

UNITED STATES INTERNATIONAL TRADE COMMISSION

Washington, D.C.

In the Matter of

**CERTAIN GRAPHICS SYSTEMS,
COMPONENTS THEREOF, AND
DIGITAL TELEVISIONS CONTAINING
THE SAME**

Inv. No. 337-TA-1318

**INITIAL DETERMINATION ON VIOLATION OF SECTION 337 AND
RECOMMENDED DETERMINATION ON REMEDY AND BOND**

Administrative Law Judge Cameron Elliot

(July 7, 2023)

Pursuant to the Notice of Investigation and Rule 210.42(a) of the Rules of Practice and Procedure of the United States International Trade Commission, this is my Initial Determination in the matter of *Certain Graphics Systems, Components Thereof, and Digital Televisions Containing the Same*, Investigation No. 337-TA-1318.

TABLE OF CONTENTS

Page

I.	INTRODUCTION.....	1
A.	Procedural Background.....	1
B.	The Parties	4
C.	The Asserted Patents and Claims.....	5
D.	Products at Issue	6
1.	Domestic Industry Products.....	6
2.	Accused Products.....	9
II.	STANDARDS OF LAW	15
A.	Standing	15
B.	Claim Construction	15
C.	Infringement.....	18
D.	Domestic Industry	21
1.	Technical Prong	21
2.	Economic Prong.....	22
E.	Invalidity	24
1.	35 U.S.C. § 101.....	24
2.	35 U.S.C. § 102.....	24
3.	35 U.S.C. § 103.....	25
III.	IMPORTATION AND JURISDICTION	27
IV.	U.S. PATENT NO. 7,742,053	31
A.	Level of Ordinary Skill in the Art.....	31
B.	Claims-at-Issue	31
C.	Claim Construction	32
1.	“at least one memory device . . . operative to store a plurality of pixel command threads and . . . a plurality of vertex command threads”.....	32
2.	“select a command thread”	39
3.	“command processing engine”	41
D.	Infringement.....	43
1.	Claim 5.....	43
a.	[5a] “command thread”	43
b.	[5a] “at least one memory device . . .”.....	50

	c.	[5b] “an arbiter . . .”	57
	d.	[5c] “a plurality of command processing engines . . .”	61
2.		Claim 6	65
3.		Claim 9	66
4.		Indirect Infringement	66
E.		Domestic Industry – Technical Prong	67
1.		Claim 5	68
	a.	[5a] “at least one memory device . . .”	68
	b.	[5b] “an arbiter . . .”	73
	c.	[5c] “a plurality of command processing engines . . .”	78
2.		Claim 6	80
3.		Claim 9	80
F.		Validity	80
1.		Claim 5	81
2.		Claim 6	88
3.		Claim 9	88
V.		U.S. PATENT NO. 8,854,381	89
A.		Level of Ordinary Skill in the Art	89
B.		Claims-at-Issue	89
C.		Claim Construction	90
	1.	“a plurality of engines configured to receive . . . a plurality of tasks and to load state data associated with each of the plurality of tasks”	91
	2.	“configured to . . . execute a first task . . . while executing a second task”	95
	3.	“each of the plurality of tasks”	98
D.		Infringement	100
	1.	Claim 15	100
	a.	[15a] “a memory”	101
	b.	[15d] “a plurality of engines configured to receive, from a scheduling module associated with a second processing unit, a plurality of tasks and to load state data associated with each of the plurality of tasks”	103

	c.	[15e] “a shader core configured to receive the plurality of tasks from at least one of the plurality of engines and to execute a first task from the plurality of tasks while executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks”	113
	2.	Claim 16.....	121
	3.	Claim 17.....	123
	4.	Claim 19.....	123
	5.	Claim 20.....	124
	6.	Indirect Infringement	125
E.		Domestic Industry – Technical Prong.....	125
	1.	Claim 15.....	126
	a.	[15d] “a scheduling module associated with a second processing unit”	126
	b.	[15d] “a plurality of engines . . . to load state data associated with each of the plurality of tasks”	132
	c.	[15e] “a shader core configured to receive the plurality of tasks from at least one of the plurality of engines and to execute a first task from the plurality of tasks while executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks”	136
	2.	Claim 16.....	137
	3.	Claim 17.....	137
	4.	Claim 19.....	138
	5.	Claim 20.....	138
F.		Validity	139
	1.	Jiao	139
	a.	Claim 15.....	140
	b.	Claim 16.....	159
	c.	Claim 17.....	160
	d.	Claim 19.....	161
	e.	Claim 20.....	163
	2.	Stuttard.....	164
	a.	Claim 15.....	164

	b.	Claim 16.....	171
	c.	Claim 17.....	172
	d.	Claim 19.....	172
	e.	Claim 20.....	173
	3.	Patentable Subject Matter	174
VI.		DOMESTIC INDUSTRY - ECONOMIC PRONG	179
VII.		CONCLUSIONS OF LAW	180
VIII.		RECOMMENDED DETERMINATION ON REMEDY AND BOND	181
	A.	Limited Exclusion Order.....	182
	B.	Cease and Desist Order.....	183
	C.	Bond.....	186
IX.		INITIAL DETERMINATION AND ORDER	186

TABLE OF ABBREVIATIONS

CDX	Complainant's Demonstrative Exhibit
CIB	Complainant's Revised Initial Post-Hearing Brief
CPB	Complainant's Pre-Hearing Brief
CPX	Complainant's Physical Exhibit
CRB	Complainant's Reply Post-Hearing Brief
CX	Complainant's Exhibit
Hr'g Tr.	Hearing Transcript
JX	Joint Exhibit
RDX	Respondents' Demonstrative Exhibit
RIB	Respondents' Initial Post-Hearing Brief
RPB	Respondents' Pre-Hearing Brief
RPX	Respondents' Physical Exhibit
RRB	Respondents' Reply Post-Hearing Brief
RX	Respondents' Exhibit

I. INTRODUCTION

A. Procedural Background

Complainants Advanced Micro Devices, Inc. and ATI Technologies ULC (together, “AMD” or “Complainants”) filed the complaint underlying this investigation on May 5, 2022. The complaint alleged respondents TCL Industries Holdings Co., Ltd.; TCL Industries Holdings (H.K.) Limited; TCL Electronics Holdings Limited f/k/a TCL Multimedia Technology Holdings, Ltd.; TCL Technology Group Corporation; TTE Corporation; TCL Holdings (BVI) Limited; TCL King Electrical Appliances (Huizhou) Co., Ltd.; Shenzhen TCL New Technologies Co., Ltd.; TCL MOKA International Limited; TCL Smart Device (Vietnam) Co., Ltd.; Manufacturas Avanzadas SA de CV; TCL Electronics Mexico, S de RL de CV; and TCL Overseas Marketing Ltd. (together, “TCL”) and respondent Realtek Semiconductor Corp. (“Realtek”) import or sell in connection with an importation certain graphics systems and televisions containing the same that infringe one or more claims of U.S. Patent Nos. 7,742,053 (“the 053 patent”), 8,760,454 (“the 454 patent”), 11,184,628 (“the 628 patent”), 8,468,547 (“the 547 patent”), and 8,854,381 (“the 381 patent”) (together, the “Asserted Patents”).

By publication of a notice in the *Federal Register* on June 7, 2022, the U.S. International Trade Commission ordered that:

Pursuant to subsection (b) of section 337 of the Tariff Act of 1930, as amended, an investigation be instituted to determine whether there is a violation of subsection (a)(1)(B) of section 337 in the importation into the United States, the sale for importation, or the sale within the United States after importation of products identified in paragraph (2) by reason of infringement of one or more of claims 1-9 of the '053 patent; claims 2-11 of the '454 patent; claims 7-12 of the '628 patent; claims 16-21 of the '547 patent; and claims 15-20 of the '381 patent, and whether an industry in the United States exists or [is] in the process of being established as required by subsection (a)(2) of section 337[.]

87 Fed. Reg. 34718 (June 7, 2022).

[REDACTED]

On June 1, 2022, the investigation was assigned to Chief Administrative Law Judge Clark Cheney. EDIS Doc. ID 772059. On June 22, 2022, the investigation was reassigned to me. EDIS Doc. ID 773680.

On June 28, 2022, I set a target date of November 7, 2023 for completion of this investigation. Order No. 6. Also on June 28, 2022, I set a *Markman* hearing date of September 6, 2022, and the evidentiary hearing for March 13-17, 2023. Order No. 7. The *Markman* hearing occurred on September 6, 2022, with the parties' disputes resolved on October 18, 2022. Order No. 31.

On July 11, 2022, AMD moved (1318-003) for partial termination of the investigation by withdrawal of claims 2-11 of the 454 patent, thereby terminating the patent from the investigation. The motion was granted via initial determination on July 14, 2022. Order No. 10. On August 4, 2022, the Commission determined not to review Order No. 10. EDIS Doc. ID 777095.

On August 26, 2022, non-party TTE Technology, Inc. moved (1318-011) to intervene in the investigation. The motion was granted via initial determination on August 30, 2022 with TTE Technology, Inc. entered as an additional respondent (collectively, with all others, "Respondents"). Order No. 17. On September 27, 2022, the Commission determined not to review Order No. 17. EDIS Doc. ID 781075.

On September 16, 2022, AMD moved again (1318-016) for partial termination of the investigation by withdrawal of claims 17-21 of the 547 patent. The motion was granted via initial determination on September 20, 2022. Order No. 23. Also on September 16, 2022, AMD moved (1318-017) to correct the names of two of the TCL respondents. This motion was granted via initial determination on September 20, 2022. Order No. 24. On October 11, 2022, the Commission determined not to review Order Nos. 23 and 24. EDIS Doc. ID 781939.

[REDACTED]

On October 4, 2022, AMD moved (1318-024) for summary determination of effective service of the complaint and notice of investigation upon Realtek. The motion was granted via initial determination on October 19, 2022. Order No. 33. On November 18, 2022, the Commission determined not to review Order No. 33. EDIS Doc. ID 784871.

On January 13, 2023, AMD moved (1318-047) for summary determination that the economic prong is satisfied and that respondents have not met their burden regarding a license defense. On February 1, 2023, the portion of the motion regarding the license defense was denied. Order No. 59. On February 6, 2023, the portion of the motion concerning economic prong was granted via initial determination. Order No. 62. On March 15, 2023, the Commission determined to review and, upon review, to affirm Order No. 62. EDIS Doc. ID 792553.

Also on January 13, 2023, Realtek moved (1318-049) for summary determination that the 547 patent is invalid and of no importation for certain accused products. On January 30, 2023, the motion was denied. Order No. 57.

On January 23, 2023, AMD moved a third time (1318-050) for partial termination of the investigation by withdrawal of claim 16 of the 547 patent, thereby terminating that patent from the investigation. The motion was granted via initial determination on January 24, 2023. Order No. 56. On February 22, 2023, the Commission determined not to review Order No. 56. EDIS Doc. ID 790860.

On February 6, 2023, AMD moved a fourth time (1318-056) for partial termination of the investigation by withdrawal of claims 1-4 and 7 of the 053 patent and claims 8, 11, and 12 of the 628 patent. The motion was granted via initial determination on February 7, 2023. Order No. 64. On March 7, 2023, the Commission determined not to review Order No. 64. EDIS Doc. ID 791955.

[REDACTED]

On March 9, 2023, AMD moved a fifth time (1318-065) for partial termination of the investigation by withdrawal of claim 8 of the 053 patent and claim 18 of the 381 patent. The motion was granted via initial determination on March 14, 2023. Order No. 70. On March 30, 2023, the Commission determined not to review Order No. 70. EDIS Doc. ID 793435.

On March 31, 2023, AMD moved a sixth time (1318-068) for partial termination of the investigation by withdrawal of claims 7, 9, and 10 of the 628 patent, thereby terminating the 628 patent from the investigation. The motion was granted via initial determination on April 3, 2023. Order No. 72. On April 19, 2023, the Commission determined not to review Order No. 72. EDIS Doc. ID 794577.

The evidentiary hearing took place on March 13-17, 2023. Pursuant to the procedural schedule and Order No. 71, the parties submitted initial and reply post-hearing briefs on April 7, 2023 and April 21, 2023, respectively. On May 23, 2023, the parties filed a joint motion (1318-071) for leave to file a statement on a certain issue in the investigation, which remains pending. The motion (1318-071) is granted, but its contents need not be considered given the resolution of that issue below.

As of the date of this initial determination, no other motions remain pending.

B. The Parties

Complainant Advanced Micro Devices, Inc. is a Delaware corporation, with a principal place of business in Santa Clara, CA. CIB at 3. Complainant ATI Technologies ULC is incorporated in Canada, with a principal place of business in Ontario, Canada, and is a “wholly-owned indirect subsidiary” of Advanced Micro Devices, Inc. *Id.* Together, the companies, according to an AMD witness, “provide world class processors to the marketplace. So that is CPUs, central processing units; graphics processing units, GPUs; accelerated processing unit[s],

[REDACTED]

or APUs; and then more recently . . . FPGAs, field programmable gate arrays; and DPUs as well.” Hr’g Tr. (O’Neil) at 21:17-21. AMD explains that Advanced Micro Devices, Inc. is the sole owner of the 381 patent, and ATI Technologies ULC is the sole owner of the 053 patent. CIB at 3.

Respondent Realtek Semiconductor Corporation “is headquartered in Taiwan and primarily engages in the research, development, and the sale of various types of integrated circuits.” RIB at 2. Respondents add, “Realtek’s other products include software and hardware application design, testing, repairs, and technical consultations to various types of integrated circuits; the sale of various types of silicon intellectual property; and adjunct trade and sales related to its core businesses.” *Id.* at 3. The remaining respondents, collectively referred to as TCL, “are located in China, Hong Kong, Vietnam, and Mexico.” RIB at 3. According to Respondents, they “design, manufacture, distribute and sell televisions and other consumer electronic products.” *Id.*

C. The Asserted Patents and Claims

The 053 patent, entitled “Multi-Thread Graphics Processing System,” issued on June 22, 2010 to Laurent LeFebvre, Andrew Gruber, and Stephen Morein. The 053 patent reports an assignment on its face to ATI Technologies ULC and claims priority to an application filed September 29, 2003.

The 381 patent, entitled “Processing Unit that Enables Asynchronous Task Dispatch,” issued on October 7, 2014 to Michael Mantor and Rex McCrary. The 381 patent reports an assignment on its face to Advanced Micro Devices, Inc. and claims priority to a provisional application filed September 3, 2009.

The Asserted Patents cover apparatuses, methods, and systems pertaining to graphics processing and parallel computing. The claims of the 053 patent encompass “graphics processing system[s]” comprising hardware components that perform operations involving “pixel command

Contains Confidential Business Information Subject to Protective Order

threads” and “vertex command threads.” *See* 053 patent at cls. 1, 5. The patent explains, “[i]n a graphics processing system, it is important to manage and control multiple command threads relating to texture applications.” *Id.* at 1:22-24. The claims of the 381 patent encompass apparatuses, systems, and methods “for enabling asynchronous task dispatch,” particularly in the processing of graphics data. *See* 381 patent at 1:20, 2:66. The patent discloses a plurality of virtual engines which receive tasks and load state data associated with those tasks, and then a shader core which executes those tasks. *Id.* at Abstract.

The following patent claims are presently at issue in this investigation, according to AMD’s briefing:

Asserted Patent	Infringement Claims	Domestic Industry Claims
7,742,053	5, 6, 9	5, 6, 9
8,854,381	15-17, 19, 20	15-17, 19, 20

See generally CIB at 3, 24, 48, 83, 106.

D. Products at Issue

1. Domestic Industry Products

The domestic industry products in this investigation are graphics processing units (GPUs) designed and sold by AMD which include specialized circuitry and programming to render three-dimensional objects on a two-dimensional screen. The GPUs can be sold individually, or combined with CPUs in integrated circuit (IC) packages, or on printed circuit boards. Hr’g Tr. (Reinman) at 100:13-101:6; Hr’g Tr. (McCrary) at 39:7-40:18. GPUs are generally classified by their architecture and in this investigation each of AMD’s domestic industry products is [REDACTED]. CIB at 7. These architectures are further broken down into microarchitectures. [REDACTED]

Contains Confidential Business Information Subject to Protective Order

██████████. *Id.* AMD describes its domestic industry products as “AMD GPUs, printed circuit boards containing the same, and all supporting software, drivers (including graphics libraries and APIs), and/or frameworks,” and presents them as a list in Appendix A to its initial post-hearing brief, reproduced below (altogether, the “DI Products”):

[illegible]

[illegible]

AMD has met its burden here. Dr. Reinman has offered credible testimony, with persuasive demonstrative charts and pursuant to his personal review of the source code, that the [REDACTED]. *See* Hr’g Tr. (Reinman) at 101:7-103:24. He further testified the [REDACTED] microarchitecture is representative of the group, again through his source code analysis. *Id.* at 180:4-19. Respondents’ source code investigator and expert, Mr. Erik de la Iglesia, agrees. He testified, “Q. Is it your

[REDACTED]

opinion that the source code of the AMD domestic industry products produced by AMD are

? A. Oh, with respect to those aspects that I was asked to investigate, yes.” Hr’g Tr. (de la Iglesia) at 860:6-10.

