

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

REALTEK SEMICONDUCTOR CORP.,
Petitioner,
v.

ADVANCED MICRO DEVICES, INC.
Patent Owner.

Case No. IPR2023-00687

U.S. Patent No. 8,854,381

**PETITIONER'S OPPOSITION TO PATENT OWNER'S
CONTINGENT MOTION TO AMEND**

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. LEGAL STANDARDS	1
III. SUBSTITUTE CLAIMS 28 AND 29 IMPERMISSIBLY ADD NEW MATTER AND SHOULD BE REJECTED UNDER 35 U.S.C. §112	2
A. To Satisfy the Written Description Requirement, Each Claim Element Must Be Expressly or Inherently Disclosed in the Specification, Not Just Rendered Obvious.....	2
B. Proposed Claim 28 Introduces New Matter and Does Not Satisfy The Written Description Requirement of 35 U.S.C. §112	4
B. Proposed Claim 29 Introduces New Matter and Does Not Satisfy The Written Description Requirement of 35 U.S.C. §112	12
III. SUBSTITUTE CLAIM 28 IS OBVIOUS IN LIGHT OF JIAO.....	15
A. Claim 28	15
1. Claim 28 (Preamble): The computing system of claim 15, wherein the shader core comprises:	15
2. Claim element 28[a]: <i>multiple processing elements</i> ;.....	15
3. Claim element 28[b]: a first plurality of processing elements configured to execute the first task based on a first set of state data associated with the first task.....	16
4. Claim element 28[c]: a second plurality of processing elements configured to execute the second task based on a second set of state data associated with the second task, ...	16
5. Claim element 28[d]: wherein each of the plurality of tasks is decomposable into two or more related threads	16
6. Claim element 28[e]: and wherein at least one processing element is configured to receive tasks from each engine.	17
IV. SUBSTITUTE CLAIMS 28 AND 29 ARE OBVIOUS OVER WILT	19
A. Claim 28	19
7. Claim 28 (Preamble): The computing system of claim 15, wherein the shader core comprises:	19

8.	Claim element 28[a]: <i>multiple processing elements</i> ;	19
9.	Claim element 28[b]: a first plurality of processing elements configured to execute the first task based on a first set of state data associated with the first task	19
10.	Claim element 28[c]: a second plurality of processing elements configured to execute the second task based on a second set of state data associated with the second task,	19
11.	Claim element 28[d]: wherein each of the plurality of tasks is decomposable into two or more related threads	19
12.	Claim element 28[e]: and wherein at least one processing element is configured to receive tasks from each engine.	21
C.	Claim 29	23
1.	Claim 29 (Preamble). <i>The computing system of claim 15</i> ,	23
2.	Claim element 29[a]: wherein the plurality of engines comprises a first queue configured to receive the first task and a second queue configured to receive the second task,	23
3.	Claim element 29[b]: wherein the plurality of engines are arranged in parallel with one another,	23
4.	Claim element 29[c]: wherein each engine is configured to receive at least one of the plurality of tasks via a distinct command buffer for each engine,	23
5.	Claim element 29[d]: wherein each of the plurality of tasks is decomposable into two or more related threads	24
6.	Claim element 29[e]: wherein each engine is configured to load state data associated with at least one of the plurality of tasks.	24

PETITIONER'S EXHIBIT LIST

Exhibit #	Reference Name
1001	U.S. Patent 8,854,381 to Lefebvre et al.
1002	Prosecution History of U.S. Patent 8,854,381
1003	Declaration of Massoud Pedram, Ph.D.
1004	U.S. Patent No. 8,151,095 to Wilt ("Wilt")
1005	U.S. Patent No. 8,319,774 to Jiao et al. ("Jiao")
1006	PCT Pub. No. WO00/62182A2 to Stuttard et al. ("Stuttard")
1007	Curriculum Vitae of Massoud Pedram, Ph.D.
1008	<i>Certain Graphics Systems, Components Thereof, and Digital Televisions Containing The Same</i> , Inv. No. 337-TA-1318 (Order No. 31, Oct. 18, 2022)
1009	Deposition Transcript of Dr. John Hart
1010	Declaration of Dr. John Hart dated August 15, 2022
1011	ITC claim construction ruling (Order No. 31)
1012	Declaration of Dr. Massoud Pedram
1013	AMD Brief (177 Proceeding related to 1318 ITC investigation)
1014	Declaration of Dr. Massoud Pedram

