(650) 351-7241 Telephone
(415) 426-4744 Facsimile
ddemory@bdiplaw.com
cjohanningmeier@bdiplaw.com
rlin@bdiplaw.com
mflynno'Brien@bdiplaw.com

***Attorney in Charge for Plaintiff BiTMICRO
LLC***

B. Russell Horton
State Bar No. 10014450
George Brothers Kincaid & Horton, L.L.P.
114 West 7th Street, Ste. 1100
Austin, Texas 78701
(512) 495-1400 Telephone
(512) 499-0094 Facsimile
rhorton@gbkh.com

Attorney for Plaintiff BiTMICRO LLC

CERTIFICATE OF SERVICE

The undersigned hereby certifies that a true and correct copy of the above and foregoing document has been served on December 28, 2022 to all counsel of record who are deemed to have consented to electronic service via the Court's CM/ECF system.

/s/ Michael Flynn-O'Brien
Michael Flynn-O'Brien (*Pro Hac Vice*)

APPENDIX A

'243 Patent, col. 8:61 – 9:22 prior to Certification of Correction	Markup with changes in Certificate of Correction	'243 Patent, col. 8:61 – 9:22 after changes in Certification of Correction
FIG. 19 shows a first pin assignment and connection technique that enables a selection of specific module in a multiple module stack. The figure illustrates the controller and SDRAM modules as an example. The SDRAM module balls, such as 1907, are represented by big oblong shapes, and the pads such as 1908 are represented by smaller oblong shapes. A pad to ball connection is represented by a rectangle, such as 1906. The controller module has four active pads such as 1904, one each for the four SDRAM modules such as 1901. The controller module's active pad 00 1904 is connected to SDRAM module 0's X0 ball 1903, pad 01 to SDRAM module 1's X0 ball 1902, and so on. The X0 balls are the active balls for the SDRAM modules. All SDRAM modules are identical which include the connections as follows: X1 ball is connected to X1 pad, X2 ball to X2 pad, and so on. The technique involves using a ladder-like routing path, such as 1905 on the stacked modules. This technique	FIG. 19 shows a first pin assignment and connection technique that enables a selection of specific module in a multiple module stack. The figure illustrates the controller and SDRAM modules as an example. The SDRAM module balls, such as 1907, are represented by big oblong shapes, and the pads such as 1908 are represented by smaller oblong shapes. A pad to ball connection is represented by a rectangle, such as 1906. <u>An active pad may be also referred to herein as an "active port."</u> The controller module has four active pads such as 1904, one each for the four SDRAM modules such as 1901. The controller module's active pad 00 1904 is <u>ultimately</u> connected to ~~SDRAM module 0's X0 ball 1903, pad 01 ball 1903 of~~ <u>SDRAM module 0, active pad 01 is connected to SDRAM module 1's X0 ball 1902</u> X0 ball 1902 of SDRAM module 1, and so on. These X0 balls are the active balls ~~for the SDRAM modules for these~~ <u>SDRAM modules. Thus, as illustrated in Figure 19, there is one active ball, such as X0, ultimately coupled to one active pad per module.</u>	FIG. 19 shows a first pin assignment and connection technique that enables a selection of specific module in a multiple module stack. The figure illustrates the controller and SDRAM modules as an example. The SDRAM module balls, such as 1907, are represented by big oblong shapes, and the pads such as 1908 are represented by smaller oblong shapes. A pad to ball connection is represented by a rectangle, such as 1906. An active pad may be also referred to herein as an "active port." The controller module has four active pads such as 1904, one each for the four SDRAM modules such as 1901. The controller module's active pad 00 1904 is ultimately connected to X0 ball 1903 of SDRAM module 0, active pad 01 is connected to X0 ball 1902 of SDRAM module 1, and so on. These X0 balls are the active balls for these SDRAM modules. Thus, as illustrated in Figure 19, there is one active ball, such as X0, ultimately coupled to one active pad per module.