Dr. Pfister on the other hand did not personally investigate any source code (Hr’g Tr. (Pfister) at 771:24-772:6) and his testimony demonstrated a general unfamiliarity with the product line. For example, he seemed to believe that [REDACTED] “is also called the [REDACTED]” (*id.* at 732:25-733:1) even though [REDACTED] is only one of many [REDACTED] microarchitectures. Dr. Pfister’s testimony also acknowledges the findings of Mr. de la Iglesia, but dismisses its relevance because Mr. de la Iglesia only reviewed what Dr. Pfister instructed him to review. *Id.* at 734:2-7 (“so I directed him to the source code I wanted him to review. So it was not our burden to show that these are representative architectures.”). Even if true, the source code Mr. de la Iglesia did review showed substantial similarity between products.

Dr. Pfister admittedly refers to a variety of technical differences between the product lines, including differences of

. Hr’g Tr. (Pfister) at 734:8-19. Yet he offered no explanation as to how these differences might affect an infringement analysis. So Dr. Pfister’s opinions are properly assigned less weight than, and do not overcome, the combined opinions of Dr. Reinman and Mr. de la Iglesia, which were based on direct source code review.

Accordingly, for the purposes of this investigation, the [REDACTED] is accepted as representative of all DI Products.

2. Accused Products

The accused products in this investigation consist of GPUs, GPUs incorporated onto printed circuit boards, and digital televisions that include those GPUs from Realtek and TCL. CIB at 7. The TCL televisions contain GPUs supplied by Realtek and [REDACTED]. None of

Contains Confidential Business Information Subject to Protective Order

Realtek, [REDACTED], or TCL design their GPUs, however, which is instead done by [REDACTED]. CIB at 8. More specifically, according to AMD, the Realtek and [REDACTED] accused products all incorporate [REDACTED]. See *generally id.* at 8-13. These are further broken down into microarchitectures. [REDACTED]. *Id.* at 7. [REDACTED]. *Id.* AMD presents the list of the accused products in Appendix A to its initial post-hearing brief, which is reproduced below (“Accused Products”) (“Realtek Accused Products”) (“TCL Accused Products”):

[illegible]

[illegible][illegible]

Contains Confidential Business Information Subject to Protective Order

Age Group	Don't know	No	Yes	Yes, definitely
18-24	10%	10%	20%	60%
25-34	10%	10%	25%	45%
35-44	10%	10%	20%	60%
45-54	10%	10%	20%	60%

[illegible]

Contains Confidential Business Information Subject to Protective Order

[illegible][illegible]

CIB, App'x A; *see* RIB at 9-10 (citing JX-0011C; JX-0012C).

[REDACTED]

As with the DI Products, AMD makes several assertions on the similarity between Accused Products to streamline its infringement analysis. Specifically, AMD contends, “Realtek and

[REDACTED]

[REDACTED]

[REDACTED] “Realtek and [REDACTED]

[REDACTED]

[REDACTED] and “Realtek and [REDACTED]

[REDACTED]

[REDACTED]” CIB at 8-9 (citations omitted). AMD then refers to the testimony of Dr. Reinman and Mr. de la Iglesia regarding their extensive source code review and demonstrative diagrams to argue “[t]he evidence has shown that, with limited exceptions that are expressly stated, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]” *Id.* at 9-10 (citations omitted).

Respondents do not meaningfully dispute these claims of representativeness, apart from arguing “there are substantive differences between different models and architectures of GPUs,

[REDACTED]

[REDACTED]” RRB at 1-2. Respondents issue the same caveat as AMD, which is that where these differences affect reading the claims onto the Accused Products, they have been specifically identified and argued in following sections. *See id.* at 2.

Accordingly, subject to specific differences discussed explicitly below, [REDACTED]

[REDACTED] is accepted as representative of all Accused Products, and [REDACTED]

[REDACTED]

[REDACTED] is accepted as representative of [REDACTED]
[REDACTED].

II. STANDARDS OF LAW

A. Standing

Commission Rule 210.12 states in relevant part “[f]or every intellectual property based complaint (regardless of the type of intellectual property involved), [the complaint must] include a showing that at least one complainant is the owner or exclusive licensee of the subject intellectual property.” 19 C.F.R. § 210.12(a)(7). In determining whether this rule is met, the Commission looks to the standing requirement used by courts in patent infringement cases. *Certain Audio Processing Hardware, Software, and Products Containing the Same*, Inv. No. 337-TA-1026, Comm’n Op. at 9 (April 18, 2018) (citations omitted).

B. Claim Construction

“The construction of claims is simply a way of elaborating the normally terse claim language in order to understand and explain, but not to change, the scope of the claims.” *Embrex, Inc. v. Serv. Eng’g Corp.*, 216 F.3d 1343, 1347 (Fed. Cir. 2000). Although most of the disputed claim terms were construed in an earlier order, some of the issues presented below are only resolvable with additional claim construction.

Claim construction focuses on the intrinsic evidence, which consists of the claims themselves, the specification, and the prosecution history. *See Phillips v. AWH Corp.*, 415 F.3d 1303, 1314 (Fed. Cir. 2005) (en banc); *see also Markman v. Westview Instr., Inc.*, 52 F.3d 967, 979 (Fed. Cir. 1995) (en banc). As the Federal Circuit in *Phillips* explained, courts must analyze each of these components to determine the “ordinary and customary meaning of a claim term” as understood by a person of ordinary skill in art at the time of the invention. 415 F.3d at 1313. “Such intrinsic evidence is the most significant source of the legally operative meaning of disputed

claim language.” *Bell Atl. Network Servs., Inc. v. Covad Commc'ns Grp., Inc.*, 262 F.3d 1258, 1267 (Fed. Cir. 2001).

“It is a ‘bedrock principle’ of patent law that ‘the claims of a patent define the invention to which the patentee is entitled the right to exclude.’” *Phillips*, 415 F.3d at 1312 (quoting *Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1115 (Fed. Cir. 2004)). “Quite apart from the written description and the prosecution history, the claims themselves provide substantial guidance as to the meaning of particular claims terms.” *Id.* at 1314; see *Interactive Gift Express, Inc. v. Compuserve Inc.*, 256 F.3d 1323, 1331 (Fed. Cir. 2001) (“In construing claims, the analytical focus must begin and remain centered on the language of the claims themselves, for it is that language that the patentee chose to use to ‘particularly point [] out and distinctly claim [] the subject matter which the patentee regards as his invention.’”). The context in which a term is used in an asserted claim can be “highly instructive.” *Phillips*, 415 F.3d at 1314. Additionally, other claims in the same patent, asserted or unasserted, may also provide guidance as to the meaning of a claim term. *Id.* “Courts do not rewrite claims; instead, we give effect to the terms chosen by the patentee.” *K-2 Corp. v. Salomon S.A.*, 191 F.3d 1356, 1364 (Fed. Cir. 1999). “[T]he specification ‘is always highly relevant to the claim construction analysis. Usually it is dispositive; it is the single best guide to the meaning of a disputed term.’” *Phillips*, 415 F.3d at 1315 (quoting *Vitronics Corp. v. Conceptiontronic, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996)). “[T]he specification may reveal a special definition given to a claim term by the patentee that differs from the meaning it would otherwise possess. In such cases, the inventor’s lexicography governs.” *Id.* at 1316.

In addition to the claims and the specification, the prosecution history should be examined, if in evidence. *Phillips*, 415 F.3d at 1317; see *Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898,

[REDACTED]

913 (Fed. Cir. 2004). The prosecution history can “often inform the meaning of the claim language by demonstrating how the inventor understood the invention and whether the inventor limited the invention in the course of prosecution, making the claim scope narrower than it would otherwise be.” *Phillips*, 415 F.3d at 1317; *see Chimie v. PPG Indus. Inc.*, 402 F.3d 1371, 1384 (Fed. Cir. 2005) (“The purpose of consulting the prosecution history in construing a claim is to exclude any interpretation that was disclaimed during prosecution.”).

When the intrinsic evidence does not establish the meaning of a claim, then extrinsic evidence (*i.e.*, all evidence external to the patent and the prosecution history, including dictionaries, inventor testimony, expert testimony, and learned treatises) may be considered. *Phillips*, 415 F.3d at 1317. Extrinsic evidence is generally viewed as less reliable than the patent itself and its prosecution history in determining how to define claim terms. *Id.* “The court may receive extrinsic evidence to educate itself about the invention and the relevant technology, but the court may not use extrinsic evidence to arrive at a claim construction that is clearly at odds with the construction mandated by the intrinsic evidence.” *Elkay Mfg. Co. v. Ebco Mfg. Co.*, 192 F.3d 973, 977 (Fed. Cir. 1999).

The construction of a claim term is generally guided by its ordinary meaning. However, courts may deviate from the ordinary meaning when: (1) “the intrinsic evidence shows that the patentee distinguished that term from prior art on the basis of a particular embodiment, expressly disclaimed subject matter, or described a particular embodiment as important to the invention;” or (2) “the patentee acted as his own lexicographer and clearly set forth a definition of the disputed claim term in either the specification or prosecution history.” *Edwards Lifesciences LLC v. Cook Inc.*, 582 F.3d 1322, 1329 (Fed. Cir. 2009); *see Apple Inc. v. Wi-LAN Inc.*, 25 F.4th 960, 967 (Fed. Cir. 2022); *GE Lighting Sols., LLC v. AgiLight, Inc.*, 750 F.3d 1304, 1309 (Fed. Cir. 2014) (“the

[REDACTED]

specification and prosecution history only compel departure from the plain meaning in two instances: lexicography and disavowal.”); *Omega Eng’g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1324 (Fed. Cir. 2003) (“[W]here the patentee has unequivocally disavowed a certain meaning to obtain his patent, the doctrine of prosecution disclaimer attaches and narrows the ordinary meaning of the claim congruent with the scope of the surrender.”); *Rheox, Inc. v. Entact, Inc.*, 276 F.3d 1319, 1325 (Fed. Cir. 2002) (“The prosecution history limits the interpretation of claim terms so as to exclude any interpretation that was disclaimed during prosecution.”). Nevertheless, there is a “heavy presumption that a claim term carries its ordinary and customary meaning.” *CCS Fitness, Inc. v. Brunswick Corp.*, 288 F.3d 1359, 1366 (Fed. Cir. 2002) (citations omitted). The standard for deviating from the plain and ordinary meaning is “exacting” and requires “a clear and unmistakable disclaimer.” *Thorner v. Sony Computer Entm’t Am. LLC*, 669 F.3d 1362, 1366-67 (Fed. Cir. 2012); see *Epistar Corp. v. Int’l Trade Comm’n*, 566 F.3d 1321, 1334 (Fed. Cir. 2009) (requiring “expressions of manifest exclusion or restriction, representing a clear disavowal of claim scope” to deviate from the ordinary meaning) (citation omitted).

C. Infringement

“An infringement analysis entails two steps. The first step is determining the meaning and scope of the patent claims asserted to be infringed. The second step is comparing the properly construed claims to the device accused of infringing.” *Markman*, 52 F.3d at 976. A patentee may prove infringement either literally or under the doctrine of equivalents. Infringement of either sort must be proven by a preponderance of the evidence. *SmithKline Diagnostics, Inc. v. Helena Labs. Corp.*, 859 F.2d 878, 889 (Fed. Cir. 1988). A preponderance of the evidence standard “requires proving that infringement was more likely than not to have occurred.” *Warner-Lambert Co. v. Teva Pharm. USA, Inc.*, 418 F.3d 1326, 1341 n.15 (Fed. Cir. 2005).

[REDACTED]

Literal infringement is a question of fact. *Finisar Corp. v. DirecTV Group, Inc.*, 523 F.3d 1323, 1332 (Fed. Cir. 2008). “To establish literal infringement, every limitation set forth in a claim must be found in an accused product, exactly.” *Microsoft Corp. v. GeoTag, Inc.*, 817 F.3d 1305, 1313 (Fed. Cir. 2016) (quoting *Southwall Techs., Inc. v. Cardinal IG Co.*, 54 F.3d 1570, 1575 (Fed. Cir. 1995)). If any claim limitation is absent, there is no literal infringement of that claim as a matter of law. *Bayer AG v. Elan Pharm. Research Corp.*, 212 F.3d 1241, 1247 (Fed. Cir. 2000).

Doctrine of equivalents is also a form of infringement. One rubric for evaluating if a claimed feature is not literally, but nonetheless equivalent to, a claimed feature is known as the function-way-result test. Under this test, the accused feature is equivalent to the claim limitation when “it performs substantially the same function in substantially the same way to obtain the same result.” *Duncan Parking Techs., Inc. v. IPS Grp., Inc.*, 914 F.3d 1347, 1362 (Fed. Cir. 2019) (quoting *Graver Tank & Mfg. Co. v. Linde Air Prods. Co.*, 339 U.S. 605, 608 (1950)). Another is known as the insubstantial differences test, where “[a]n element in the accused device is equivalent to a claim limitation if the only differences between the two are insubstantial.” *Voda v. Gordia Corp.*, 536 F.3d 1311, 1326 (Fed. Cir. 2008) (citing *Honeywell Int’l Inc. v. Hamilton Sundstrand Corp.*, 370 F.3d 1131, 1139 (Fed. Cir. 2004)). The Supreme Court has further instructed, “the proper time for evaluating equivalency . . . is at the time of infringement, not at the time the patent was issued.” *Warner-Jenkinson Co., Inc. v. Hilton Davis Chem. Co.*, 520 U.S. 17, 37 (1997).

In addition to direct infringement, Section 271 of the Patent Act defines two categories of indirect infringement, inducement of infringement and contributory infringement. 35 U.S.C. § 271. For indirect infringement violations under Section 337, the direct infringement element may occur after importation, so long as all the other elements of indirect infringement are met at the time of importation. *See Certain Vision-Based Driver Assistance System Cameras and*

[REDACTED]

Components Thereof, Inv. No. 337-TA-907, Comm’n Op. at 19 (Dec. 1, 2015) (citing *Suprema, Inc. v. Int’l Trade Comm’n*, 796 F.3d 1338, 1348 (Fed. Cir. 2015)). It is well settled that “[a]bsent direct infringement of the patent claims, there can be neither contributory infringement ... nor inducement of infringement.” *Met-Coil Sys. Corp. v. Korners Unltd., Inc.*, 803 F.2d 684, 687 (Fed.Cir.1986) (citations omitted).

As to the first category, “[w]hoever actively induces infringement of a patent shall be liable as an infringer.” 35 U.S.C. § 271(b); see *DSU Med. Corp. v. JMS Co.*, 471 F.3d 1293, 1305 (Fed. Cir. 2006) (en banc) (“To establish liability under section 271(b), a patent holder must prove that once the defendants knew of the patent, they actively and knowingly aided and abetted another’s direct infringement.”) (citations omitted). “The mere knowledge of possible infringement by others does not amount to inducement; specific intent and action to induce infringement must be proven.” *Id.* (citations omitted). A defendant’s belief regarding patent validity is not a defense to a claim of induced infringement. *Commil USA, LLC v. Cisco Sys., Inc.*, 135 S. Ct. 1920 (2015). A defendant’s willful blindness on the question of infringement will satisfy the knowledge requirement. *Global-Tech Appliances, Inc. v. SEB S.A.*, 563 U.S. 754, 765, 768-771 (2011). The requisite intent to induce infringement “may be established through circumstantial evidence” and “may be inferred from all of the circumstances.” See *Broadcom Corp. v. Qualcomm, Inc.*, 543 F.3d 683, 699 (Fed. Cir. 2008) (citing *DSU Med. Corp. v. JMS Co.*, 471 F.3d 1293 (Fed. Cir. 2006); *Water Techs. Corp. v. Claco, Ltd.*, 850 F.2d 660, 669 (Fed. Cir. 1988)).

As to the second category, “a party who sells a component with knowledge that the component is especially designed for use in a patented invention, and is not a staple article of commerce suitable for substantial noninfringing use, is liable as a contributory infringer.” *Wordtech Sys., Inc. v. Integrated Networks Solutions, Inc.*, 609 F.3d 1308, 1316 (Fed. Cir. 2010).

[REDACTED]

Contributory infringement is premised upon a finding that: (1) the entity sells, offers to sell, or imports into the United States a component of a product; (2) the component has no substantial non-infringing use; (3) the component constitutes a material part of the claimed invention; (4) the entity was aware of the patent and knew that the product may be covered by a claim of the patent; and (5) the use of the component in the product directly infringes the claim. *See Certain Gaming & Entm't Consoles, Related Software, & Components Thereof*, Inv. No. 337-TA-752, Initial Determination (Remand) at 7-8 (Mar. 22, 2013) (citing 35 U.S.C. § 271(c); *Arris Group v. British Telecomm. PLC*, 639 F.3d 1368, 1376 (Fed. Cir. 2011)), *non-reviewed*, Notice (May 23, 2013). As with inducement, willful blindness on the question of infringement will satisfy the knowledge requirement. *Global-Tech*, 563 U.S. at 765, 768-771.

D. Domestic Industry

In an investigation based on a claim of patent infringement, section 337 requires that an industry in the United States, relating to the articles protected by the patent, exist or be in the process of being established. 19 U.S.C. § 1337(a)(2). Under Commission precedent, the domestic industry requirement has been divided into (i) a “technical prong” (which requires articles covered by the asserted patent) and (ii) an “economic prong” (which requires certain levels of activity with respect to the protected articles or patent itself). *See Certain Video Game Systems and Controllers*, Inv. No. 337-TA-743, Comm’n Op. at 6-7 (April 14, 2011) (“*Video Game Systems*”).