TABLE OF AUTHORITIES

	Page(s)
CASES	
<i>Apple Inc. v. Qualcomm Inc.</i> , IPR2018-01275, Paper 44 (PTAB February 12, 2020)	4
<i>Ariad Pharms., Inc. v. Eli Lilly & Co.</i> , 598 F.3d 1336 (Fed. Cir. 2010)	2, 4
<i>Corus Realty Holdings, Inc. v. Zillow Grp., Inc.</i> , 860 F. App'x 728 (Fed. Cir. 2021)	5
<i>In re Man Mach. Interface Techs. LLC</i> , 822 F.3d 1282 (Fed. Cir. 2016)	5
<i>Juno Therapeutics, Inc. v. Kite Pharma, Inc.</i> , 10 F.4th 1330 (Fed. Cir. 2021)	6
<i>Lockwood v. American Airlines, Inc.</i> , 107 F.3d 1565 (Fed. Cir. 1997)	2, 3, 4, 14
<i>Martin v. Mayer</i> , 823 F.2d 500 (Fed. Cir. 1987)	3
<i>Nautilus, Inc. v. Icon Health & Fitness</i> , IPR2017-01408, Paper 51 (PTAB December 3, 2018)	4
<i>PowerOasis, Inc. v. T-Mobile USA, Inc.</i> , 522 F.3d 1299 (Fed. Cir. 2008)	3, 14
<i>SuperGuide Corp. v. DirecTV Enterprises, Inc.</i> , 358 F.3d 870 (Fed. Cir. 2004)	6
<i>Tronzo v. Biomet</i> , Inc., 156 F.3d 1154, 1159 (Fed. Cir. 1998)	3
<i>TurboCare Div. of Demag Delaval Turbomach. v. Gen. Elec. Co.</i> , 264 F.3d 1111 (Fed. Cir. 2001)	1, 3

OTHER AUTHORITIES

37 C.F.R. § 42.6(a).....	29
37 C.F.R. § 42.6(e).....	28
37 C.F.R. §42.24(a).....	29
37 C.F.R. §42.121(a)(2).....	1
37 C.F.R. §42.121(b)(1).....	1
<i>In re Jones</i> , 10 Fed. Appx 822, 828 (Fed. Cir. 2001).....	14

I. INTRODUCTION

AMD's Contingent Motion to Amend should be denied. As a threshold issue, proposed claims 28 and 29 add new matter and fail to satisfy the written description requirement. AMD contends that a POSITA would recognize that its substitute claims represent one "possible" way of implementing what is disclosed in the original application. But this falls well short of satisfying §112, which requires that each claim limitation be actually or inherently disclosed in the original application. The written description requirement cannot be satisfied by mere "probabilities or possibilities" and a disclosure is inadequate "that merely renders the invention obvious."

While obviousness is inapplicable to the written description analysis, obviousness *does* render the claims unpatentable over the prior art. Proposed claim 28 is rendered unpatentable by Jiao and claims 28 and 29 are unpatentable over Wilt. AMD's Motion to Amend should also be denied for these additional reasons.

II. LEGAL STANDARDS

A motion to amend "may not enlarge the scope of the claims of the patent or introduce new matter." 35 U.S.C. § 316(d)(3). *Accord*, 37 C.F.R. §42.121(a)(2). New matter is any addition to the claims that lacks written description support in the original application. *TurboCare Div. of Demag Delaval Turbomach. v. Gen. Elec. Co.*, 264 F.3d 1111, 1118 (Fed. Cir. 2001); 37 C.F.R. §42.121(b)(1).

The disclosure in the original application must “reasonably convey[] to those skilled in the art that the inventor had possession of the claimed subject matter as of the filing date” to satisfy the written description requirement. *Ariad Pharms., Inc. v. Eli Lilly & Co.*, 598 F.3d 1336, 1351 (Fed. Cir. 2010). “Possession” of the claimed subject matter must be shown by descriptive means such as words, structures, and figures that set forth fully the claimed invention. *Lockwood v. American Airlines, Inc.*, 107 F.3d 1565, 1572 (Fed. Cir. 1997). “It is *not sufficient* for purposes of the written description requirement of §112 that the disclosure, when combined with the knowledge in the art, would lead one *to speculate* as to modifications that the inventor might have envisioned, but failed to disclose.” *Id.* (emphasis added).

III. SUBSTITUTE CLAIMS 28 AND 29 IMPERMISSIBLY ADD NEW MATTER AND SHOULD BE REJECTED UNDER 35 U.S.C. §112

A. To Satisfy the Written Description Requirement, Each Claim Element Must Be Expressly or Inherently Disclosed in the Specification, Not Just Rendered Obvious

AMD’s Motion to Amend should be denied because claim elements 28[e] and 29[e] add new matter and fail to satisfy the written description requirement. For example, AMD and Dr. Hart argue that if various disclosures of the original application are “combined” they “*indicate that one way*” to perform a recited function is to use the configuration required by new claim element 28[e]. Paper 13 at 9. Similarly, AMD and Dr. Hart admit that the configuration required by new

claim element 29[e] is only “**one possible design** based on the disclosure.” *Id.* at

11. These statements effectively admit that the substitute claims do not satisfy the written description requirement.

To satisfy the written description requirement, each limitation of a claim must be expressly or inherently disclosed in the specification, which does not include obvious potential permutations of the disclosed embodiments:

While the meaning of terms, phrases, or diagrams in a disclosure is to be explained or interpreted from the vantage point of one skilled in the art, ***all the limitations must appear in the specification. The question is not whether a claimed invention is an obvious variant of that which is disclosed in the specification.***

Lockwood, 107 at 1572 (emphasis added). The written description requirement is “not a question of whether one skilled in the art *might* be able to construct the patentee's device from the teachings of the disclosure. ... Rather, it is a question whether the application necessarily discloses that particular device.” *Martin v. Mayer*, 823 F.2d 500, 505 (Fed. Cir. 1987). “This requires that the written description actually or inherently disclose the claim element.” *PowerOasis, Inc. v. T-Mobile USA, Inc.*, 522 F.3d 1299, 1306-07 (Fed. Cir. 2008), *citing TurboCare*, 264 F.3d at 1118–20 (the claimed subject matter must at least be inherent in the original disclosure, *i.e.*, it must “necessarily be present”); *Tronzo v. Biomet, Inc.*, 156 F.3d 1154, 1159 (Fed. Cir. 1998) (where claim limitations were generic to the

shape of a recited “cup,” the written description requirement was not met because the original application did not disclose all cup shapes literally or “inherently”).