'243 Patent, col. 8:61 – 9:22 prior to Certification of Correction	Markup with changes in Certificate of Correction	'243 Patent, col. 8:61 – 9:22 after changes in Certification of Correction
enables the controller module's active pad to be routed to the desired specific module in the stack. The example embodiment illustrated in FIG. 19 comprises a repeating pattern for the balls and pads. Passive balls are connected within each module to passive pads with an offset distance equaling one periodic distance of the repeating pattern. The importance is the approach using repeating pattern and offsetting of the passive ball to passive pad within each module, other offset distances can be used also in other embodiment of the present invention, e.g. using an offset of two or any multiple of the periodic distance of the repeating pattern. The example embodiment does not limit the scope of the present invention.	All SDRAM modules are ~~identical~~ <u>identical, and in the example shown in FIG. 19, include the connections as follows: include the following: an X1 passive ball is connected to an X1 passive pad, an X2 passive ball to X2 passive pad, and so on. The connection between a passive ball on one surface of a SDRAM module and a passive pad on another surface of the same SDRAM module, such as the X1 passive ball and X1 passive pad connection in FIG. 19, is named "passive port", as as passive ports 1909 illustrated in FIG. 19.</u> The technique involves using a ladder-like routing path, such as 1905 on the stacked modules. This technique enables the controller module's active pad to be routed to the desired specific module in the stack. <u>Thus, an active signal applied at an active pad, such as active pad 02, shown with reference number 1904 on controller module 1910, is routed through at least one passive port connecting a passive ball and passive pad, such as the passive ports 1905 connecting passive ball X2 to passive pad X2 and passive ball X1 to passive pad X2, respectively, so that this active signal eventually reaches an active ball, such as X0 ball SDRAM Module 2, in</u>	All SDRAM modules are identical, and in the example shown in FIG. 19, include the following: an X1 passive ball is connected to an X1 passive pad, an X2 passive ball to X2 passive pad, and so on. The connection between a passive ball on one surface of a SDRAM module and a passive pad on another surface of the same SDRAM module, such as the X1 passive ball and X1 passive pad connection in FIG. 19, is named "passive port", as as passive ports 1909 illustrated in FIG. 19. The technique involves using a ladder-like routing path, such as 1905 on the stacked modules. This technique enables the controller module's active pad to be routed to the desired specific module in the stack. Thus, an active signal applied at an active pad, such as active pad 02, shown with reference number 1904 on controller module 1910, is routed through at least one passive port connecting a passive ball and passive pad, such as the passive ports 1905 connecting passive ball X2 to passive pad X2 and passive ball X1 to passive pad X2, respectively, so that this active signal eventually reaches an active ball, such as X0 ball SDRAM Module 2, in Figure 19. The example embodiment illustrated in

'243 Patent, col. 8:61 – 9:22 prior to Certification of Correction	Markup with changes in Certificate of Correction	'243 Patent, col. 8:61 – 9:22 after changes in Certification of Correction
	Figure 19. The example embodiment illustrated in FIG. 19 comprises a repeating pattern for the balls and pads. Passive balls are connected within each module to passive pads with an offset distance equaling one periodic distance of the repeating pattern. <u>For example in Figure 19, an X1 passive ball is connected to an X1 passive pad, X2 passive ball to X2 passive pad, etc.</u> ~~The importance is the approach using~~ By using repeating pattern and offsetting of the passive ball to passive pad within each module, other offset distances can be used ~~also in other alternative embodiments of the present invention, e.g. using invention.~~ <u>For example, an alternative embodiment may include using an offset of two or any multiple of the periodic distance of the repeating pattern. The example embodiment does not limit This alternative embodiment is not intended to limit the scope of the present invention in any way.</u>	FIG. 19 comprises a repeating pattern for the balls and pads. Passive balls are connected within each module to passive pads with an offset distance equaling one periodic distance of the repeating pattern. For example in Figure 19, an X1 passive ball is connected to an X1 passive pad, X2 passive ball to X2 passive pad, etc. By using repeating pattern and offsetting of the passive ball to passive pad within each module, other offset distances can be used in alternative embodiments of the present invention. For example, an alternative embodiment may include using an offset of two or any multiple of the periodic distance of the repeating pattern. This alternative embodiment is not intended to limit the scope of the present invention in any way.