1. Technical Prong

The technical prong of the domestic industry requirement is satisfied when the complainant in a patent-based section 337 investigation establishes that it is practicing or exploiting the patents at issue. *See* 19 U.S.C. §§ 1337 (a)(2), (3); *Certain Microsphere Adhesives, Process for Making Same and Prods. Containing Same, Including Self-Stick Repositionable Notes*, Inv. No. 337-TA-366, Comm’n Op. at 8 (U.S.I.T.C. Jan. 16, 1996). “In order to satisfy the technical prong of the

[REDACTED]

domestic industry requirement, it is sufficient to show that the domestic industry practices any claim of that patent, not necessarily an asserted claim of that patent.” *Certain Ammonium Octamolybdate Isomers*, Inv. No. 337-TA-477, Comm’n Op. at 55 (U.S.I.T.C. Aug. 28, 2003). Historically, the Commission permits the complainant’s products, and those of its licensees, to be considered for technical prong purposes. *See Certain Magnetic Tape Cartridges and Components Thereof*, Inv. No. 337-TA-1058, Comm’n Op. at 28-29 (April 9, 2019).

The test for claim coverage for the purposes of the technical prong of the domestic industry requirement is the same as that for infringement. *See Certain Doxorubicin and Preparations Containing Same*, Inv. No. 337-TA-300, Initial Determination at 109 (U.S.I.T.C. May 21, 1990), *aff’d*, Views of the Commission at 22 (U.S.I.T.C. Oct. 31, 1990); *Alloc, Inc. v. Int’l Trade Comm’n*, 342 F.3d 1361, 1375 (Fed. Cir. 2003). “First, the claims of the patent are construed. Second, the complainant’s article or process is examined to determine whether it falls within the scope of the claims.” *Certain Doxorubicin and Preparations Containing Same*, Inv. No. 337-TA-300, Initial Determination at 109. As with infringement, the technical prong of the domestic industry can be satisfied either literally or under the doctrine of equivalents. *Certain Dynamic Sequential Gradient Devices and Component Parts Thereof*, Inv. No. 337-TA-335, ID at 44, Pub. No. 2575 (U.S.I.T.C. May 15, 1992). In short, the patentee must establish by a preponderance of the evidence that the domestic product practices one or more claims of the patent.

2. Economic Prong

The “economic prong” of the domestic industry requirement is satisfied when there exists in the United States, in connection with products practicing at least one claim of the patent at issue: (A) significant investment in plant and equipment; (B) significant employment of labor or capital; or (C) substantial investment in its exploitation, including engineering, research and development, and licensing. 19 U.S.C. § 1337(a)(3). Establishment of the “economic prong” is not dependent

[REDACTED]

on any “minimum monetary expenditure” and there is no need for complainant “to define the industry itself in absolute mathematical terms.” *Certain Stringed Musical Instruments and Components Thereof*, Inv. No. 337-TA-586, Comm’n Op. at 25-26 (May 16, 2008) (“*Stringed Instruments*”). However, a complainant must substantiate the significance of its activities with respect to the articles protected by the patent. *Certain Printing and Imaging Devices and Components Thereof*, Inv. No. 337-TA-690, Comm’n Op. at 30 (Feb. 17, 2011) (“*Imaging Devices*”). Further, a complainant can show that its activities are significant by showing how those activities are important to the articles protected by the patent in the context of the company’s operations, the marketplace, or the industry in question. *Id.* at 27-28. That significance, however, must be shown in a quantitative context. *Lelo Inc. v. Int’l Trade Comm’n*, 786 F.3d 879, 886 (Fed. Cir. 2015). The Federal Circuit noted that when the ITC first addressed this requirement, it found the word “‘significant’ denoted ‘an assessment of the *relative* importance of the domestic activities.’” *Id.* at 883-4 (internal citation omitted) (emphasis added). In general, “[t]he purpose of the domestic industry requirement is to prevent the ITC from becoming a forum for resolving disputes brought by foreign complainants whose only connection with the United States is ownership of a U.S. patent.” *Certain Battery-Powered Ride-On Toy Vehicles*, Inv. No. 337-TA-314, USITC Pub. No. 2420, Initial Determination at 21 (Aug. 1991); *see Certain Vacuum Insulated Flasks and Components Thereof*, Inv. No. 337-TA-1216, Notice at 3-4 (Oct. 21, 2021) (“Given the nature and extent of [complainant’s] investments in plant and equipment as a whole, [complainant] is not a mere importer.”).

E. Invalidity

1. 35 U.S.C. § 101

Section 101 states:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

35 U.S.C. § 101. To determine patent eligibility under § 101, courts apply the two-step *Alice* test and first “determine whether the claims at issue are directed to a patent-ineligible concept,” and then if so, “examine the elements of the claim to determine whether it contains an ‘inventive concept’ sufficient to ‘transform’ the claimed abstract idea into a patent-eligible application.” *Alice Corp. Pty. v. CLS Bank Intern.*, 573 U.S. 208, 217-18, 221 (2014). “The ‘directed to’ inquiry applies a stage-one filter to claims, considered in light of the specification, based on whether ‘their character as a whole is directed to excluded subject matter.’” *Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327, 1335 (Fed. Cir. 2016) (citing *Internet Patents Corp. v. Active Network, Inc.*, 790 F.3d 1343, 1346 (Fed. Cir. 2015); *Genetic Techs. Ltd. v. Merial L.L.C.*, 818 F.3d 1369, 1375 (Fed. Cir. 2016)). To save a patent at the second step, an inventive concept must be evident in the claims. *Synopsys, Inc. v. Mentor Graphics Corp.*, 839 F.3d 1138, 1151-52 (Fed. Cir. 2016).

2. 35 U.S.C. § 102

Pursuant to 35 U.S.C. § 102, a patent claim is invalid as anticipated if:

- (a) the invention was known or used by others in this country, or patented or described in a printed publication in this or a foreign country, before the invention thereof by the applicant;
- (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of the application for patent in the United States;

35 U.S.C. § 102 (pre-AIA). “A patent is invalid for anticipation if a single prior art reference discloses each and every limitation of the claimed invention. Moreover, a prior art reference may

[REDACTED]

anticipate without disclosing a feature of the claimed invention if that missing characteristic is necessarily present, or inherent, in the single anticipating reference.” *Schering Corp. v. Geneva Pharm., Inc.*, 339 F.3d 1373, 1377 (Fed. Cir. 2003) (citations omitted); *see Santarus, Inc. v. Par Pharm., Inc.*, 694 F.3d 1344, 1354 (Fed. Cir. 2012). “A century-old axiom of patent law holds that a product ‘which would literally infringe if later in time anticipates if earlier.’” *Upsher-Smith Labs., Inc. v. PamLab, L.L.C.*, 412 F.3d 1319, 1322 (Fed. Cir. 2005) (citing *Schering Corp.*, 339 F.3d at 1322). Anticipation, and all other grounds of patent invalidity, must be proved by clear and convincing evidence. *Microsoft Corp. v. I4I Ltd. P’ship*, 564 U.S. 91, 95, (2011).

3. 35 U.S.C. § 103

Section 103 of the Patent Act states:

A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

35 U.S.C. § 103(a) (pre-AIA). “Obviousness is a question of law based on underlying questions of fact.” *Scanner Techs. Corp. v. ICOS Vision Sys. Corp. N.V.*, 528 F.3d 1365, 1379 (Fed. Cir. 2008). The underlying factual determinations include: “(1) the scope and content of the prior art, (2) the level of ordinary skill in the art, (3) the differences between the claimed invention and the prior art, and (4) objective indicia of non-obviousness.” *Id.* (citing *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1, 17-18 (1966)). These factual determinations are often referred to as the “*Graham* factors.”

The critical inquiry in determining the differences between the claimed invention and the prior art is whether there is a reason to combine the prior art references. *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 418-21 (2007). In *KSR*, the Supreme Court rejected the Federal Circuit’s rigid

application of the teaching-suggestion-motivation test. While the Court stated that “it can be important to identify a reason that would have prompted a person of ordinary skill in the relevant field to combine the elements in the way the claimed new invention does,” it described a more flexible analysis:

Often, it will be necessary for a court to look to interrelated teachings of multiple patents; the effects of demands known to the design community or present in the marketplace; and the background knowledge possessed by a person having ordinary skill in the art, all in order to determine whether there was an apparent reason to combine the known elements in the fashion claimed by the patent at issue As our precedents make clear, however, the analysis need not seek out precise teachings directed to the specific subject matter of the challenged claim, for a court can take account of the inferences and creative steps that a person of ordinary skill in the art would employ.

Id. at 418. Since *KSR*, the Federal Circuit has announced that, where a patent challenger contends that a patent is invalid for obviousness based on a combination of prior art references, “the burden falls on the patent challenger to show by clear and convincing evidence that a person of ordinary skill in the art would have had reason to attempt to make the composition or device . . . and would have had a reasonable expectation of success in doing so.” *PharmaStem Therapeutics, Inc. v. ViaCell, Inc.*, 491 F.3d 1342, 1360 (Fed. Cir. 2007); *see KSR*, 550 U.S. at 399 (“The proper question was whether a pedal designer of ordinary skill in the art, facing the wide range of needs created by developments in the field, would have seen an obvious benefit to upgrading Asano with a sensor.”). In addition to demonstrating that a reason exists to combine prior art references, the challenger must demonstrate that the combination of prior art references discloses all of the limitations of the claims. *Velandier v. Garner*, 348 F.3d 1359, 1363 (Fed. Cir. 2003) (explaining that a requirement for a finding of obviousness is that “all the elements of an invention are found in a combination of prior art references”).

An obviousness determination should also include a consideration of how “secondary considerations” such as “commercial success, long felt but unsolved needs, failure of others, etc.,

[REDACTED]

might be utilized to give light to the circumstances surrounding the origin of the subject matter sought to be patented.” *Graham*, 338 U.S. at 17-18. “For [such] objective evidence to be accorded substantial weight, its proponent must establish a nexus between the evidence and the merits of the claimed invention.” *In re GPAC Inc.*, 57 F.3d 1573, 1580 (Fed. Cir. 1995); *see Merck & Cie v. Gnosis S.P.A.*, 808 F.3d 829, 837 (Fed. Cir. 2015). “Where the offered secondary consideration actually results from something other than what is both claimed and novel in the claim, there is no nexus to the merits of the claimed invention.” *In re Huai-Hung Kao*, 639 F.3d 1057, 1068 (Fed. Cir. 2011); *see Apple Inc. v. Samsung Elecs. Co., Ltd.*, 839 F.3d 1034, 1054-1056 (Fed. Cir. 2016).

III. IMPORTATION AND JURISDICTION

In its initial post-hearing brief, AMD contends, “[t]he evidence shows that Respondents have each imported, sold for importation, or sold after importation, at least one Accused Product, directly and/or indirectly through their corporate affiliates and distribution networks, thereby satisfying the importation requirement pursuant to 19 U.S.C. § 1337(a)(1)(B).” CIB at 16. AMD notes that “[w]ith regards to TCL, importation is undisputed.” *Id.* (citing, *inter alia*, JX-0011C; JX-0012C); JX-0074C (Gong) at 42:9-13, 55:10-14, 109:21-25, 39:18-21, 43:22-23, 114:12-115:3). [REDACTED]

[REDACTED] Respondents offer no dispute here. *See generally* RIB at 10-11; RRB at 2.

Accordingly, the importation requirement under 19 U.S.C. § 1337(a)(1)(B) is satisfied, and the Commission has *in rem* jurisdiction over the TCL Accused Products.

As to the approximately [REDACTED] Realtek Accused Products, AMD contends “[i]mportation is also undisputed with regards to Realtek. Realtek imports, sells for importation, and/or sells in the United States after importation Accused Products” according to a letter and stipulation filed as part

Contains Confidential Business Information Subject to Protective Order

of the present investigation. CIB at 17 (citing EDIS Doc. ID 789311; JX-0348C at ¶ 5). Respondents, however, do oppose importation for three models of Accused Products. RIB at 11. Respondents argue, “[t]here is no evidence in the record showing that Realtek has imported, sold for importation, or sold after importation at least [REDACTED] [REDACTED]” and “AMD did not introduce any evidence that these products, or [REDACTED] [REDACTED], have actually been imported into the United States, either by Respondents or by any third party. Nor did AMD introduce any evidence that Realtek sold these products ‘for importation.’” *Id.*; see RRB at 2 (“These were not the subject of the parties’ importation stipulation . . . and were not addressed by Complainants.”).

For these three models, AMD counters, “Realtek is incorrect and ignores the vast evidence of Realtek’s sale for importation of these products—up to [REDACTED] products annually to [REDACTED], with Realtek’s heavy involvement in the design and manufacture, and their knowledge or constructive knowledge of subsequent exportation to the United States.” CRB at 2 (citing CIB at 147-150). AMD adds:

Realtek, with a specific support team of at least fourteen individuals, tailors its products and presentations to [REDACTED]

[REDACTED] to purchase [REDACTED]

See, e.g., CX-1640C.0009-0026; CX-1535C at CX.1535C.1; CX-1656C at CX-1656C.1; CX-1655C; JX-0022C at JX-0022C.70 (Rog. No. 9); JX-0048C (Chen Tr.) at 58:6-59:3; CX-1640C at CX-1640C.24 (presentation to [REDACTED]). Realtek acknowledges that [REDACTED] consumers will be using these devices to watch United States baseball. CX-1543C at CX-1543C.00017.


Id.

AMD has shown it is more likely than not that Realtek has sold [REDACTED]

[REDACTED] for importation into United States. Presentations from Realtek to [REDACTED]

Contains Confidential Business Information Subject to Protective Order

demonstrate a collaboration and communication such that Realtek knows which of its GPU products end up in which [REDACTED] products:



CX-1543C.3; see CX1640C.15, 16; CX-1543C.42 (coordinating [REDACTED]
[REDACTED]). Other slides—also authored by Realtek—strongly suggest the United States is an understood market for these devices:

[REDACTED]

CX-1543C.17. When combined with the fact that massive commercial quantities of [REDACTED] [REDACTED] have been sold at least to one distributor, [REDACTED] [REDACTED], it is more likely than not that at least some of these [REDACTED] units were imported and sold within the United States.

As for [REDACTED], only [REDACTED] have been sold by [REDACTED]. CX-1359C. This is not a commercial quantity, so it cannot be assumed that at least some [REDACTED] have reached their intended commercial destinations (*e.g.*, the United States) as opposed to being held by intermediate vendors for evaluation. *See, e.g.*, CIB at 148 (describing “pattern of smaller orders preceding orders for large volumes” where smaller orders are just samples). AMD has presented no other evidence that any of these [REDACTED] were known to be headed for the United States, even with non-party [REDACTED] being a U.S.-based company with headquarters in [REDACTED]. *See* CRB at 2. Nonetheless, because every other Accused Product has been imported, and infringement of every Accused Product, including [REDACTED] [REDACTED], rises or falls based on analysis of the representative code, the importation requirement has

[REDACTED]

been satisfied even as to [REDACTED]. *See Electrical Connectors and Cages, Components Thereof, and Products Containing the Same*, Inv. No. 337-TA-1241, Comm’n Op. at 10 (Sep. 30, 2022) (public) (“the importation of a representative product within each of these product categories is sufficient to satisfy the importation requirement for all products in each of these representative categories”).

Accordingly, the importation requirement under 19 U.S.C. § 1337(a)(1)(B) is satisfied, and the Commission has *in rem* jurisdiction over the Accused Products.

IV. U.S. PATENT NO. 7,742,053

A. Level of Ordinary Skill in the Art

In Order No. 31, I determined that a person having ordinary skill in the art of the 053 patent at the time of invention:

[W]ould have a bachelor degree in electrical engineering, computer engineering, computer science, or a related field, and at least two years of experience in ‘parallel computing or graphics processing including developing, designing, or programming software or hardware for graphics processing or parallel processing units, graphics accelerators[,] or other graphics processing systems,’ with additional education substituting for experience and vice versa.

Order No. 31 at 9. The parties do not challenge this definition (*see generally* CIB; CRB; RIB; RRB), and it is applied throughout this initial determination.

B. Claims-at-Issue

Claims 5, 6, and 9 of the 053 patent are at issue in this investigation, either through allegations of infringement or domestic industry technical prong. *See* CIB at 3. They are reproduced below, along with the parties’ agreed limitation annotations:

5. A graphics processing system comprising:

[5a] at least one memory device comprising a first portion operative to store a plurality of pixel command threads and a second portion operative to store a plurality of vertex command threads;

Contains Confidential Business Information Subject to Protective Order

[5b] an arbiter, coupled to the at least one memory device, operable to select a command thread from either of the plurality of pixel command threads and the plurality of vertex command threads; and

[5c] a plurality of command processing engines, coupled to the arbiter, each operable to receive and process the command thread.

6. The graphics processing system of claim 5, wherein the plurality of command processing engines comprises at least one arithmetic logic unit.

....

9. The graphics processing system of claim 5, wherein a single memory device comprises the first portion and the second portion of the at least one memory device.

053 patent at cls. 5 (annotated), 6, 9.

C. Claim Construction

As part of the *Markman* process, the following claim terms of the 053 patent were construed, either as-agreed between the parties or determined by Order No. 31:

Claim Term	Construction
“command threads”	“a sequence of commands applicable to the corresponding element,” with no requirement that they comprise state bits
“memory device”	Does not mean “memory with dedicated storage areas for pixel and vertex threads”
“coupled to”	Does not mean “directly connected”

See Order No. 31 at 20, 21, 22; *see also id.* at 15-17, 19, 21 (resolving additional disputes but not advancing a specific construction). Within their briefing, the parties identify a number of additional claim construction disputes, discussed below.

1. “at least one memory device . . . operative to store a plurality of pixel command threads and . . . a plurality of vertex command threads”

For “at least one memory device . . . operative to store a plurality of pixel command threads and . . . a plurality of vertex command threads,” the parties dispute whether a complete thread must

[REDACTED]

be stored at a given moment in time, or whether portions of a complete thread are sufficient so long as the entire thread is stored over a period of time. AMD argues for the latter, and Respondents the former. *See* CIB at 21; RIB at 18.