“One shows that one is ‘in possession’ of the invention by describing the invention, with all its claimed limitations, not that which makes it obvious.” *Lockwood*, 107 F.3d at 1572. *Accord*, *Ariad*, 598 F.3d at 1352 (a disclosure is inadequate “that merely renders the invention obvious”). AMD may not “pick and choose” elements of different disclosures or embodiments and combine them to show that the proposed substitute claim would have been obvious based on the disclosure. *Apple Inc. v. Qualcomm Inc.*, IPR2018-01275, Paper 44 at 40-44 (PTAB February 12, 2020). *See also* *Nautilus, Inc. v. Icon Health & Fitness*, IPR2017-01408, Paper 51 at 57 (PTAB December 3, 2018) (“obviousness is not the standard for written description”).

B. Proposed Claim 28 Introduces New Matter and Does Not Satisfy The Written Description Requirement of 35 U.S.C. §112

AMD proposes to substitute claim 28 for claim 16, which depends from claim 15. Substitute claim 28 is as follows:

28. The computing system of claim 15, wherein the shader core comprises:

- [a] multiple processing elements;
- [b] a first plurality of processing elements configured to execute the first task based on a first set of state data associated with the first task; and

- [c] a second plurality of processing elements configured to execute the second task based on a second set of state data associated with the second task,
- [d] wherein each of the plurality of tasks is decomposable into two or more related threads
- [e] and *wherein at least one processing element is configured to receive tasks from each engine.*

Claim elements 28[b] and [c] are recited in current claim 16, whereas claim elements 28[a], [d] and [e] are added. Proposed claim element 28[e] introduces new matter and lacks written description support. Claim element 28[e] requires that “at least one” processing element is “configured to” receive multiple tasks from “each engine.” However, there is no disclosure of processing elements being specifically configured to receive two or more tasks from specific engines, much less from *each* engine. Ex. 1014, ¶11. As AMD admits, it relies on an alleged “*combination*” of disclosures that only “*indicate*” that this might be “*one way*” of implementing the invention. This falls far short of satisfying the written description requirement.


The claim language “configured to” means “designed to” and requires more than mere capability or the mere possibility or statistical probability that something could occur. *Corus Realty Holdings, Inc. v. Zillow Grp., Inc.*, 860 F. App’x 728, 733 (Fed. Cir. 2021); *see also In re Man Mach. Interface Techs. LLC*, 822 F.3d 1282, 1286 (Fed. Cir. 2016) (“designed to” or “configured to” has a narrower

meaning than “capable of”); *Aspex Eyewear*, 672 F.3d at 1349 (same). The plain meaning of “at least one” is “one or more.” *SuperGuide Corp. v. DirecTV Enterprises, Inc.*, 358 F.3d 870, 886 (Fed. Cir. 2004). Accordingly, to satisfy the written description requirement, the original disclosure must expressly or necessarily disclose that one or more processing elements are specifically designed to receive multiple tasks from *each* of the “plurality of engines.”¹ It does not.

Claim 28 is directed to the “shader core,” which the ’381 patent defines as “an array of processing elements.” Ex. 1001 at 1:59-61. In the ’381 patent, the “array of processing elements” of the shader core is basically a black box. The patent contains no disclosure of any specific arrangement of the processing elements that the patentee had in his possession at the time of the invention, especially in relation to the claimed “plurality of engines.” For example, Figures 3A, 3B and 4 only show a “shader core” box without any processing elements or

¹ The written description requirement requires that the specification sufficiently demonstrate that the inventors possessed the full scope of the claimed invention. *Juno Therapeutics, Inc. v. Kite Pharma, Inc.*, 10 F.4th 1330, 1336 (Fed. Cir. 2021). Thus, the specification must provide written description support for a single processing element being configured to receive multiple tasks from each engine, as well as a subset of the processing elements. The specification supports neither.

other components. The only Figure that provides *any* detail of the processing elements 220 is Figure 2:



However, Figure 2 does not show *any* configuration between any processing element (or set of processing elements) and any engine, much less *each* engine. Figure 2 only shows a two-dimensional array of a scalable number of processing elements. These disclosures indicate that the patentee had *no intended configuration at all* for how any individual processing element or groups of processing elements relate to each engine.

The rest of the specification fares no better. AMD cites to Paragraphs 27-32, 35, 50-54 of the original application, but nothing in those paragraphs supports claim element 28[e]. *See, e.g.*, Ex. 1002 at 6-35, ¶¶ 27-29 (no disclosure of the processing elements), 31-32, 50, 53-54 (no disclosure of the relationship of processing elements to engines), 35 (disclosing only that a plurality of engines share resources of a single shader core).