The balance of evidence supports Respondents' interpretation. First, the 053 patent is clear, as confirmed in Order No. 31, that a command thread is made up of commands. The patent states, "a command thread is a sequence of commands applicable to the corresponding element, such as [a] pixel command thread relative to processing of pixel elements and a vertex command thread relative to vertex processing commands." 053 patent at 2:41-45; *see id.* at 2:52-56, 3:20-23, 3:65-4:1, 4:23-27; Order No. 31 at 19-20. Second, the patent's claims exclusively refer to "command threads" as the units of information which are stored, selected, or provided amongst components of the graphics processing system (*i.e.*, the currency of the system). *See* 053 patent at cls. 1-9. The patent's specification also consistently refers to entire threads as what are passed between and operated on by the system's components, both in the present invention and in the prior art:

A graphics processing system comprises at least one memory device storing a plurality of pixel command threads and a plurality of vertex command threads. An arbiter coupled to the at least one memory device is provided that selects a command thread from either the plurality of pixel or vertex command threads based on relative priorities of the plurality of pixel command threads and the plurality of vertex command threads. The selected command thread is provided to a command processing engine capable of processing pixel command threads and vertex command threads.

id. at Abstract;

In the prior art embodiments of FIG. 1, the first buffer 104 receives an input command 124 and outputs a completed command thread 126 to the second arbiter 102. In one embodiment, the command thread may include an indicator, such as a flag, indicating when the access to the ALU resources has been completed for the associated command. The arbiter 102 receives the input command 124 and thereupon provides, in due course, the command thread to either an appropriate texture fetch buffer 110 or an ALU buffer 108. Thereupon, the steps are repeated

[REDACTED]

where an output thread command 128 is provided to another ALU (not shown) or texture fetch processor (not shown) and returned to the buffer 108 or 110. The buffer 110 also produces the output 132 which is a command thread. The output 132 may be provided to another arbiter 103 to be provided further along the graphics processing pipeline.

id. at 1:49-64;

The system and method further includes an arbiter operably coupled to the reservation station such that the arbiter retrieves a first command thread of the plurality of command threads stored therein. The arbiter may be any implementation of hardware, software or combination thereof such that the arbiter receives the command thread and thereupon provides the command thread to a command processing engine. The system and method further includes the command processing engine coupled to receive the first command thread from the arbiter such that the command processor may perform at least one processing command from the command thread. Whereupon, a command processing engine provides the first command thread back to the associated reservation station.

id. at 2:45-58;

With respect to FIG. 4, a pixel command thread 324 may be retrieved by the arbiter 306 and a vertex command thread 326 may also be retrieved. The arbiter 306 then provides one thread 328, which may be either 324 or 326 to the graphics processing engine 310, such as a texture engine, and provides the other thread 330 to the ALU 308.

Upon execution of the command, the ALU 308 then returns the command thread 332 to the appropriate reservation station 302 or 304. As illustrated in FIG. 4, the ALU 308 is coupled to both reservation station 302 and reservation station 304 for providing the thread back thereto. The same data transfer occurs when the graphic processing engine 310 performs the commands and returns the command thread 334 back to the originating reservation station 302 or 304. It is also noted, that in the present embodiment, multiple command operations may be performed by a particular unit 308 or engine 310, but in order to switch a command thread from ALU 308 to a graphics processing engine 310, that command thread must be returned back to the appropriate reservation station 302 or 304 and re-retrieved by the arbiter 306 and thereupon provided to the other unit 308 or engine 310 respectively.

id. at 4:28-48; *see id.* at 4:65-5:5, 5:43-56.

In the pertinent context of storage by a “memory device,” the specification is no different.

The elements most akin to the claims’ “memory device” are the “reservation stations,” as these are

disclosed as storing a plurality of command threads and being coupled to the arbiter:




FIG. 4

053 patent at Fig. 4; *id.* at 2:36-41 (“Generally, the present invention includes a multi-thread graphics processing system and method thereof including a reservation station having a plurality of command threads stored therein. A reservation station may be any type of memory device capable of reserving and storing command threads.”), cl. 5 (“at least one memory device . . . an arbiter, coupled to the at least one memory device”). These reservation stations, consistent with the rest of the specification, are disclosed as storing complete command threads without reference to storage over time or in portions of threads:

In the ALU resource division 120, the command thread 118 may be received from an input command 124 as selected by the arbiter 101. The command thread 118 may then be withdrawn from the reservation stations 104 and 108 for the purpose of being provided to an ALU (not shown) and the command threads within the texture fetch resource division 122 maybe withdrawn from the reservation stations 106 and 110 to be provided to a texture fetch processors (not shown).

053 patent at 1:40-48.



Generally, the present invention includes a multi-thread graphics processing system and method thereof including a reservation station having a plurality of command threads stored therein. A reservation station may be any type of memory device capable of reserving and storing command threads. . . . The system and method further includes the command processing engine coupled to receive the first command thread from the arbiter such that the command processor may perform at least one processing command from the command thread. Whereupon, a command processing engine provides the first command thread back to the associated reservation station.

id. at 2:36-58;

Therefore, when all of the commands in the command thread have been executed and the command thread is retrievable from the reservation station 202, the command thread may be provided to a further processing element (not illustrated) within a graphics processing pipeline.

id. at 3:23-28;

FIG. 4 illustrates another embodiment of a multi-thread command processing system 300 having a first reservation station 302, a second reservation station 304, an arbiter 306, an ALU 308 and a graphics processing engine 310. In this embodiment, the first reservation station 302 is a pixel reservation station such that the command threads 312, 314 and 316 contain pixel-based commands therein. Furthermore, in this embodiment the second reservation station 304 is a vertex reservation station is directed towards vertex command threads illustrated as command threads 318, 320 and 322.

id. at 3:63-4:4;

It is also noted, that in the present embodiment, multiple command operations may be performed by a particular unit 308 or engine 310, but in order to switch a command thread from ALU 308 to a graphics processing engine 310, that command thread must be returned back to the appropriate reservation station 302 or 304 and re-retrieved by the arbiter 306 and thereupon provided to the other unit 308 or engine 310 respectively.

id. at 4:41-48;

The method further includes providing the second command thread to the graphics processing engine, step 430. The next step, step 432, is prior to writing the selected command thread to either the first reservation station or the second reservation station, interleaving the selected command thread and the second selected command thread. Thereupon, the method further includes performing a second command in response to the second selected command thread, step 434.

[REDACTED]

id. at 6:23-30. And when the patent uses expansive language meant to capture variations on the embodiments disclosed, it refers to “storage capacity . . . to accommodate the storage [of] any suitable corresponding number of command threads” which again implies complete threads.

Otherwise storage capacity would not seem to be an issue:

It should be understood that there exists implementations of other variations and modifications of the invention and its various aspects, as may be readily apparent to those of ordinary skill in the art, and that the invention is not limited by the specific embodiments described herein. For example, the storage capacity of the reservation stations may be adequately adjusted to accommodate the storage [of] any suitable corresponding number of command threads.

Id. at 6:64-7:4. Thus, the intrinsic evidence strongly supports finding “at least one memory device comprising a first portion operative to store a plurality of pixel command threads and a second portion operative to store a plurality of vertex command threads” means storing complete command threads at any given moment in time, and not partial command threads.

AMD highlights two specification passages in opposition. CIB at 21-22 (citing 053 patent at 3:4-8, 4:49-51). Neither is particularly persuasive. One reads, “[i]n one embodiment, each command thread within the reservation station 302 and 304 may be stored across two physical pieces of memory.” 053 patent at 4:49-51. AMD argues that Respondents improperly read out this embodiment with their construction. CIB at 21-22. Not so, however. A reservation station that consists of two physical pieces of memory is no less a “memory device,” in the parlance of the claims, than if it were just one piece. In either circumstance, “each command thread” is “stored.” 053 patent at 4:49-51.

The second passage reads, “[a]s noted above, in one embodiment the reservation station 202 operates similar to a first in first out (FIFO) memory device, although command threads 208-212 may be retrieved from any location, whereas completed commands exit the memory device in a FIFO manner.” 053 patent at 3:4-8. AMD argues that individual commands exiting the memory

[REDACTED]

in a FIFO manner “indicat[es] that individual commands of a command thread may also enter in an FIFO manner, and thus allow[s] for only a portion of the commands of a command thread to be stored by the memory device at one time.” CIB at 21.

AMD’s inference is not totally unreasonable, but also not particular supportive of the argument, because loading a FIFO buffer one command at a time can still result in a moment where the entire thread is stored. Moreover, this single paragraph is outweighed by the rest of the specification which overwhelmingly discusses entire threads leaving the reservation station and being returned. For example, the patent states:

Once a thread is selected by the arbiter 306, the thread is marked as invalid and submitted to the appropriate execution unit 308 or 312. Upon the execution of the associated command of the command thread, the thread is thereupon returned to the station 302 or 304 at the same storage location with its status updated, once all possible sequential instructions have been executed.

....

Once again, using the exemplary embodiment of FIG. 4, the command threads, such as 312-316 and/or 318-322 may be provided from the graphics processing engine 310 and written back thereto or in another embodiment may be provided to the ALU 308 by the arbiter 306 and, upon execution of an arithmetic command, provided back to the associated reservation station, 302 or 304 respectively.

053 patent at 4:21-27, 6:47-54. It is also entirely possible that AMD’s excerpt reflects a typographical error when it refers to “commands” and not “command threads.” Later paragraphs repeat the same first-in-first-out procedure—but for *entire threads*:

In one embodiment, each station 302, 304 maintains the state of each thread, such as threads 312-322. In one embodiment, the thread lives in a given location in the station 302, 304, *in the order that the thread is received* therein.

....

In one embodiment, the arbiter 204 retrieves the command threads 208-212 based on a priority scheme. For example, the priority may be based on specific commands that have been executed within a command thread or specific commands which are to be executed within a command for the effective utilization of the arbiter 204 and

[REDACTED]

the command processing engine 206. In an alternative embodiment, the arbiter 204 *may always retrieve the oldest available thread*.

Id. at 4:14-17 (emphasis added), 3:29-36 (emphasis added).

Beyond the intrinsic evidence, AMD highlights testimony of Respondents' expert, Dr. Massoud Pedram, that to satisfy the limitation, “all of the instructions within the required plurality of vertex and pixel command threads need not be stored in the memory device at the same time.” CIB at 9 (citing Hr'g Tr. (Pedram) at 982:11-17); CRB at 8-9 (same). This is certainly helpful to AMD's position. But this question and answer occurred during cross examination, when Dr. Pedram was not given an opportunity to explain his opinion. *See* Hr'g Tr. (Pedram) at 982:11-17. And it is otherwise unclear whether his opinion was that a person of ordinary skill in the art at the time of the invention would understand the patent's reference to storage of a “command thread” to mean any storage of any one of the commands in that thread, which would be clearly relevant to claim construction, rather than something more specific to the prior art he was being cross examined about at that moment in the hearing. *See id.*

Overall then, the balance of evidence shows “at least one memory device . . . operative to store a plurality of pixel command threads and . . . a plurality of vertex command threads” means operative to store a plurality of complete sequences of pixel commands and a plurality of complete sequences of vertex commands.

2. “select a command thread”

The parties also identify “select a command thread” within claim 5 as having a disputed meaning. AMD contends “the claimed arbiter can satisfy the requirement of selecting ‘a’ command thread from the plurality of pixel and vertex command threads so long as the claimed arbiter makes a selection of *one or more* command threads of the plurality of command threads.” CIB at 22 (emphasis in original); CRB at 3-7. Respondents on the other hand contend, “this claim

[REDACTED]

term should be construed to require that the arbiter be ‘operable to select a single command thread,’ which does not preclude the arbiter from *also* [being] capable of selecting a plurality of command threads from a larger plurality,” or, put another way, “the claim still requires the arbiter to have *the minimum capability of selecting a single command thread.*” RIB at 11 (emphasis in original), 15 (“the arbiter must be capable of selecting a single thread; it may be able to select more than one thread, but it must be able to select, at a minimum, just one”); RRB at 2-4.

Although it is a close issue, AMD has the more persuasive position. The Federal Circuit has held clearly “[t]hat ‘a’ or ‘an’ can mean ‘one or more’ is best described as a rule, rather than merely as a presumption or even a convention.” *Baldwin Graphic Sys., Inc. v. Siebert, Inc.*, 512 F.3d 1338, 1342 (Fed. Cir. 2008). Even Respondents acknowledge, with citation to *Baldwin*, that limiting “a” to just one or single item “is an exception to the rule and it ‘arises where the language of the claims themselves, the specification, or the prosecution history [*i.e.*, the intrinsic evidence] necessitate a departure from the rule.’” RIB at 4 (citing *Baldwin*, 512 F.3d at 1342-43). Yet, as AMD observes, Respondents point to no specification excerpts or prosecution statements to support their exception—*e.g.*, not even generalized descriptions of the purpose or benefits of single thread selection. *See* CIB at 22; *see generally* RIB at 11-17; RRB at 2-4. Indeed, Respondents essentially rely on nothing in the claims apart from the claim language at issue. *See generally* RIB at 11-17; RRB at 2-4. They do cite also to limitation 5[c], reciting a command processing engine “operable to receive and process the command thread.” RIB at 14. But as AMD argues, this ordinary antecedent use of “the” has no bearing on whether “a” is met when more than thread is selected by the arbiter. CRB a 5-6. Respondents’ reliance on statements made during *Markman* proceedings and various expert testimony does not overcome this obvious lack of intrinsic support. *KCJ Corp. v. Kinetic Concepts, Inc.*, 223 F.3d 1351, 1356 (Fed. Cir. 2000) (“the article ‘a’ receives

[REDACTED]

a singular interpretation only in rare circumstances when the patentee evinces a clear intent to so limit the article”) (citation omitted). Respondents’ restrictive scope is therefore not justified.

Accordingly, “select a command thread” does not mean selecting one and only one thread.

3. “command processing engine”

As part of *Markman* proceedings, the parties disputed “whether the command processing engine may be hardware or software implemented on hardware . . . or is limited to hardware.” Order No. 31 at 18. Order No. 31 deferred the issue after finding “[t]he intrinsic evidence does not conclusively resolve this dispute,” and that additional evidence, particularly expert testimony, would be helpful. *Id.* at 18-19.

AMD continues to argue for a hardware-only requirement to “command processing engine.” CIB at 22. For support, AMD looks primarily to testimony from Dr. Reinman and a final written decision from the PTAB resolving an *inter partes* review of the 053 patent. *Id.* at 22-23 (citing Hr’g Tr. (Reinman) at 251:1-10; *LG Electronics, Inc. v. ATI Technologies ULC*, IPR2015-00325, Paper 62 at 26 (P.T.A.B. Apr. 14, 2016) (“053 FWD”)).

AMD’s position is rejected. The cited testimony and written decision, while used to advocate for a hardware-only interpretation, fairly establish that “command processing engines” could also be software to those of ordinary skill in the art at the time of the invention. Dr. Reinman’s testimony that “this would be something that would be designed in hardware in order to be fast enough for a graphics processing system” implies that software is still an option, albeit not preferable. Hr’g Tr. (Reinman) at 251:5-10. And the PTAB decision reports that AMD had itself argued it had practiced the invention in software as an act of actual reduction to practice:

Here, based on the evidence in this trial, we are not persuaded that ATI has demonstrated that the claimed system was actually reduced to practice. In its Patent Owner Response and Sur-reply, ATI does not contend that it constructed a hardware device in accordance with the challenged claims. Nor has ATI established that an implementation based on the RTL code constitutes the subject matter of the

challenged claims in this record. At best, ATI has shown that the RTL code constitutes software files in a hardware description language that describes the logical design and behavior of the system. ATI, however, has not demonstrated that the design inferred by the RTL code was synthesized into a set of interconnected hardware devices. It is well-established that an equivalent of a physical embodiment does not satisfy the first requirement of actual reduction to practice. *Eaton v. Evans*, 204 F.3d 1094, 1097–98 (Fed. Cir. 2000) (holding that a party cannot obviate the first requirement of constructing a physical embodiment through evidence of testing an equivalent).

053 FWD at 19; *see id.* at 24 (dismissing “ATI’s argument that the inventive process for the claimed invention ended after chip design, ‘so it is logical that RTL is a valid reduction to practice.’”), 30 (identifying C++ emulation code and RTL code as having been used for reduction to practice), 33 (identifying desirable software-related evidence). Thus, as a matter of intrinsic evidence via the patentee’s own statements, the claims, including “command processing engines,” can be practiced via software.

With that said, the PTAB did conclude, as noted by AMD, that the claims could not be practiced solely in software and required hardware elements. It stated, “[c]alling the RTL software code an embodiment or implementation of the claimed system does not change the fact that the claimed hardware elements do not exist.” 053 FWD at 22; CIB at 23. The determination’s persuasive value is greatly diminished, however, by a lack of explanation as to why the claims absolutely required “physical” construction (*i.e.*, a hardware “command processing engine”). *See generally* 053 FWD. And the determination even acknowledged the petitioner’s argument that “command processing engine,” specifically, could be “hardware and/or software,” and accepted that party’s argument that “arbiter” could also be “any computer hardware, software, or combination thereof.” *See* 053 FWD at 6-8 (summarizing petitioner LG claim construction arguments and holdings).

[REDACTED]

Accordingly, the balance of evidence shows a person of ordinary skill would have understood a “command processing engine” could be implemented in software, hardware, or a mix of the two.

D. Infringement

According to AMD’s post-hearing briefing, the use, manufacture, or sale of Accused Products infringe claims 5, 6, and 9. *See* CIB at 12, 13. Of these, claim 5 is independent, and claims 6 and 9 each depend from claim 5. For the reasons discussed below, AMD has not shown infringement of any claim by the Accused Products.

1. Claim 5

For reference, claim 5 of the 053 patent requires:

5. A graphics processing system comprising:

at least one memory device comprising a first portion operative to store a plurality of pixel command threads and a second portion operative to store a plurality of vertex command threads;

an arbiter, coupled to the at least one memory device, operable to select a command thread from either of the plurality of pixel command threads and the plurality of vertex command threads; and

a plurality of command processing engines, coupled to the arbiter, each operable to receive and process the command thread.

053 patent at cl. 5. The parties dispute whether any of these limitations are found in the Accused Products.

a. [5a] “command thread”

The term “command thread” appears throughout claim 5 of the 053 patent. As discussed above, Order No. 31 set the construction of this term as “a sequence of commands applicable to the corresponding element.” Order No. 31 at 20. Within the accused products, AMD identifies a data unit referred to as [REDACTED] as the claimed “command thread.” CIB at 27.

Contains Confidential Business Information Subject to Protective Order

AMD argues that there are two types: [REDACTED]; and each constitutes the claimed “pixel command thread” and “vertex command thread,” respectively. *Id.* (citing Hr’g Tr. (Reinman) at 241:15-242:7). AMD provides the following background on fragments and vertices, and the process of [REDACTED]:

At trial, Dr. Reinman explained that a “fragment” is a type of data element, which will be a pixel on the screen unless it is hidden or obscured by other fragments. *Id.* at 106:25-107:5. Likewise, a “vertex” is another type of data element processed by a GPU, which makes up one point of a triangle of a 3D object within a 3D space. *Id.* at 105:3-106:4. As Respondents’ non-infringement expert, Dr. Pfister explained,

[REDACTED] Hrg. Tr. (Day 3) (Pfister) at 569:14-572:10; 795:4-19.

[REDACTED] Hrg. Tr. (Day 1) (Reinman) at 108:8-21; CX-0773C at CX-0773C.25.

[REDACTED] CX-0773C at CX-0773C.26-27.

[REDACTED] *Id.* at CX- 0773C.25.

[REDACTED] *Id.* at 31 (emphasis added); *see also* CX-1730C at CX-1730C.39

[REDACTED]).

[REDACTED] Hrg. Tr. (Day 1) (Reinman) at 112:4-23, 252:16-23, 254:25-255:21; CX-0773C at CX-0773C.25-27, 30-31; *see also* Hrg. Tr. (Day 2) (Pfister) at 571:14-575:1

[REDACTED] . Specifically,


[REDACTED] Hrg. Tr. (Day 1) (Reinman) at 112:20-23; *see also* Hrg. Tr. (Day 1) (Pfister) at 795:1-3 ([REDACTED]).

[REDACTED] Hrg. Tr. (Day 1) (Reinman) at 112:7-9.

[REDACTED]


Id. at 27-28. Dr. Reinman provides the following demonstrative illustration to show the composition of [REDACTED]

[REDACTED]:



CDX-0005C.180; Hr’g Tr. (Reinman) at 253:4-18. Respondents’ expert, Dr. Pfister, provides a similar demonstrative showing [REDACTED]

[REDACTED]



RDX-0005C.16 (citing, *inter alia*, RX-0491C.31). Dr. Pfister describes this as [REDACTED]

[REDACTED] Hr’g Tr. (Pfister) at 569:11-13. AMD argues [REDACTED]

[REDACTED] constitute a single

[REDACTED]

thread because [REDACTED]

[REDACTED] CRB at 11 (emphasis in original); CIB at 28. AMD finds this consistent with an embodiment of the 053 patent and Dr. Pfister’s description of [REDACTED] in a single vertex. CIB at 30-31 (citing 053 patent at 2:41-45; Hr’g Tr. (Pfister) at 795:4-19); *see id.* at 27 (citing, *inter alia*, Hr’g Tr. (Pfister) at 569:14-572:10).

AMD acknowledges however, [REDACTED] but labels this as [REDACTED] as opposed to how the term is used in the patent. CIB at 28-29 (citing, *inter alia*, Hr’g Tr. (Reinman) at 112:20-23). AMD also suggests that, based on the 053 patent specification, a “command thread” may pertain to multiple corresponding elements. *Id.* at 30 (citing 053 patent at 2:41-45); CRB at 14 (“the ’053 Patent provides no limit of the number of commands in the sequence of commands, much less on the number of elements that could be applicable to a specific command thread.”). In response to Respondents’ initial brief, AMD further argues that there is no difference between “instructions” and “commands,” a point allegedly agreed to by both of Respondents’ experts (CRB at 12 (citing Hr’g Tr (Pfister) at 773:4-774:13; Hr’g Tr. (Pedram) at 981:20-982:9)); and that its theory does not depend on treating just [REDACTED] as a “command thread” (*id.* (citing CIB at 31-38)).

Respondents nonetheless dispute that [REDACTED] can meet the construction of “command thread”—“a sequence of commands applicable to the corresponding element”—for a variety of reasons. *See* RIB at 20-25; RRB at 7-9. First, they argue, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] *Id.* (citing Hr’g Tr.

[REDACTED]

(Pfister) at 571:1-572:10, 597:12-598:5; Hr’g Tr. (Reinman) at 381:20-382:5) (emphasis in original). Respondents next contend that [REDACTED]

[REDACTED] in contravention of the ordered construction. RIB at 23 (citing Hr’g Tr. (Reinman) at 279:4-14). Relatedly, Respondents argue that any reliance on [REDACTED]

[REDACTED] See *id.* at 23-24 (citing Hr’g Tr. (Pfister) at 598:20-599:8, 601:15-24). Lastly, Respondents criticize the general reliability of Dr. Reinman based on his taking a position that, [REDACTED]
[REDACTED]
[REDACTED] *Id.* at 24 (citing Hr’g Tr. (Reinman) at 280:5-284:1).

AMD has shown by a preponderance of the evidence that [REDACTED] in the accused products, at least some of the time, qualify as “a sequence of commands applicable to the corresponding element” as construed by Order No. 31 and taught in the patent. As reproduced above, the experts agree that [REDACTED] in the Accused Products is, at least regularly, [REDACTED]. Hr’g Tr. (Reinman) at 112:4-23, 253:14-18; Hr’g Tr. (Pfister) at 570:17-572:24, 575:2-9. There is nothing in the patent suggesting that those units of data must be a particular vertex or pixel, or subcomponents of either, to qualify as an “element.” Indeed, the 053 patent is almost silent on the identity of “element,” as it uses the term just four times and in varying ways:

In a graphics processing system, it is important to manage and control multiple command threads relating to texture applications. In a typical graphics processing system, *the processing elements*, such as vertices and/or pixels, are processed

[REDACTED]

through multiple steps providing for the application of textures and other processing instructions, such as done through one or more arithmetic logic units (ALU).

....

Generally, the present invention includes a multi-thread graphics processing system and method thereof including a reservation station having a plurality of command threads stored therein. A reservation station may be any type of memory device capable of reserving and storing command threads. Furthermore, a command thread *is a sequence of commands applicable to the corresponding element*, such as pixel command thread *relative to processing of pixel elements* and a vertex command thread relative to vertex processing commands.

....

Therefore, when all of the commands in the command thread have been executed and the command thread is retrievable from the reservation station 202, the command thread may be provided *to a further processing element* (not illustrated) within a graphics processing pipeline.

053 patent at 1:22-28 (emphasis added), 2:36-45 (emphasis added), 3:23-28 (emphasis added).

Even then, the first two passages largely match AMD’s interpretation, while none match Respondents’. So, at a minimum, the term “element” should be construed broadly.

And Dr. Pfister agreed [REDACTED]

[REDACTED]. CIB at 27 (citing, *inter alia*, Hr’g Tr. (Pfister) at 795:4-19). Contrary to Respondents’ assertions, the questions to which Dr. Pfister gave these answers were directed to [REDACTED]

[REDACTED]. See RRB at 13. This leaves little reason to find that [REDACTED]

[REDACTED]

There is also little reason to find that a [REDACTED]

[REDACTED] fails to meet the ordered construction. The patent uses “instruction” sparingly—only twice in relevant context: “[i]n a typical graphics processing system, the processing elements, such as vertices and/or pixels, are processed through

[REDACTED]

multiple steps providing for the application of textures and other processing instructions, such as done through one or more arithmetic logic units (ALU)” and “[u]pon the execution of the associated command of the command thread, the thread is thereupon returned to the station 302 or 304 at the same storage location with its status updated, once all possible sequential instructions have been executed.” 053 patent at 1:24-28, 4:23-27. The first instance indicates a broad, generalized scope; and the second could be fairly interpreted as either a distinct computation (Respondents’ position) or as synonym for “command” (AMD’s position). In any event, the latter does not “specifically distinguish[] between ‘commands’ and ‘instructions’” as Respondents contend (RIB at 23; RRB at 7), and the overall weight of expert testimony supports finding the terms interchangeable. Hr’g Tr. (Reinman) at 112:4-23 (“so that represents the fact that it’s a single instruction, one command, that is being applied to multiple pieces of data”); Hr’g Tr. (Pfister) at 773:1-774:10 (agreeing “[a]n instruction is a command that represents the operation to be performed by the processor”); Hr’g Tr. (Pedram) at 981:20-982:9 (“Q. Sir, yes or no: An instruction is a type of command? A. Yes, it can be.”). Dr. Reinman’s answer of “[t]hat’s one embodiment, sir, yes” to the question, “[s]o the ’053 Patent says you can execute one command by executing a sequence of instructions; right?” (Hr’g Tr. (Reinman) at 278:25-279:2) does not meaningfully take away from this nearly undisputed conclusion. Nor does his opinion that jobs, tasks, shader programs, or waves could be “command threads.” *Id.* at 280:5-284:1.

As observed above, the intrinsic evidence supports a broad meaning. Indeed, Respondents’ complaint over “almost everything could be characterized as sequence of commands” would not seem problematic to the aims of the patent. RIB at 25. And Respondents’ remaining point regarding treatment of [REDACTED] is better addressed in the following section concerning the claimed “at least one memory device” (*see* CRB at 12 (“Respondents incorrectly

Contains Confidential Business Information Subject to Protective Order

claim that AMD's position is that [REDACTED].
Rather, AMD's position is that a memory device that [REDACTED] is a memory device operable to store pixel command threads and vertex command threads.")).

Accordingly, AMD has shown the accused products include the claimed "command threads" in the form of quads or warps.

b. [5a] "at least one memory device . . ."

AMD contends, "at least one memory device comprising a first portion operative to store a plurality of pixel command threads and a second portion operative to store a plurality of vertex command threads" is present in the Accused Products through [REDACTED]

[REDACTED] CIB at 25. To be specific, AMD contends that the limitation is met in either of two ways: [REDACTED]. *Id.* AMD asserts that there is no dispute [REDACTED] exist in the Accused Products and are "memory devices." *Id.* (citing Hr'g Tr. (Pfister) at 580:21-581:2, 784:5-788:14). As for the hierarchy, AMD argues it is considered a unified subsystem within ARM documentation (*id.* at 26 (citing CX-0773C.35)) and operates as described by Dr. Reinman:

[REDACTED]

Hr'g Tr. (Reinman) at 243:8-244:1. AMD further contends that: there is no requirement for the memory portions to be dedicated or fixed; only sequences of instructions need to have a "corresponding vertex or pixel element"; and Respondents have not shown [REDACTED]

[REDACTED] See CRB at 15.

[REDACTED]

Respondents dispute the limitation is met. First, they contend [REDACTED]

[REDACTED] cannot meet the limitation because [REDACTED]

[REDACTED] RIB at 26 (citing CPB at 25; Hr’g Tr. (Pfister) at 586:13-23, 601:6-24); RRB at 12. They add, “there is no evidence or testimony that [REDACTED] [REDACTED] RIB at 26

As for [REDACTED], Respondents argue they

[REDACTED] RIB at 26 (citing CPB at 26-27). Respondents add that [REDACTED] [REDACTED] (*id.* at 27 (citing Hr’g Tr. (Reinman) at 382:6-11); RRB at 10 ([REDACTED] [REDACTED])), and the limitation cannot be met by [REDACTED] [REDACTED] (RIB at 27; RRB at 10). Related to the first point, [REDACTED] do not have “a corresponding vertex or pixel element” as required by the claim, according to Respondents. *Id.* (“[REDACTED] [REDACTED]”).

Respondents also find fault with [REDACTED]

[REDACTED] According to Respondents, “as agreed by both side’s experts, [REDACTED]

[REDACTED] RIB at 27-28 (citing Hr’g Tr. (Pfister) at 582:10-583:6, 573:21-23); RRB at 10 (citing Hr’g Tr. (Pfister) at 570:8-572:10, 582:10-583:6, 573:21-23; Hr’g Tr. (Reinman) at 289:10-14) ([REDACTED] [REDACTED])).

[REDACTED]

Respondents then argue that because [REDACTED]

[REDACTED] there can be no direct infringement; and “AMD did not present any indirect infringement theory with respect to end users.” RIB at 27 (citing Hr’g Tr. (Pfister) at 586:24-588:2). In fact, they allege AMD has completely waived any assertion that [REDACTED] exist in the products for failure to present that assertion in its pre-hearing brief. RIB at 27 (citing CPB at 26-27); RRB at 11.

Relatedly, Respondents argue [REDACTED]

[REDACTED] . . . there is no infringement by any Realtek Product. . . . And AMD never advanced an indirect infringement theory in its contentions based on [REDACTED].” *Id.* at 28 (citing Hr’g Tr. (Pfister) at 580:12-20); RRB at 12.

AMD has not met its burden to show that the Accused Products include “at least one memory device comprising a first portion operative to store a plurality of pixel command threads and a second portion operative to store a plurality of vertex command threads” by a preponderance of evidence. To begin, there is no dispute that the Accused Products operate [REDACTED]

[REDACTED] And, as it is not asserted that [REDACTED]

[REDACTED], then it can be logically inferred that [REDACTED]

[REDACTED]. *See, e.g., Liquid Dynamics Corp. v. Vaughan Co., Inc.*, 449 F.3d 1209, 1220 (“While direct evidence may have made LD's case of infringement stronger, it does not render the case presented to the jury unreasonable as a matter of law. Indeed, the jury could have reasonably inferred from the evidence submitted that each of the tanks infringed the claims.”); CX-1211C.58 (“6.5.1.1 Warp Creation”).

[REDACTED]

Moving on, AMD has persuasively shown that [REDACTED]

[REDACTED]

[REDACTED]. Hr'g Tr. (Reinman) at 243:8-244:1; Hr'g Tr. (Pfister) at 784:22-785:1; CX0-0773C.31

([REDACTED]

[REDACTED]

[REDACTED]). Contrary to Respondents' suggestion, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

CX-0773C.35-36; *see id.* at CX-0773C.22 ([REDACTED])

[REDACTED]

[REDACTED]). [REDACTED]

[REDACTED] JX-0192C ([REDACTED]) at 49:19-50:5; JX-0231C (Vineyard)

([REDACTED])

[REDACTED]).

Respondents confirm as much, stating, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] RRB

at 10 (emphasis added); *see id.* ([REDACTED])

[REDACTED]

[REDACTED]) (emphasis removed and added)). And there is no reason to find [REDACTED]

[REDACTED] do not qualify as “at least one memory device.”

AMD has failed, however, to sufficiently show how “a first portion operative to store a plurality of pixel command threads and a second portion operative to store a plurality of vertex command threads,” is met. In its post-hearing briefing, AMD’s theory is that the “first portion”

and “second portion” are whichever portions of memory the pixel command threads and vertex command threads happen to be located:

Third, Dr. Pfister further criticized Dr. Reinman’s opinion for not identifying the specific first or second portions of the memory device(s) which are operative to store the command threads. Hrg. Tr. (Day 2) (Pfister) at 586:24-587:22. Yet Dr. Pfister identified no support in the record which imports a limitation where the first and second portions have to be in any specific location of a memory device. For example, dependent claim 8 makes it clear that the first and second portions may each encompass the entirety of a memory device, and principles of claim differentiation provide that independent claim 5 is necessarily broader, and encompasses the embodiments of claim 8. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1315 (Fed. Cir. 2005). As Dr. Reinman explained, the requirement of first and second portions provides that the command threads “are stored in different address spaces within the overall memory . . . so that . . . they wouldn’t be overriding each other’s contents,” not that the claim element requires dedicated or specific areas of memory for storing pixel command threads versus vertex command threads. *See* Hrg. Tr. (Day 1) (Reinman) at 243:8-244:19. In effect, Respondents’ insistence that the plurality of vertex and pixel command threads must be stored in a specific location within memory is a rehash of their claim construction proposal for the term “memory device” as meaning “[m]emory with dedicated storage areas for pixel and vertex command threads.” Order No. 31 at 20. As the ALJ noted in the Markman order rejecting Respondents’ proposed construction for “memory device,” “nothing in the claims requires” “either that a particular memory portion only stores a particular type of command thread, or that the memory portion cannot be programmed to store other data – as AMD puts it, that the memory in the device cannot be ‘dynamically allocated.’” *Id.* at 21. Dr. Pfister’s opinion ignores the ALJ’s guidance, and should be rejected.