AMD appears to rely on paragraphs 30 and 50-52 of the original application, which describe two ways to partition the shader core to “simultaneously *process* tasks from two or more different virtual engines.” Such partitioning only relates to

how the shader core can *process* two tasks from two different engines simultaneously (*e.g.*, how two tasks might be issued to processing elements within the shader core for simultaneous execution). This partitioning for simultaneous processing does not expressly or inherently disclose *any* configuration for the *receiving* of tasks by any processing elements from *any* engine, much less a configuration to receive a plurality of tasks from *each* engine.


But even if the description of how tasks are issued to processing elements within the shader core somehow suggests a relationship between the processing elements and each of the engines—and it does not—the partitioning described fails to support claim element 28[e] for additional reasons. The first type is called “spatial partitioning”: “a first ... task from a first virtual engine is issued to a first set of processing elements (SIMDs) of the shader core, and a second ... task from a second virtual engine is issued to a second set of processing elements (SIMDs) of the shader core.” Ex. 1002 ¶30. Because “spatial partitioning” involves *different* sets of processing elements executing tasks from *different* engines, it has no possible relevance to claim element 28[e]’s configuration of at least one processing element receiving tasks from *each* engine.

The other type of partitioning is called “temporal partitioning” and is described as “different processing tasks of different virtual engines share a percentage of time of processing elements 220 of shader core 214.” *Id.* ¶ 52. *See*


id. ¶ 30. Temporal partitioning is disclosed as one way for the shader core to “simultaneously process tasks from two or more different virtual engines.”

Because temporal partitioning is an example of *simultaneous* execution of multiple tasks from *different* engines, it necessarily *excludes* Dr. Hart’s examples showing the entire shader core processing tasks from one engine at one point in time, and a different engine at a different point in time, which is also inconsistent with the invention’s goal of efficient multithreaded processing. Ex. 2008 ¶¶ 36-38.² Once again, the partitioning described does not necessarily disclose the configuration required by claim element 28[e].

In the alternative, Dr. Hart embarks on an elaborate color-coding exercise to show another processing scenario that could *hypothetically* arise, as a matter of statistical possibility, during operation:



Actual disclosure in Figure 2



Dr. Hart’s made-up figure

² This hypothetical scenario also fails to show how a single processing element or a subset of processing elements are “configured to” receive tasks from each engine.

Ex. 1001, Fig. 2 (excerpt); Ex. 2008 ¶ 40 (colors representing tasks from different engines). As this color coding exercise inherently demonstrates, Dr. Hart’s hypothetical scenario is neither expressly nor inherently disclosed in the specification. In this hypothetical, the first and second columns of processing elements only receive tasks from one engine, while the third column of processing elements *potentially* receives tasks from different engines over time. But the ’381 patent does not expressly or inherently disclose such a configuration—that is just *one way* that different tasks *might* be processed, over time, as a matter of *possibility or probability*. For example, it is also possible that the first two columns of processing elements always receive tasks from the “blue” and “green” engines, with the result that *no* processing element receives tasks from *each* engine. There is simply no disclosure in the ’381 patent of any design or configuration in which one or more processing elements are specifically configured or designed to *receive* multiple tasks from *each* engine.

AMD tries to gloss over this lack of written description by arguing that “[c]ombined, these all indicate that *one* way to execute tasks simultaneously ... is for the processing elements of the shader core to receive and execute tasks from one engine at one time, and to receive and execute tasks from another engine at another time, and that over time the processing elements will receive tasks from each engine.” Paper 13, 8-9. But AMD cannot point to any express or necessary

disclosure of such a configuration—it is merely *one possible way* in which tasks from the engines *might be* received by processing elements, as a matter of probability. This does not satisfy the written description requirement.

AMD’s argument is essentially one of capability. As discussed above, the claim language “configured to” means “designed to” and requires more than mere capability or the mere possibility that something could occur. Regardless of whether the disclosed shader core has the *capability* to receive tasks from each engine, there is no disclosure of one or more processing elements being *configured to* receive multiple tasks from each engine.

The lack of written description for proposed claim 28 does not stop there. The “at least one processing element” of claim element 15[e] could be in addition to “the first plurality of processing elements” and “the second plurality of processing elements” of claim elements 15[b] and [c]. Moreover, the “tasks” recited in claim element 15[e] need not be part of “the plurality of tasks” of independent claim 15 and may be independent of the “first task” and the “second task” of claim 15 and claim elements 15[b] and [c]. There is no disclosure of any of these additional configurations, all of which are within the scope of claim 28.

For each of the foregoing reasons, proposed claim 28 impermissibly adds new matter and lacks written description.