CIB at 37; *see id.* at 32 (also referring to testimony, “are stored in different address spaces within the overall memory . . . so that . . . they wouldn’t be overriding each other’s contents”). AMD relies on the following expert testimony:

Q. What is the basis for your opinion that they not only store a plurality of pixel command threads and vertex command threads, but that they are operative store – they have a first portion – excuse me, operative to store a plurality of pixel command threads, and a second portion operative to store a plurality of vertex command threads?

A. 

[REDACTED]

Hr'g Tr. (Reinman) at 245:25-246:13; *see id.* at 411:11-22.

The problem, as Respondents point out, is that this [REDACTED] theory does not appear in AMD's prehearing brief. RIB at 27; RRB at 11; *see generally* CPB. It does not appear within the brief's section dedicated to limitation [5a] (*id.* at 23-30) nor does it appear in the section's citations to Dr. Reinman's reports, even assuming that sort of incorporation by reference is fair notice (CX-0593C at ¶¶ 221-226; CX-0597C at ¶¶ 90-102, 103-111, 142-143, 181-186, 189-191). AMD instead appears to assert that [REDACTED] that satisfy the claim limitation. CPB at 30. AMD does not, however, explain what or where those portions are within [REDACTED]. *See id.*

AMD's new [REDACTED] theory is thus not found in its prehearing brief, even under a generous reading. *See* CPB at 23-30. This is an ordinary violation of Ground Rule 9.2, and thus the theory is waived. Order No. 4. Notably, AMD does not dispute the allegations of waiver in its reply brief. *See generally* CRB. "First portion" and "second portion" are important claim elements, both within claim 5 and other claims of the 053 patent. Order No. 31 at 20. There is no reason for AMD to have waited until the evidentiary hearing to reveal this theory, as simple as it is.

And the omission of the theory seems to have been deliberate, because Dr. Reinman opined at his deposition, before prehearing briefs were due, that [REDACTED]

[REDACTED]

[REDACTED] Hr'g Tr. (Reinman) at 291:2-17 (quoting deposition testimony of Dr. Reinman as impeachment). By contrast, Dr. Reinman agreed at the hearing that [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] *Id.* at 288:20-290:8. In short, as Respondents argue, Dr. Reinman essentially conceded that [REDACTED]


[REDACTED] RIB at 27 (emphasis omitted). As AMD has offered no other explanation as to how “first portion” and “second portion” are met in the accused memory devices that is consistent with its prehearing brief, the limitation has not been shown.

c. [5b] “an arbiter . . .”

AMD contends the next limitation, “an arbiter, coupled to the at least one memory device, operable to select a command thread from either of the plurality of pixel command threads and the plurality of vertex command threads,” is met in the Accused Products through a component referred to by AMD as [REDACTED]

[REDACTED]

[REDACTED] CIB at 38-39 (citing, *inter alia*, CX-1211C.13, 78, 91). This component is shown below [REDACTED]



By way of background, Order No. 31 resolved a dispute over whether the arbiter must somehow “actively select” its threads or may merely “select” them. Order No. 31 at 17. It was determined that a plain and ordinary meaning of “select” should apply; and that meaning “presupposes some degree of data processing or computational activity; an arbiter that simply stores the command threads ‘in a display list,’ without routing them, as in [prior art], would presumably not ‘select’ anything and so would not satisfy the claim language.” *Id.*

With that breadth in mind, AMD has met its burden for this limitation. The evidence shows



[REDACTED]

[REDACTED]. CIB at 38-39 (citing Hr’g Tr. (Reinman) at 244:21-245:24, 247:14-248:18; CX-1084C.16; CX-1730C.11, 38-39; CX-0773C.30-31; JX-0192C ([REDACTED]) at 128:14-129:14); CX-1211C.13 (“[REDACTED]”), 90-99. This qualifies as the “routing” discussed in Order No. 31.

Respondents’ objections are not persuasive. First, they argue [REDACTED]. RIB at 28-29; RRB at 13. This is based on two faulty premises already rejected above: [REDACTED]. See RIB at 29.

Respondents also argue that selection of [REDACTED] is not the same as selecting the command thread, and that only the latter will satisfy the claim. RIB at 29-30. This is reasonable in the context of [REDACTED]. But certainly there is nothing preventing the “selection” of a command thread via the selection of [REDACTED]. Dr. Reinman has provided credible testimony that [REDACTED] is both disclosed in the 053 patent and would have been known to those of ordinary skill in the art at the time of the invention. Hr’g Tr. (Reinman) at 236:1-238:8 (discussing, *inter alia*, [REDACTED]); 053 patent at 4:49-5:18 (same). Respondents do not contest either point. See RRB at 13-14. And while they briefly argue “AMD now contends [this theory], for the first time in its Post-Trial Brief” (RRB at 14), [REDACTED] was plainly disclosed before the hearing (CPB at 32 ([REDACTED])).

[REDACTED]

[REDACTED]

[REDACTED])).

Relatedly, Respondents argue that because [REDACTED]

[REDACTED]

[REDACTED] *i.e.*, no selection of a “thread.” RIB at 30 (citing Hr’g Tr. (Pfister) at 598:20-599:8, 601:15-24; Hr’g Tr. (de la Iglesia) at 833:16-836:18); RRB at 13-14. But [REDACTED] constitutes a claimed “command thread,” so what happens *after* [REDACTED] is selected [REDACTED] is irrelevant to this claim element. The system may [REDACTED]

[REDACTED]

[REDACTED]. Indeed, Dr. Reinman offered testimony that [REDACTED]

[REDACTED]. Hr’g Tr. (Reinman) at 242:25-244:1; CX-1724C.7 (“[REDACTED]

[REDACTED]”), 9-10 ([REDACTED]); CX-1211C.90-99 ([REDACTED]). And Mr. de la Iglesia testified that [REDACTED]

[REDACTED] Hr’g Tr. (de la Iglesia) at 834:9-20, 835:5-7 ([REDACTED])

[REDACTED]). So the weight of the evidence shows that [REDACTED]

[REDACTED] which is not very different from at least one 053 patent embodiment. 053 patent at 4:14-27.

Contains Confidential Business Information Subject to Protective Order

Lastly, Respondents argue that [REDACTED]
[REDACTED]. RIB at 30; RRB at 14. Respondents find support
in Dr. Reinman's testimony that [REDACTED]
[REDACTED]. Hr'g Tr. (Reinman) at 382:6-18. Yet which component
ultimately fetches or retrieves the instructions from memory is not germane to [REDACTED]
[REDACTED].

Accordingly, this limitation has been shown in the Accused Products.

d. [5c] "a plurality of command processing engines . . ."

Within the Accused Products, AMD identifies [REDACTED]
[REDACTED] as the "plurality of command processing engines, coupled to the arbiter, each
operable to receive and process the command thread." CIB at 43 (citing Hr'g Tr. (Reinman) at
250:12-255:21). AMD continues, [REDACTED]
[REDACTED] *Id.*
(citing Hr'g Tr. (Reinman) at 254:25-255:21; Hr'g Tr. (Pfister) at 774:16-776:21; JX-0192C
([REDACTED]) at 34:16-35:9, 114:7-10). AMD continues to explain how [REDACTED]
[REDACTED]

[REDACTED]



Id. at 44-45.

AMD has met its burden. The evidence shows [REDACTED] (the claimed “command processing engine”) [REDACTED] (the claimed “command thread”) and is coupled to [REDACTED] (the claimed “arbiter”) which tells [REDACTED] (the claimed “at least one memory device”) [REDACTED] (the claimed “receive and process the command thread”). *See, e.g.*, CX-1211C.13 ([REDACTED]), 17 ([REDACTED]); CX-1724C.9 ([REDACTED]) [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED], 36-37 ([REDACTED]) ([REDACTED]) [REDACTED]

Respondents oppose on a number of grounds. First, they suggest that AMD’s [REDACTED] theory has been waived for failure to appear in the pre-hearing brief under Ground Rule 9.2 RIB at 31 (citing CPB at 36-37), 34 (“To the extent the ALJ permits AMD to advance its untimely

Contains Confidential Business Information Subject to Protective Order

[REDACTED] theory raised for the first time at the hearing . . ."); RRB at 15. To the contrary,

AMD's pre-hearing brief states:

The evidence will show that the command processing engines [REDACTED]
[REDACTED] are each operable to receive and process the command thread. The evidence
will show, for example, that [REDACTED] comprises [REDACTED]
[REDACTED]

CPB at 36. This is fair notice that [REDACTED] constitute the claimed "command processing
engines."

Respondents argue the theory nevertheless fails because [REDACTED]
[REDACTED]

RIB at 34 (citing Hr'g Tr. (Reinman) at 387:10-24). But, as explained above, the evidence shows
[REDACTED]

[REDACTED] And the cited testimony from Dr.
Reinman is inapposite. Hr'g Tr. (Reinman) at 387:10-24.

Respondents then contend that because [REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED] RIB at 35-36 (citing Hr'g Tr. (Reinman) at
387:16-20) (emphasis in original); RRB at 16. Not so. As AMD correctly reasons:

[REDACTED]

Indeed, Claim 5 only requires that the plurality of command processing engines be
operable to process two types of command threads, pixel and vertex. JX-0001 at

Claim 5. This does not mean that to infringe, a command processing engine must be able to process every type of command in a pixel command thread and vertex command thread. See JX-0001 at 2:52-56.

CRB at 18-19. The patent expressly discloses this technique of a single thread being passed between more than one command processing engine:

It is also noted, that in the present embodiment, multiple command operations may be performed by a particular unit 308 or engine 310, but in order to switch a command thread from ALU 308 to a graphics processing engine 310, that command thread must be returned back to the appropriate reservation station 302 or 304 and re-retrieved by the arbiter 306 and thereupon provided to the other unit 308 or engine 310 respectively.

053 patent at 4:41-48 (discussing Figure 4);




FIG. 4

Id. at Fig. 4. So Respondents’ statement that “what Dr. Reinman is now accusing of infringement are two alleged ‘command processing engines’ that are dedicated to processing *different* command threads” (RRB at 16 (emphasis in original)) is inaccurate. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

JX-0192C ([REDACTED]) at 34:20-35:9.

Lastly, Respondents view AMD as pursuing a theory which includes not only [REDACTED] as the “command processing engines,” but also “texture units.” RIB at 36. Respondents argue this is a problem because [REDACTED]

[REDACTED]

[REDACTED] *Id.* (citing Hr’g Tr. (de la Iglesia) at 836:19-837:9; Hr’g Tr. (Reinman) at 385:24-386:3). Even if true as a technical matter, [REDACTED]

[REDACTED] *See generally* CIB; CRB at 20. Thus, infringement does not turn on whether texture units can or cannot process command threads. The point is irrelevant.

Accordingly, limitation [5c] has been shown in the Accused Products.

2. Claim 6

For reference, claim 6 of the 053 patent requires:

6. The graphics processing system of claim 5, wherein the plurality of command processing engines comprises at least one arithmetic logic unit.

053 patent at cl. 6. AMD contends, as testified to by Dr. Reinman, that [REDACTED]

[REDACTED] constitute the claimed “at least one arithmetic logic unit.” CIB at

[REDACTED]

45; Hr’g Tr. (Reinman) at 259:3-8. Respondents do not dispute practice of claim 6 apart from its dependency on claim 5. RIB at 37; *see* RRB at 17-21. In light of the cited evidence and opinion, the additional elements of claim 6 are met in the Accused Products.

3. Claim 9

For reference, claim 9 of the 053 patent requires:

9. The graphics processing system of claim 5, wherein a single memory device comprises the first portion and the second portion of the at least one memory device.

053 patent at cl. 9. AMD contends, as testified to by Dr. Reinman, [REDACTED]

[REDACTED]

[REDACTED] CIB at 46; Hr’g Tr. (Reinman) at 259:10-260:8. Respondents generally dispute the limitation insofar as AMD may point to [REDACTED] as a “single memory device.” RIB at 37; RRB at 17.

AMD has not met its burden. As determined above, AMD has only shown that pixel command threads and vertex command threads are more than likely stored *somewhere* within [REDACTED] without further specificity. AMD has not timely presented how “first portion” and “second portion” are accomplished within [REDACTED]. Thus, Dr. Reinman’s assertion that [REDACTED] is a “single memory device” that can by itself store multiple pixel threads and multiple vertex threads is also inadequate. Hr’g Tr. (Reinman) at 259:10-260:8.

Accordingly, AMD has not shown that claim 9 is met in the Accused Products.

4. Indirect Infringement

With respect to that class of Accused Products which [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] and to the extent they are found to not directly infringe due to [REDACTED], AMD contends “Realtek indirectly infringes the asserted claims of the ’053 Patent on the basis that Realtek knowingly induced and/or actively contributed to other parties’ direct infringement, and that these Realtek Products have no substantial non-infringing use as stand-alone devices [REDACTED].” CIB at 46 (citing Hr’g Tr. (Reinman) at 89:23-92:12; CX-0773C.11, 18, 21-23; CX-0555C.38; CX-0535C.44-56). AMD claims “all of these products [REDACTED] [REDACTED]” *Id.* (citing Hr’g Tr. (Reinman) at 126:23-127:4); *see id.* at 46-47 (citing Hr’g Tr. (Levitt) at 469:13-23, 531:6-9; Hr’g Tr. (Pfister) at 784:5-785:1).

As Realtek observes, however, AMD points to no evidence “of any specific software or components that Realtek’s accused products are allegedly combined with . . . that result in direct infringement.” RIB at 38. It is not enough for AMD to identify [REDACTED] in general as the missing component that [REDACTED], because AMD must also show that the resulting combination directly infringes. *See Met-Coil Sys. Corp., Inc.*, 803 F.2d at 687. In particular, if [REDACTED] is necessary to show infringement because it contains [REDACTED] [REDACTED] But AMD does not even attempt to prove that. *See* CIB at 46-48. AMD has therefore not shown indirect infringement involving any Accused Product.

E. Domestic Industry – Technical Prong

According to AMD’s post-hearing briefing, the DI Products practice claims 5, 6, and 9 of the 053 patent. *See* CIB at 14. For the reasons discussed below, AMD has shown by a preponderance of the evidence that the DI Products practice these claims.

[REDACTED]

1. Claim 5

Claim 5 of the 053 patent is presented above in connection with infringement. The parties dispute each of the three limitations of the claim.

a. [5a] “at least one memory device . . .”

For, “at least one memory device comprising a first portion operative to store a plurality of pixel command threads and a second portion operative to store a plurality of vertex command threads,” AMD contends, as testified to by Dr. Reinman, that “each DI Product includes at least one memory device comprising a memory hierarchy, including at least [REDACTED]

[REDACTED]” CIB at 49 (citing, *inter alia*, Hr’g Tr. (Reinman) at 184:14-186:19, 261:19-265:9). AMD points to its [REDACTED]

[REDACTED] *Id.* at 50 (citing, *inter alia*, Hr’g Tr. (Reinman) at 108:4-111:7, 112:4-113:1, 182:3-183:9).

As for the claimed “first portion” and “second portion,” AMD contends “the memory devices of the [REDACTED]

[REDACTED]

CIB at 51 (citing, *inter alia*, Hr’g Tr. (Reinman) at 112:4-113:1, 182:3-183:9, 261:19-265:9, 282:11-17; CX-0502C.40-43), 51-52 (citing CX-0502C.38-48).

As with infringement, AMD has shown that the DI Products include a memory hierarchy, including [REDACTED]

[REDACTED], which all function together as one system and, in total, contain [REDACTED]

Contains Confidential Business Information Subject to Protective Order

[REDACTED]. See CIB at 52-53. Respondents complain that “AMD’s contentions and Dr. Reinman’s testimony pointed to nearly every possible memory structure in the AMD Products as constituting the claimed ‘memory device’” (RIB at 40), but Respondents have been able to effectively address AMD’s arguments on this element, and the identification of a broad range of components does not offend the plain and ordinary meaning of the claim language.

Moving on, Respondents contend AMD has not met its burden because it failed to identify what constitutes the “first portion” and “second portion” within any particular memory device or the overall memory hierarchy. They argue that AMD has essentially skipped over these elements:

AMD’s arguments regarding the technical prong of DI also fail for each individual claim element. For claim element 5[a], AMD only asserts that [REDACTED]

[REDACTED] CPTB at 52. But this merely restates the claim language without actually identifying any “first portion” for storing a plurality of vertex command threads or a “second portion” for storing a plurality of pixel command threads. AMD cannot simply point to [REDACTED] and allege that the first portion and second portion required by claim element 5[a] must be there somewhere. See *DMI, Inc. v. Deere & Co.*, 802 F.2d 421, 429 (Fed. Cir. 1986) (“Conjecture and speculation are not competent evidence”); *Reckitt Benckiser LLC v. Aurobindo Pharma Ltd.*, 239 F. Supp. 3d 822, 830-31 (D. Del. 2017) (finding no infringement for a claim requiring two portions, when plaintiff merely assumed that the accused products must have both portions in order to function based on unsupported expert testimony).

RRB at 18. As noted above, “first portion” and “second portion” are important claim elements, both within claim 5 and for others, especially asserted claim 9. Order No. 31 at 20. There is some truth to Respondents’ characterization. See generally CIB at 49-53.

However, AMD’s briefing does cite to the testimony of Dr. Reinman where he explained that [REDACTED], and thus are in

their own respective “portions”; *i.e.*, the “first portion” and “second portion” are

happen to be located:

A. . . .



. . . .

Q. So you’ve explained how the memory devices you identified stored plurality pixel command threads and a plurality of vertex command threads. What is the basis of your opinion, Dr. Reinman, that the memory system and devices that you’ve identified are operative, or have a first portion operative to store pixel command threads, and a second portion operative to store a plurality of vertex command threads?

A.