B. Proposed Claim 29 Introduces New Matter and Does Not Satisfy The Written Description Requirement of 35 U.S.C. §112

AMD proposes to substitute claim 29 for claim 17, which depends from claim 15. Substitute claim 29 is as follows:

29. The computing system of claim 15,
- [a] wherein the plurality of engines comprises a first queue configured to receive the first task and a second queue configured to receive the second task,
 - [b] wherein the plurality of engines are arranged in parallel with one another,
 - [c] wherein each engine is configured to receive at least one of the plurality of tasks via a distinct command buffer for each engine,
 - [d] wherein each of the plurality of tasks is decomposable into two or more related threads, and
 - [e] *wherein each engine is configured to load state data associated with at least one of the plurality of tasks.*

Claim element 29[a] is current claim 17, whereas claim elements 29[b], [c], [d] and [e] are added.

Proposed claim element 29[e] introduces new matter and lacks written description support. There is no disclosure in the application for the '381 patent of a configuration where *each* engine is *configured to* (designed to) load state data associated with at least one of the plurality of tasks. The only disclosure is that *the plurality of engines—as a whole*—loads state data. The disclosure never refers to

each engine loading state data, nor is such a requirement necessarily disclosed.

Even AMD admits that claim element 29[e] is only a “*possible* design” that the patentee might have “*intended* to convey.” This falls far short of establishing that the inventor was in possession of the invention at the time of application.

The specification only twice refers to state data being loaded by the plurality of engines. The Abstract states, “The *plurality* of virtual engines is configured to ... load a set of state data associated with each of the plurality of tasks.” Ex. 1001, Abstract. This identical statement is repeated, without any further detail, in the Brief Summary. Ex. 1001, 3:3-6. The operation is mentioned one other time, but without reference to *any* of the engines performing the operation. *Id.*, 3:33-34 (“In a second operation, a set of state data associated with each of the plurality of tasks is loaded.”). There is no other disclosure.

All other disclosures in the specification refer to the loading of state data within the shader core (*e.g.*, into the LDS), which relates to the execution of tasks. Ex. 1002, Abstract and ¶¶ 5-8, 13, 16, 53. The state data loaded by the plurality of engines in claim element 15[d] is distinct from the state data loaded by the shader core, which is separately claimed in claim element 15[e]. None of the Figures show any configuration of engines related to the loading of state data.

Accordingly, the specification only discloses that *the plurality* of engines loads state data, which can be satisfied by *only one* of the plurality of engines

loading state data. Since the claimed element “each engine is configured to load state data” is neither expressly nor inherently disclosed, the written description requirement is not satisfied.

As explained *supra* at Sections II. and III.A., it is well-established that to satisfy the written description requirement, each limitation of a claim must be expressly or inherently disclosed in the specification. *See Lockwood*, 107 at 1572; *PowerOasis*, 522 F.3d at 1306-07. For example, in *In re Jones*, the Federal Circuit agreed with the Board that the term “helical channel with closed ends” lacked written description because while it was ***possible*** for the channel to have closed ends, nothing in the specification expressly disclosed closed ends or indicated that closed ends were necessarily present. 10 Fed. Appx 822, 828 (Fed. Cir. 2001). The written description requirement cannot be satisfied by “probabilities or possibilities.” *Id.*

AMD’s motion to amend ***admits*** that there is no written description of ***each*** engine of the invention loading state data:

“As ***a POSITA would understand*** this statement is ***intended to convey*** that each engine loads the state data associated with the task it receives and transmits to the shader core. At a minimum, a POSITA would know that was one ***possible*** design ***based on*** the disclosure.”

Paper 13 at 11-12 (citations omitted) (emphasis added). The fact that it is only a ***“possible”*** design based on the disclosure admits what the specification makes

clear—that the claimed design is neither expressly nor inherently disclosed and the written description requirement is not met.

Accordingly, there is no disclosure that supports claim element 15[e], “wherein each engine is configured to load state data associated with at least one of the plurality of tasks.” Proposed substitute claim 29 impermissibly adds new matter and does not satisfy the requirements of 35 U.S.C §112.

III. SUBSTITUTE CLAIM 28 IS OBVIOUS IN LIGHT OF JIAO

A. Claim 28

1. Claim 28 (Preamble): *The computing system of claim 15, wherein the shader core comprises:*

For all of the reasons provided in the Petition and Petitioner’s Reply, the claim elements of claim 15 are obvious in light of Jiao.

2. Claim element 28[a]: *multiple processing elements;*

For all of the reasons provided in the Petition and Petitioner’s Reply with respect to claim 16, Jiao renders obvious a shader core having multiple processing elements. *See* Ex. 1005, Fig. 4, 1:43-47, 8:15-54, 9:33-35. Specifically, Jiao’s computational core “comprises one or more execution units (EUs) 420a through 420h.” *Id.*

3. Claim element 28[b]: *a first plurality of processing elements configured to execute the first task based on a first set of state data associated with the first task*

For all of the reasons provided in the Petition and Petitioner's Reply, Jiao renders obvious this claim element, which is already recited in claim 16.

4. Claim element 28[c]: *a second plurality of processing elements configured to execute the second task based on a second set of state data associated with the second task,*

For all of the reasons provided in the Petition and Petitioner's Reply, Jiao renders obvious this claim element, which is already recited in claim 16.

5. Claim element 28[d]: *wherein each of the plurality of tasks is decomposable into two or more related threads*

AMD's proposed construction for "task" in its POR is that a task must be at least two or more related threads. POR, 21-23; Ex. 1005, ¶¶ 26-27 ("a POSITA would know, a task comprises more than one thread"). As explained in Petitioner's Reply, AMD is wrong. If true, however, then all of the tasks disclosed in the cited prior art (including Jiao and Wilt) would necessarily be decomposable into two or more related threads for processing in a multi-threaded processor. A task comprising two or more threads must be decomposable into two or more related threads or it cannot be executed. Stated differently, element 28[d] is inherent in every "task" under AMD's definition of a "task."

In addition, the computational core in Jiao has many parallel EUs, and each EU "is capable of multi-threaded operations simultaneously." Ex. 1005, Fig. 4,


1:43-47, 8:15-54, 9:33-35. The “engines” of Jiao (e.g., the vertex shader, geometry shader, and pixel shader) deliver tasks to the execution units 420, and “the EU pool control unit 206 handles the assigning of those tasks to available threads within the various execution units 420.” Ex. 1005, 9:25-54. The EU pool control unit “is responsible for assigning ... tasks to threads of the various execution units 420, and also, performs an associated ‘bookkeeping’ of the tasks and threads.” *Id.* Because “tasks” can encompass a very wide range of programmable operations to meet the needs of different computer programs (e.g., *id.*, 5:47-51), it is inherent, or at least obvious to a POSITA, that each of a plurality of “tasks” is decomposable into two or more related threads. *See also* Paper 1, 43; Ex. 1003, ¶103 (testifying that a POSITA “would have understood that Jiao’s disclosure inherently discloses a higher level task that has given rise to the low level thread”).

6. Claim element 28[e]: *and wherein at least one processing element is configured to receive tasks from each engine.*

As explained in the Petition, the claimed “plurality of engines” corresponds to the vertex shader 304, geometry shader 306, and pixel shader 310 in computational core 204. Paper 1, 41. The “shader core” of Jiao consists of multiple execution units (EUs) arranged in parallel. “While eight (8) EUs 420 are shown or implied in FIG. 4 (labeled EU0 through EU7), it should be appreciated that the number of EUs need not be limited to eight, but may be greater or fewer in number for some embodiments.” Ex. 1005, 8:15-27, Fig. 4. Thus, each EU 420 is

a “processing element” in the shader core of Jiao.

Jiao discloses or at least renders obvious that at least one EU (a processing element) is configured to receive tasks from each engine. For example, Jiao discloses that as “various shader tasks are generated by the vertex shader 304, geometry shader 306, and pixel shader 310, they are delivered to the respective execution units 420 to be carried out.” Ex. 1005, 9:35-39. Specifically, “the EU pool control unit 206 is responsible for assigning vertex shader, geometry shader, and pixel shader tasks to threads of the various execution units 420,” depending on “how much free space is available for each execution unit.” *Id.*, 9:39-10:9. Since the EU pool control unit assigns tasks from *each* of the plurality of engines to *each* of the Execution Units (e.g., eight EUs), and each Execution Unit is a “processing element,” it is obvious that at least one processing element (EU) is configured to receive tasks from each engine. This is also shown in Figure 4, in which each processing element (EU0 to EU7, comprising 8 processing elements) is shown as receiving tasks from each of the three engines VS 304, GS 306, and PS 310. Fig. 4 (excerpt annotated with one obvious possibility). Accordingly, Jiao at least renders obvious substitute claim 28.



IV. SUBSTITUTE CLAIMS 28 AND 29 ARE OBVIOUS OVER WILT

A. Claim 28

7. **Claim 28 (Preamble):** *The computing system of claim 15, wherein the shader core comprises:*

See discussion of claim 15 for Wilt in the Petition and Petitioner's Reply.

8. **Claim element 28[a]:** *multiple processing elements;*

See discussion of claim 16 for Wilt in the Petition and Petitioner's Reply.

Specifically, claim 16 already requires two pluralities of processing elements with the claimed shader core. *See* Fig. 2 and Ex. 1004, 4:1 (multiple PPUs, each of which is “a highly parallel processor,” contains Core 0 ... Core C-1, each of which is capable of executing tens or hundreds of threads concurrently).

9. **Claim element 28[b]:** *a first plurality of processing elements configured to execute the first task based on a first set of state data associated with the first task*

See discussion of claim 16 for Wilt in the Petition and Petitioner's Reply.

10. **Claim element 28[c]:** *a second plurality of processing elements configured to execute the second task based on a second set of state data associated with the second task,*

See discussion of claim 16 for Wilt in the Petition and Petitioner's Reply.