Hr’g Tr. (Reinman) at 263:24-265:9; *see* CIB at 51, 52. This is a sufficient allegation for first and second memory portions in the DI Products. And it is not disputed on technical grounds by Dr. Pfister.

AMD technical documentation would seem to support this scenario, as



[REDACTED]

[REDACTED]

CX-0079C.7-8; *see id.* at 11, 17 (discussing that [REDACTED])

[REDACTED], 23

[REDACTED]

(discussing [REDACTED]

[REDACTED];

[REDACTED]

CX-0502C.40; *see id.* at 12 (“[REDACTED]

[REDACTED]”), 37, 42 (“[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]”), 43-45, 55-56. In

particular, the AMD documentation states [REDACTED]

[REDACTED]

[REDACTED]

CX-0080C.28. Together, this evidence shows it is more likely than not that a plurality of pixel command threads and a plurality of vertex command threads are [REDACTED] stored in the [REDACTED]. Notably, unlike with infringement, Respondents do not allege AMD has waived any theory here. *See* RIB at 38-44; RRB at 11, 17-21.

Respondents also complain that “AMD did [not] meet its burden to show that claim 5[a] is met because it did not tie any of its multiple ‘arbiters’ to any of its hundreds of potential ‘memory devices.’” RIB at 40. This is discussed below in connection with limitation [5b] where the relevant claim language appears. Finally, Respondents contend, “AMD presented no theory of indirect infringement through end-users operating software, so there is no infringement of claim 5 by Respondents.” RRB at 11. To the extent Respondents intended to argue that there is no practice by AMD because end-user software is required to infringe the claims, software is not required as the claim recites a memory device which is merely “operative” to store the variety of command threads. 053 patent at cl. 5.

Accordingly, AMD has shown the DI Products practice limitation [5a].


b. [5b] “an arbiter . . .”

For “an arbiter, coupled to the at least one memory device, operable to select a command thread from either of the plurality of pixel command threads and the plurality of vertex command threads,” AMD contends, as testified to by Dr. Reinman, the DI Products’ [REDACTED] meets the

limitation. CIB at 53-54 (citing, *inter alia*, Hr’g Tr. (Reinman) at 180:20-183:9, 265:11-268:1).

AMD provides the following figure to show the interaction between the [REDACTED] (a.k.a.


[REDACTED] :



CX-0502C.40; *see* CIB at 54; Hr’g Tr. (Reinman) at 265:11-268:1. AMD argues that the [REDACTED]

[REDACTED]

[REDACTED] :



[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

CIB at 54-56. The evidence AMD cites supports its theory and demonstrates the limitation is met.

Respondents oppose on a number of grounds which are not persuasive. First is the position that the claimed arbiter must be able to select a single thread, and [REDACTED] is unable to do so. RIB at 41. This, however, is based on a faulty claim construction, as explained above.

Respondents also take issue with the activity of the [REDACTED], viewing it as a

RIB at 41; RRB at 19-

20. Respondents assert:

[REDACTED]

RIB at 41. Respondents contend Dr. Reinman was impeached on this very point. *Id.* at 42 (citing Hr’g Tr. (Reinman) at 409:20-410:6).

Respondents’ position is not supported by the cited testimony of its expert or AMD documentation. Respondents argue that [REDACTED]

[REDACTED], but Dr. Pfister testified the number is only [REDACTED] (Hr’g Tr. (Pfister) at 739:14-740:12). And only the latter is consistent with the figure reproduced above, where [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]. Thus, it is reasonable to assume that [REDACTED] are

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

CX-0502C.45; *id.* at 46 (“[REDACTED]

[REDACTED]

[REDACTED]”); *see* Hr’g Tr. (Reinman) at 409:11-19. Another document explains, “[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED].” CX-0078C.7.

This [REDACTED] qualify as the “[REDACTED] found sufficient above, in the context of infringement. Order No. 31 at 17. Dr. Reinman’s offhand testimony to the contrary (Hr’g Tr. (Reinman) at 409:20-23 (“[REDACTED] [REDACTED]”)) is certainly relevant and supportive of Respondents’ briefing. And AMD’s attempt to rehabilitate (CRB at 24) was not successful. But Dr. Reinman’s statement is ultimately outweighed by the testimony of Dr. Pfister and AMD documentation.

Finally, as introduced under limitation [5a], Respondents argue that AMD has failed to tie any of the alleged “memory device[s]” to the alleged “arbiter.” RIB at 40. Not so. As shown above, the [REDACTED] within the [REDACTED] receives its pixel and vertex command threads ([REDACTED] [REDACTED] [REDACTED]). This meets the construction of “coupled.” Order No. 31 at 22.

Accordingly, AMD has shown the DI Products practice limitation [5b].

c. [5c] “a plurality of command processing engines . . .”

For “a plurality of command processing engines, coupled to the arbiter, each operable to receive and process the command thread,” AMD contends, as testified to by Dr. Reinman, the DI Products’ [REDACTED] are the “command processing engines.” CIB at 56-57 (citing Hr’g Tr. (Reinman) at 180:20-183:9, 268:2-270:3). AMD explains how each DI Product includes one or more [REDACTED] [REDACTED] [REDACTED]

[REDACTED]

[REDACTED]

Id. at 57. There is no indication that any of these [REDACTED] processes only vertex or only pixel command threads. Thus, AMD has shown the limitation is met in the DI Products.

Respondents present a two-part dispute for this limitation. First, they insist that the full set of [REDACTED] RIB at 43-44. Second, they view the limitation as not met because the [REDACTED]. *Id.*; *see* RRB at 20-21. This, however, is not a requirement of the claims. “[A] plurality of command processing engines, coupled to the arbiter, each operable to receive and process the command thread” means the plurality, as a collective, is operable to receive and process the command thread previously recited as selected by the arbiter. 053 patent at cl. 5. A preponderance of the evidence shows [REDACTED] are so operable—*i.e.*, capable. RIB at 51 (“‘Operable’ is construed to mean ‘capable.’”) (citing *In re Imes*, 778 F.3d 1250, 1253-54 (Fed. Cir. 2015); *ParkerVision, Inc. v. Qualcomm Inc.*, 903 F.3d 1354, 1361 (Fed. Cir. 2018)).

Accordingly, AMD has shown the DI Products practice limitation [5c].

[REDACTED]

2. Claim 6

Claim 6 of the 053 patent is presented above in connection with infringement. AMD contends, as testified to by Dr. Reinman, that [REDACTED]

[REDACTED]

[REDACTED]” CIB at 57; Hr’g Tr. (Reinman) at 270:4-21. Respondents do not dispute practice of claim 6 apart from its dependency on claim 5. RIB at 38-44; *see* RRB at 17-21. In light of the cited evidence and opinion, the additional elements of claim 6 are met in the DI Products.

3. Claim 9

Claim 9 of the 053 patent is presented above in connection with infringement. AMD contends, as testified to by Dr. Reinman, that the “[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]” CIB at 58; Hr’g Tr. (Reinman) at 270:22-271:20, 288:20-289:5. Respondents do not dispute practice of claim 9 apart from its dependency on claim 5. RIB at 38-44; *see* RRB at 17-21. In light of the cited evidence and opinion, the additional elements of claim 9 are met in the DI Products.

Accordingly, the limitation has been shown in the DI Products.

F. Validity

As presented in their briefing, Respondents identify the following invalidity theories for the 053 patent:

Contains Confidential Business Information Subject to Protective Order


Claims	Theory
5, 6, 9	Rendered obvious under 35 U.S.C. § 103 by Stuttard (RX-0416)

See generally RIB at 44-61. Pub. No. WO 00/62182 (“Stuttard”) published on October 19, 2000, and is prior art to the 053 patent under pre-AIA 35 U.S.C. § 102(b). RIB at 45; RX-0416. AMD does not contest this prior art status. *See generally* CIB; CRB. Accordingly, Stuttard qualifies as prior art to the 053 patent at least under pre-AIA 35 U.S.C. § 102(b).

1. Claim 5

Respondents contend claim 5 is rendered obvious by prior art reference Stuttard. Stuttard is entitled “Parallel Data Processing Apparatus,” and teaches “[a] data processing apparatus comprises a SIMD (single instruction multiple data) array (10) of processing elements. The processing elements are operably divided into a plurality of processing blocks, the processing blocks being operable to process respective groups of data items.” RX-0416 at Title, Abstract.

For limitation [5a], “at least one memory device comprising a first portion operative to store a plurality of pixel command threads and a second portion operative to store a plurality of vertex command threads,” Respondents and their expert, Dr. Pedram, look to Stuttard’s cache memory 1024 located within the “Thread Manager” block, as shown below:



RX-0416 at Fig. 4; *see* RIB at 49 (citing RX-0416.132; Hr’g Tr. (Pedram) at 902:4-1). For further background, the Thread Manager 102 provides threads to two controllers, which subsequently provide them to a number of “processing blocks 106” as in the figure below:




FIG. 3

RX-0416 at Fig. 3; *see id.* at 8:15-10:3.

To show that the Thread Manager's cache unit 1024 would contain both pixel command threads and vertex command threads, Respondents primarily point to the following passages:

The thread manager 102 is shown in more detail in Figure 4, and comprises a cache memory unit 1024 for storing instructions fetched for each thread.

....

The threads are assigned priorities relative to one another. Of course, although the example described here has eight threads, any number of threads can be controlled in this manner. At any particular moment in time, each thread may be assigned to any one of a number of tasks. For example, thread zero may be assigned for general system control, thread 1 assigned to execute 2D (two dimensional) activities, and threads 2 to 7 assigned to executing 3D activities (such as calculating vertices, primitives or rastering).

[REDACTED]

RX-0416 at 11:3-5, 16-25; RIB at 49 (citing same), 51 (citing same); RRB at 26-27, 28. Importantly, Respondents reference testimony from Dr. Pedram that a POSITA would know that the final phrasing, “3D activities (such as calculating vertices, primitives, or rastering),” means “calculating vertices requires vertex command thread processing and doing rasterization at least requires doing pixel command thread processing.” RIB at 49 (citing Hr’g Tr. (Pedram) at 902:13-903:22). Moreover, because there are eight threads, and threads 2 through 7 are disclosed as assigned to 3D activities, Respondents and Dr. Pedram reason that this group would include more than one pixel thread and more than one vertex thread “because . . . there could be any number of these coming in at any time being present.” *Id.*

Respondents view AMD as ignoring this teaching of Stuttard when presenting its opposition, and otherwise requiring an “at the same” time requirement which does not exist in the claim. RIB at 51; RRB at 25. Nevertheless, assuming the “pixel command threads” and “vertex command threads” must be stored at the same time, Respondents argue Stuttard’s system is “capable” of such (RIB at 50-51) and it otherwise would have been obvious to a POSITA to do so to achieve a number of benefits (*see id.* at 50 (citing Hr’g Tr. (Pedram) at 903:23-905:5), 51-52 (citing Hr’g Tr. (Pedram) at 904:23-905:5; RX-0416 at 12:1-7, 10:19-31, 10:35-11:1, 11:16-33, 12:1-7); RRB at 27 (citing same)).

Respondents have not met their clear and convincing burden to show the memory limitation is either disclosed in Stuttard or an obvious modification thereof. First, it is a correct claim construction that “at least one memory device comprising a first portion operative to store a plurality of pixel command threads and a second portion operative to store a plurality of vertex command threads” means the memory device(s) are operable to store both pluralities at the same time. A plain reading of the claim and the patent’s specification support this construction for the



same reasons they support the storage of complete threads (*i.e.*, all instructions at the same time), as argued by Respondents and determined above. *See, e.g.*, RIB at 18 (arguing that without storing all instructions at the same time “there would be no minimum capability required of the memory device”); 053 patent at 2:65-3:3, 3:40-44, 3:50-53, 3:62-4:51, 5:10-16, 5:43-56, Fig. 2, Fig. 4, Fig. 5.

As to Stuttard’s disclosure, the passage most relied on by Respondents and excerpted above— “[f]or example, thread zero may be assigned for general system control, thread 1 assigned to execute 2D (two dimensional) activities, and threads 2 to 7 assigned to executing 3D activities (such as calculating vertices, primitives or rastering)” (RX-0416 at 11:19-25)—only teaches that “[a]t any particular moment in time,” some threads may be for 3D activities while others may be for system control or 2D activities. It does not *clearly* teach that different kinds of 3D threads, namely vertex and pixel, are stored in different portions of the memory device at the same time, let alone a plurality of each.

And there is reason to find that vertex and pixel threads are not so stored at the same time. As pointed out by AMD and its expert, Dr. John Hart, Stuttard discloses distinct phases of 3D graphics processing (vertex operations, then binning, then pixel operations) and its processing blocks (each containing an array of processing elements (“PEs”)) acting “block parallel,” which either means or strongly suggests that vertex threads run to completion before pixel threads become stored for processing. *See* CIB at 61-62 (citing, *inter alia*, RX-0416 at 3:21-22, 41:12-24, Fig. 9; Hr’g Tr. (Hart) at 1250:7-1251:17, 1252:23-1253:19), 64-65; CRB at 25-26, 28; *see also* Hr’g Tr. (Hart) at 1251:23-1252:11, 1253:23-1254:1, 1254:19-1255:11). As a result, AMD reasons:

[T]here is no instance in where the alleged memory device of Stuttard would store both a plurality of pixel command threads and vertex command threads. *See id.*

[REDACTED]

During the vertex processing phase, the alleged memory device of Stuttard will store vertex command threads, but not pixel command threads. *Id.* Likewise, during the pixel processing phase, the alleged memory device of Stuttard will store pixel command threads, but not vertex command threads.

CIB at 65. Indeed, Stuttard states for its Thread Manager 102 shown in Figure 4 that “[i]n this particular example the same instruction stream is supplied to all of the processing blocks in the system.” RX-0416 at 11:33-35; *see id.* at 22:8-10 (describing the pipeline as “[f]rom this point, a given processing block operates on data associated with a single region of the display.”), 22:26-32 (in a final stage of the pipeline, “the processing for the primitive can be carried out only on those pixels occurring inside the primitive. The remainder of the PEs in the processing block do not take any further part in the processing until that primitive is processed.”). Respondents and Dr. Pedram do little to dispute this phased operation beyond relying on the “each thread may be assigned to any one of a number of tasks” passage evaluated above, which is not conclusive one way or the other. *See generally* RIB at 49-50, 51; RRB at 21-22, 24-25 (citing RX-0416 at 11:15-25); Hr’g Tr. (Pedram) at 983:4-23, 989:20-996:6, 1000:13-1001:22. So the possibility that Stuttard employs a phased approach to pixel thread processing and vertex thread processing dispels a clear and convincing case that the limitation is disclosed.

There is no clear and convincing case of single-reference obviousness either. Identifying a motivation for a POSITA to alter or combine the prior art is critical to the inquiry. *See InTouch Techs., Inc. v. VGo Communs., Inc.*, 751 F.3d 1327, 1347, 1351 (Fed. Cir. 2014). Dr. Pedram refers only vaguely to “efficiencies” and without reference to the specific memory which must do the storage—cache unit 1024:

Q. Did you also form an opinion as to whether or not this particular memory limitation, the full limitation shown, was obvious?

A. Yes.

[REDACTED]

A person of skill in the art looking at the Stuttard system and knowing that the command processing engine, these processing blocks of Stuttard, which is capable of doing both pixel and vertex data processing, must have received the corresponding command threads from someplace, and that someplace is of necessity the memory, which stores these threads.

So, therefore, the memory that stores the threads must be storing the pixel command threads and the vertex command threads, and a plurality of each, because there's at least eight different processing blocks in the Stuttard system. So at the same time, a person skilled in the art would understand that the advantages of being able to store pixel and vertex command threads at the same time, we think the memory system in terms of the efficiencies that it offers, run time efficiencies of the system overall. For the record, the references here, if you can go back, is page 20, line 23, to page 21, line 7. It's a long excerpt that describes the functioning of these processing blocks. If you look at that, you will see that it does -- the processing blocks all bring both pixel and vertex data processing. And then there is additional reference, page 12, lines 1 to 7.

Q. That reference is regarding the efficiency advantages that you discussed?

A. Correct, sir.

Q. All right. Moving to the next limitation. . . .

Hr'g Tr. (Pedram) at 903:23-905:6 (citing RX-0416 at 20:23-21:7, 12:1-7). This testimony is not particularly persuasive on its own (nor are its cited portions of Stuttard), but it also fails to accommodate Stuttard's possible phased operation, which would not seem to benefit from storing vertex and pixel threads together at the same time. According to Dr. Hart, "you would never have pixel threads in there while you were processing vertex threads. . . . And you would never have vertex threads in there while you were processing pixel threads because you would be finishing out the current frame with those pixel threads." Hr'g Tr. (Hart) at 1255:3-8. Respondents cite additional passages from Stuttard to supposedly supply a motivation, but they are not germane to storing vertex and pixel threads together. *See* RIB at 51-52 (citing RX-0416 at 12:1-7, 10:19-31, 10:35-11:1, 11:16-33), RRB at 27 (citing same). Similarly, Respondents view the 053 patent as teaching that "storing a plurality of vertex command threads and a plurality of pixel command

threads in different portions at the same time was well known in the art to a POSITA.” RIB at 52 (citing 053 patent at 1:24-28, 42-64, 2:39-41). But, again, the passages do not relate to storing vertex and pixel threads together. RIB at 52 (citing [053 patent] at 1:24-28, 1:42-64, 2:39-41); RRB at 28 (citing 053 patent at 1:24-28, 1:32-64, 2:39-41). The *Graham* factors, including indicia of non-obviousness, need not be considered here. Nor do the remaining claim limitations, which are challenged by AMD on similar grounds. See CIB at 61, 67-69; CRB at 29-30; RIB at 52 (“[t]he only dispute, again, is whether Stuttard’s arbiter can select between a plurality of vertex command threads and a plurality of pixel command threads ‘at the same time.’”).