11. **Claim element 28[d]:** *wherein each of the plurality of tasks is decomposable into two or more related threads*

As explained above with respect to Jiao, element 28[d] is inherent in every “task” as disclosed in Wilt under AMD’s definition of “task” as comprising more

than one thread. In addition, each PPU in Wilt has multiple multithreaded processing cores that receive and process “tasks.” Ex. 1004, 4:1-19. “Each processing core 208 is capable of executing a large number (e.g., tens or hundreds) of threads concurrently.” *Id.* “Cores 208 receive processing tasks to be executed.” *Id.* “When a new processing task arrives, work distribution unit 210 assigns the task to a core 208.” *Id.* “Cores 208 can be programmed to execute processing tasks relating to a wide variety of applications, including but not limited to linear and nonlinear data transforms, filtering of video and/or audio data, modeling operations (e.g., applying laws of physics to determine position, velocity and other attributes of objects), image rendering operations (e.g., vertex shader, geometry shader, and/or pixel shader programs), and so on.” *Id.*, 4:33-40. Such “tasks” encompass a very wide range of programmable operations. It is therefore inherent, or at least obvious to a POSITA, that when a processing “task” is assigned to a multithreaded processing “core” with tens or hundreds of threads for execution, that each of a plurality of “tasks” is decomposable into two or more related threads.

Regardless, whether a task is decomposable into one thread or “two or more” threads is merely a matter of degree or complexity of the task at hand and trivially obvious. A POSITA would have understood that while some tasks may be very simple and decompose into a single thread, most tasks will be more complex and decompose into two or more threads. Ex. 2007, 11:7-16. A POSITA would

have understood that graphics processing, as described in Wilt, would never be comprised of *only* single threaded tasks, which is why the processing cores in Wilt are designed to execute tasks using tens or hundreds of threads, as explained above. Ex. 1014 ¶33. *See also* Ex. 1004, 5:59-6:6 (describing embodiments where a single instruction results in parallel execution of a large number of threads). *See also id.*, 3:60-4:25, 4:61-5:4, 6:13-27, 8:15-29, 10:21-51).

12. Claim element 28[e]: *and wherein at least one processing element is configured to receive tasks from each engine.*

Wilt's front end 212, work distribution unit 210, and core interface 303 within each PPU collectively serve as a single "engine," and the set of cores 208 within and across all the PPUs (PP1 ... PPU U-1) collectively serve as the "shader core." *See* Ex. 1004, Fig. 2-3. Wilt discloses that PPUs can communicate with each other through the bidirectional communication path 113 and the memory bridge 105. *See* Ex. 1004, 4:7-23 and Fig. 3. Wilt explains that "Work distribution unit 210 can implement a variety of algorithms for distributing work." *Id.* Wilt gives an example of how each core 208 provides a "ready" signal when it has sufficient resources to accept a new task, and that "[w]hen a new processing task arrives, work distribution unit 210 assigns the task to a core 208 that is asserting the ready signal." *Id.* A work distribution unit can "hold" a task if "no core 28" is asserting a ready signal. *Id.* Wilt expressly states that "those skilled in the art will

recognize that other algorithms may also be used [for distributing tasks to the cores] and that the particular manner in which work distribution unit 210 distributes incoming processing tasks is not critical to the present invention.” *Id.*

Because bidirectional communication path 113 and memory bridge 105 connect each engine within each PPU, a POSITA would understand from Wilt’s explicit teaching that when no cores in a PPU have available resources, that its work distribution unit 210 may send a processing task back through the said communication path and memory bridge so that it can run on a different PPU which has cores signaling “ready.” *Id.*; Ex. 2007, 42:21-45:6. In this case, each of the cores 208 of a PPU is configured to receive tasks not only from the single engine associated with that PPU but also from each engine of every other PPU. *Id.*; Ex. 2007, 42:21-45:6. The bidirectional distribution of tasks is further suggested by Wilt’s teaching that “each PPU 202 includes a host interface that communicates with the rest of system 100 via communication path 113” and the Host interface 206 can “generate packets ... for transmission on communication path 113” just as it “also receives all incoming packets ... from communication path 113.” Ex. 1004, 3:49-67. Such a configuration is at least obvious, as it is a possible configuration taught by Wilt. Accordingly, Wilt at least renders obvious “at least one processing element is configured to receive tasks from each engine.”

C. Claim 29

1. Claim 29 (Preamble). *The computing system of claim 15,*

See discussion of claim 15 for Wilt in the Petition and Petitioner's Reply.

2. Claim element 29[a]: *wherein the plurality of engines comprises a first queue configured to receive the first task and a second queue configured to receive the second task,*

See discussion of claim 17 for Wilt in the Petition and Petitioner's Reply.

3. Claim element 29[b]: *wherein the plurality of engines are arranged in parallel with one another,*

Wilt's front end 212, work distribution unit 210, and core interface 303 of each PPU collectively serve as a single "engine," while cores 208 within and across all the PPUs collectively serve as the "shader core." In Wilt's system, the PPUs (PPU 0 ... PPU U-1) are arranged in parallel with one another, and the "one or more parallel processing units (PPUs) are part of the "parallel processing subsystem." Ex. 1004, Fig. 2, 2:27-28, 3:36-48. Because there is a single "engine" within each PPU, the plurality of engines in Wilt are arranged in parallel.

4. Claim element 29[c]: *wherein each engine is configured to receive at least one of the plurality of tasks via a distinct command buffer for each engine,*

Wilt's host interface 206 in each PPU satisfies the claimed "distinct command buffer" for each "engine." Ex. 1004, 3:56-64. In the '381 patent, the "command buffer" is disclosed as receiving tasks from the "scheduling module" and providing those tasks to the "engines." *See* Ex. 1001, 2:8-9, 10:25-39 and Fig.

4. This is disclosed by Wilt. Within each PPU 202, the host interface 206 receives all incoming packets (the claimed “tasks”) from the application software running on the CPU 102 (the claimed “scheduling module”) via the communication path 113.³ Ex. 1004, 3:49-67. The host interface then sends the received tasks to the “engine” within each PPU (*i.e.*, the front end unit 212, work distribution unit 210, and core interface 303 of each PPU). “For example, commands related to processing tasks may be directed to a front end unit 212.” *Id.*, 3:60-61. Therefore, the host interface and corresponding “engine” in each PPU discloses or renders obvious the claimed “wherein each engine is configured to receive at least one of the plurality of tasks via a distinct command buffer for each engine.”

5. Claim element 29[d]: *wherein each of the plurality of tasks is decomposable into two or more related threads*

See discussion above with respect to claim element 28[d].

6. Claim element 29[e]: *wherein each engine is configured to load state data associated with at least one of the plurality of tasks*

A POSITA would have understood that each of Wilt’s “engines” (within

³ The “packets (or other signals)” refer to data packets and control signals for the communication path 113 (which is a PCI-E link in a disclosed embodiment in Wilt.) A POISTA would understand that these packets contain the “commands related to [defining] processing tasks.” *See* Ex. 1004, 3:49-67.

each PPU) are configured to load state data associated with at least one of the plurality of tasks. Ex. 1003, ¶74.⁴ Wilt uses the phrase “state parameters” to refer to state data. *Id.*; Ex. 1004 at 10:25-32. In Wilt, tasks are received by the host interface 206 in each PPU, which passes the tasks to the front end unit 212, then the work distribution unit 210, and finally, the core interface unit 303, which loads the *state parameters* (state data) associated with the task before passing the tasks to the cores 208 for processing. Paper 1, 27. “In some embodiments, core interface 303 receives data to be processed (e.g., primitive data, vertex data, and/or pixel data) as well as *state parameters* and commands defining how the data is to be processed (e.g., what program is to be executed) from work distribution unit 210. Core interface 303 can *load* data to be processed into shared memory 306 and *[state] parameters* into parameter memory 308.” Ex. 1004, 8:15-21. Thus, where each “engine” is comprised of the front end unit 212, work distribution unit 210, and core interface unit 303 within each PPU, each “engine” is configured to load state data associated with at least one of the plurality of tasks it receives.

⁴ The '381 patent broadly defines “state data” as including, for example, “texture handles, shader constants, transform matrices, or the like.” Ex. 1001, 1:53-55.

Dated: April 10, 2024

Respectfully submitted,

/Jeffrey L. Johnson/

Jeffrey L. Johnson
Reg. No. 53,078
BAKER BOTTS, LLP
910 Louisiana Street
Houston, TX 77002-4995
Telephone: 713-229-1234
Facsimile: 713-229-1522
Email: Jeffrey.Johnson@bakerbotts.com
DLBBRealTEKIPRs@bakerbotts.com

Christopher J. Higgins
Reg. No. 66,422
ORRICK, HERRINGTON & SUTCLIFFE LLP
Columbia Center
2100 Pennsylvania Avenue NW
Washington, D.C. 20037
Telephone: (202) 339-8400
Facsimile: (202) 339-8500
Email: 0CHPTABDocket@orrick.com

Steve Baik
Reg. No. 42,281
WHITE HAT LEGAL
P. O. Box 382
San Jose, CA 95002
Telephone: (650) 618-5282
Email: sbaik@whitehat.legal

*Counsel for Petitioner Realtek
Semiconductor Corp.*

CERTIFICATE OF SERVICE

In accordance with 37 C.F.R. § 42.6(e), the undersigned certifies that on April 10, 2024, a complete and entire copy of this **PETITIONER'S OPPOSITION TO PATENT OWNER'S CONTINGENT MOTION TO AMEND** was served on AMD via electronic mail at the following addresses:

William A. Meunier
wameunier@mintz.com

Michael T. Renaud
mtrenaud@mintz.com

Adam S. Rizk
arizk@mintz.com

Reza Dokhanchy
rdokhanchy@mintz.com

AMD_IPRS@mintz.com

/Jeffrey L. Johnson/
Jeffrey L. Johnson (Reg. No. 53,078)
Lead Counsel for Petitioner

**CERTIFICATE OF COMPLIANCE WITH TYPE-VOLUME
LIMITATION, TYPEFACE REQUIREMENTS, AND TYPE STYLE
REQUIREMENTS**

1. This **Petitioner's Opposition to Patent Owner's Contingent Motion to Amend** complies with the page limit of 25 pages, excluding the parts excepted by 37 C.F.R. §42.24(a).

2. This **Petitioner's Opposition to Patent Owner's Contingent Motion to Amend** complies with the general format requirements of 37 C.F.R. § 42.6(a) and has been prepared using Microsoft® Word in 14-point Times New Roman.

Respectfully submitted,

BAKER BOTTS L.L.P.

/Jeffrey L. Johnson
Jeffrey L. Johnson (Reg. No. 53,078)
Lead Counsel for Petitioner

Date: April 10, 2024