Accordingly, Respondents have not shown limitation [5a], and thus claim 5, would have been obvious over Stuttard.

2. Claim 6

Respondents contend “wherein the plurality of command processing engines comprises at least one arithmetic logic unit” is disclosed in Stuttard where it describes an “arithmetic logic unit (ALU) 214” as within the processing unit 1061. RIB at 56 (citing RX-0416 at 47:24-30, Figs. 3, 11). The feature is disclosed as alleged. RX-0416 at 47:24-30, Figs. 3, 11. And Respondents correctly represent that “Dr. Hart did not dispute that Stuttard discloses an ALU as part of its processing blocks.” RIB at 56; *see generally* CIB at 58-69; CRB at 25-31.

Accordingly, Respondents have shown the additional limitation of claim 6 is disclosed in Stuttard, but claim 6 is not rendered invalid in light of its dependence on claim 5.

3. Claim 9

Respondents contend “wherein a single memory device comprises the first portion and the second portion of the at least one memory device” is met by cache unit 1024 in Stuttard, either as disclosed or obvious in view thereof. RIB at 56-57. Respondents summarily state, “[b]ecause this

[REDACTED]

cache unit comprises a single device, this claim is satisfied for the same reasons as element 5[a].”
Id. at 57 (citing RX-0416 at Fig. 4; Hr’g Tr. (Pedram) at 914:8-25, 915:2-5).

As determined above, Respondents have not clearly and convincingly shown that cache unit 1024 is operable to store vertex command threads and pixel command threads at the same time in a single memory device. Accordingly, claim 9 has not been disclosed or shown to be obvious.

V. U.S. PATENT NO. 8,854,381

A. Level of Ordinary Skill in the Art

In Order No. 31, I determined that a person having ordinary skill in the art of the 381 patent at the time of invention:

[W]ould have a bachelor degree in electrical engineering, computer engineering, computer science, or a related field, and at least two years of experience in ‘parallel computing or graphics processing including developing, designing, or programming software or hardware for graphics processing or parallel processing units, graphics accelerators[,] or other graphics processing systems,’ with additional education substituting for experience and vice versa.

Order No. 31 at 9. The parties do not challenge this definition (*see generally* CIB; CRB; RIB; RRB), and it is applied throughout this initial determination.

B. Claims-at-Issue

Claims 15-17, 19, and 20 of the 381 patent are at issue in this investigation, either through allegations of infringement or domestic industry technical prong. *See* CIB at 3. These claims are reproduced below, along with the parties’ agreed limitation annotations:

15. A computing system comprising:

[15a] a memory;

[15b] a first processing unit; and

[15c] a bus coupled to the memory and the first processing unit,

Contains Confidential Business Information Subject to Protective Order

[15d] wherein the first processing unit comprises: a plurality of engines configured to receive, from a scheduling module associated with a second processing unit, a plurality of tasks and to load state data associated with each of the plurality of tasks; and

[15e] a shader core configured to receive the plurality of tasks from at least one of the plurality of engines and to execute a first task from the plurality of tasks while executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks.

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

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

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.

17. The computing system of claim 15, 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.

....

19. The computing system of claim 15, wherein the first task comprises a graphics-processing task and the second task comprises a general-compute task.

20. The computing system of claim 15, wherein the shader core is configured to receive the plurality of tasks based on a priority associated with each of the plurality of tasks.

381 patent at cls. 15-17, 19, 20 (annotated).

C. Claim Construction

As part of the *Markman* process, the following claim terms of the 381 patent were construed, either as-agreed between the parties or determined by Order No. 31:

Claim Term	Construction
“shader core”	“an array of processing elements”
“task”	“unit of processing work or portion of an application to be executed”
“engine” / “plurality of engines”	Not limited to hardware-only or software-only

“load”	Does not mean “extract and store”
--------	-----------------------------------

See Order No. 31 at 36, 38, 41, 42. Within their briefing, the parties identify a number of additional claim construction disputes, discussed below.

1. “a plurality of engines configured to receive . . . a plurality of tasks and to load state data associated with each of the plurality of tasks”

For the phrase “a plurality of engines configured to receive . . . a plurality of tasks and to load state data associated with each of the plurality of tasks,” the parties have numerous disputes. First, AMD contends that *each* engine receives at least one task, and *each* engine loads state data associated with that task. *Id.* at 76-77 (citing Hr’g Tr. (Reinman) at 131:8-131:25; Hr’g Tr. (Hart) at 1264:20-1265:9; 381 patent at 4:48-51, 6:1-6, 9:34-36, 6:65-7:19, 9:47-49, 10:25-39, Figs. 2, 3A, 3B, 4). AMD cautions, however, that it is not required that each engine receive *a plurality* of tasks as suggested at trial by Respondents’ expert, Dr. Pfister. See CIB at 78-79 (citing Hr’g Tr. (Pfister) at 706:13-15, 709:4-11). Respondents view AMD’s requirements as improperly repeating language already present in the claim. See RIB at 73; RRB at 39-40.

AMD is not persuasive here. The limitation refers to “a plurality of engines” as receiving “a plurality of tasks.” In a plain and ordinary sense, this means the plurality—as a group—do the receiving. If there was subsequent language regarding “each engine,” for example, then AMD’s position would have more merit. But there is not. Similarly, as written, it is the entire group of engines which loads the state data—to be evaluated as a group. There is no reference to what each engine must do. This clear language of the claims controls. *Interactive Gift Express*, 256 F.3d at 1331 (“In construing claims, the analytical focus must begin and remain centered on the language of the claims themselves, for it is that language that the patentee chose to use to ‘particularly point [] out and distinctly claim [] the subject matter which the patentee regards as his invention.’”);

[REDACTED]

CRB at 37 (arguing Respondents’ interpretation of related term is “not supported by plain reading of the claim”).

AMD next contends, “in view of the specification, a POSITA would understand that the claimed engine is required to “service its own command buffer” in addition to the functions already explicitly recited within the claim of “receive . . . tasks” and “load state data.” CIB at 75, 77-78 (citing, *inter alia*, 381 patent at 2:10-19, 10:25-34; Hr’g Tr. (Reinman) at 132:5-24).¹ AMD argues that a connection to a command buffer is “imposed by the intrinsic record, which differentiates the patented system relative to conventional serialized systems having a single command processor/command buffer.” CRB at 36-37.

Respondents oppose (RIB at 71-73; RRB at 38-39), and rightly so. The claim language references neither command buffers nor servicing them. And on its face, AMD’s writing-in of the structural element of a buffer is not only an improper attempt to define an “engine” by how it is used and not what it is (*see Phillips*, 415 F.3d at 1312-1313 (holding claim construction is a matter of determining the meaning of terms selected by the patentee)), but also an improper attempt to read in a limitation from the specification into the claims (*Markman*, 52 F.3d at 980 (“The written description part of the specification itself does not delimit the right to exclude. That is the function and purpose of claims.”); *Hill-Rom Servs., Inc. v. Stryker Corp.*, 755 F.3d 1367, 1371 (Fed. Cir. 2014) (“[W]e do not read limitations from the embodiments in the specification into the claims.”)). *See* Hr’g Tr. (Reinman) at 132:1-24 (discussing the specification’s recitation of the advantages of multiple command buffers). So AMD’s additional structural element is not accepted.

¹ For additional support on this issue, and on other issues throughout its briefing, AMD cites expert report content and expert deposition testimony which have not been admitted in evidence and thus have not been considered. *E.g.*, CIB at 77 (citing CX-1227C), 78 (citing CX-0674C); *see* RRB at 39 n.4.

[REDACTED]

For their part, Respondents propose requiring that each “engine” “process data.” RIB at 68. For support, they point to the testimony of Dr. Pfister and excerpts from the 381 patent describing how the engines serve as a “command processor” of the GPU and otherwise “process” specific types of tasks such as “real-time low latency” or “3D.” *Id.* (citing Hr’g Tr. (Pfister) at 609:3-13, 706:5-23, 801:19-21; 381 patent at Figs. 2-4, 6:65-7:21, 1:47-49); RRB at 37. AMD, on the other hand, argues there is no intrinsic or extrinsic support behind “requiring each engine to perform ‘processing of data’” beyond the testimony of Dr. Pfister, which is inadequate. CIB at 79 (citing *Phillips*, 415 F.3d at 1318 (“Conclusory, unsupported assertions by experts as to the definition of a claim term are not useful to a court.”)); CRB at 37.

No party argues the patentees have given a special definition to “engine” or otherwise disclaimed its scope. Thus, the plain and ordinary meaning controls. *GE Lighting Sols.*, 750 F.3d at 1309. As to the plain and ordinary meaning, there is some intrinsic evidence supporting the idea that engines “process” data or tasks. *See, e.g.*, 381 patent at 7:2-21. But the difficulty lies in what appears to be Respondents’ further requirement that “processing” means computation or execution of instructions upon that data, as opposed to simply storing data, tracking progress, or routing information, or just the claimed “receiv[ing] . . . tasks” and “stor[ing] state data.” *See* Hr’g Tr. (Pfister) at 706:9-707:20 (distinguishing the activities of memories, registers, state machines, and front end units to argue they are non-engines). The 381 patent’s simplistic references to engines or other components “processing” data and tasks could include any of these activities. Indeed, there is little description of what the engines actually do with the tasks given to them by the OS scheduler beyond handing them off to the shader core (where actual SIMD execution units reside) or other “pre-processing” elements. *See, e.g.*, 381 patent at 3:1-11, 4:46-61, 5:4-28, 5:65-6:7, 6:42-

[REDACTED]

54, 6:65-7:19, 7:55-62, 7:66-8:7, 9:31-38, 9:46-52, 10:33-41, Figs. 3A, 3B, 4. For these reasons, Respondents’ proposed construction is not accepted.

Lastly, Respondents contend that “plurality of engines” “should not be construed to encompass a hardware-only embodiment.” RIB at 69-71; RRB at 37-38. Respondents acknowledge Order No. 31’s prior determination that “engines” were not limited to either hardware-only or software-only implementations (*id.* at 69), but argue that there is “no written description of the plurality of engines being implemented only in hardware” and “no disclosures in the specification of the ‘plurality of engines’ as being anything other than ‘virtual’ engines” (*id.* at 69-71 (citing, *inter alia*, 381 patent at Abstract, 3:1-6, 5:29-32, 6:1-4, 6:67-7:4, 8:8-17, 9:45-52, 10:33-39; Hr’g Tr. (Pfister) at 609:3-611:8; Hr’g Tr. (Pedram) at 920:18-922:6), 69 n.7). AMD opposes the construction, viewing the matter to have been previously settled by Order No. 31. *See* CRB at 37 (citing Order No. 31 at 39).

Respondents’ construction is not accepted. First, the chosen phrasing is confusing and would be difficult to apply in questions of infringement or validity. In practice, “should not be construed to encompass a hardware-only embodiment” means the engine must have a software component. Respondents have not shown that the plain and ordinary meaning of “engine” requires a software component, nor do they argue inventor lexicography or prosecution disavowal. *See* RIB at 69-71; RRB at 37-38; *Thorner*, 669 F.3d at 1366-67. And it has not been shown by clear and convincing evidence that construing “engines” to cover hardware-only implementations (in addition to software-only, or a combination thereof) would violate the written description requirement of 35 U.S.C. § 112 so as to limit its scope. *Pfizer, Inc. v. Teva Pharms., USA, Inc.*, 429 F.3d 1364, 1376 (Fed. Cir. 2005) (“Nevertheless, Ranbaxy has not presented sufficient evidence for us reasonably to infer that, unless ‘saccharides’ means ‘sugars’ or at least does not

[REDACTED]

encompass polysaccharides, claim 1 and 16 would have been considered by the Patent and Trademark Office (‘PTO’) to be invalid.”); *see generally Karsten Mfg. Corp. v. Cleveland Golf Co.*, 242 F.3d 1376, 1384 (Fed. Cir. 2001) (declining to construe claim terms to uphold validity over prior art previously determined by lower court to anticipate the claim); RIB at 69 n.7 (citing *Karsten*).

2. “configured to . . . execute a first task . . . while executing a second task”

Within “execute a first task . . . while executing a second task,” the parties have disputes over the “while” and “configured” terms. Respondents contend that “while” means “simultaneous” in a plain and ordinary sense, but then specify that “alternating between two tasks” is not sufficient, nor is mere overlap. *See* RIB at 62-63. For support, Respondents primarily look to the prosecution history of the 381 patent where the patentee amended the prior phrase “substantially in parallel” to “while”:

Here, the Examiner initially rejected claim 15 as indefinite, finding that the description of “a shader core configured to execute the plurality of tasks *substantially in parallel*” was inadequate because the term “substantially” was “not definite.” JX-0009.124. In response, AMD amended the claim to strike the claim term “substantially in parallel” and replace it with “execute a first task . . . *while executing* a second task.” JX-0009.141, 146 (emphasis added). This is strong evidence of how AMD understood the invention it sought to patent. *See Phillips*, 415 F.3d at 1317. Since the term “substantially in parallel” was not sufficiently definite, AMD chose to replace [it] with “while executing,” which in light of the prosecution, specification, and the plain meaning of “while” must be construed to mean “simultaneous” execution of the first task and second task.

RIB at 63-64 (emphasis in original). Respondents argue that when claims have been narrowed in this way, they cannot be later read to cover that which was disclaimed. *See id.* at 64-66 (citations omitted). In addition, Respondents refer to embodiments in the specification where “the shader core of the invention is configured to execute two tasks simultaneously (*e.g.*, on the same clock cycle) by executing the first and second task on different processing units or different subsets of

[REDACTED]

processing elements, which allows the execution of the tasks to proceed exactly in parallel.” *Id.* at 66-67 (citing 381 patent at 1:58-2:1, 5:17-39, 10:47-50).

AMD takes issue with “execution of the first task and second task must *begin and end* at the same time” as contrary to the intrinsic evidence and the plain and ordinary meaning. CIB at 81 (emphasis in original). More specifically, AMD argues the limitation “only requires that there be some overlap in their execution, which is the plain and ordinary meaning of the term ‘while.’” *Id.* (citing Hr’g Tr. (Reinman) at 157:10-158:6; 381 patent at 5:22-25). AMD criticizes Respondents’ strict interpretation as it would “improperly exclude the embodiment in which the two tasks have different lengths, or where the tasks are of equal durations but the first task has a smaller time allocation compared to the second task.” *Id.* at 82 (citing *Vitronics*, 90 F.3d at 1583). AMD also rejects Respondents’ view of the prosecution history, urging the amendment from “substantially in parallel” to “while” was a “rephrasing” or “non-substantive.” *Id.* at 83 (citing *Festo Corp. v. Shoketsu Kinzoku Kogyo Kabushiki Co.*, 535 U.S. 722, 736-37 (2002)).

Respondents are not persuasive. They state, “[t]he Examiner correctly determined a POSITA would not be sure what ‘substantially’ in parallel or ‘substantially’ simultaneously means.” RIB at 66. But given that the examiner provided no explanation for why “substantially in parallel” was indefinite (JX-0009.124), it has not been shown that the amendment to “while” was a narrowing one (*see, e.g.*, RRB at 35-36 (citing district court and unreported Federal Circuit cases where surrounding circumstances showed amendment was narrowing)). Thus, the requisite clear and unmistakable disavowal, from broader to narrower scope, has not been shown. *Thorner*, 669 F.3d at 1366-67.

Moreover, none of Respondents’ examples from the 381 patent specification disclose or suggest any importance to task execution beginning and ending at the same time (*i.e.*, exact

[REDACTED]

overlap), which—given the art and diversity of jobs/tasks/threads that are issued to GPUs—would not seem to be beneficial anyhow. *See* RIB at 66-67 (citing 381 patent at 1:58-2:1, 5:17-39, 10:47-50); *see, e.g.*, 381 patent at 2:47-56 (contrasting tasks with varying execution time lengths), 4:48-55 (distinguishing graphics processing tasks from general compute tasks), 5:34-39 (disclosing possibility of some single tasks to fully consume the resources of the GPU), 7:5-19 (disclosing variety of task types), 10:47-51 (full use of shader core by one task); Hr’g Tr. (Pfister) at 566:16-23 (stating, generally, there could be one or two jobs sent to the GPU, broken down into “maybe hundreds of tasks”).

Additionally, AMD raises the valid point that the patentee apparently viewed the embodiments of “spatial partitioning” and “temporal partitioning” as both meeting the meaning of “substantially” parallel or simultaneous execution, but “temporal portioning” would be excluded under Respondents’ construction. *See* CRB at 33-34. The patent states:

To enable the single shader core to (substantially) simultaneously process tasks from two or more different virtual engines, resources of the shader core are partitioned in space and/or time. To achieve spatial partitioning, for example, a first (*e.g.*, low-latency) task from a first virtual engine is issued to a first subset of processing elements (SIMDs) of the shader core, and a second (*e.g.*, standard graphics) task from a second virtual engine is issued to a second subset of processing elements (SIMDs) of the shader core. To achieve temporal partitioning, for example, the first and second tasks share a percentage of time of the processing elements (SIMDs) of the shader core. In an embodiment, the GPU includes a scheduling module to schedule the tasks from the two or more different virtual engines for execution on the shader core.

381 patent at 5:14-28; *see id.* at 8:8-22.

Lastly, Respondents contend “configured to” claim language refers to the design of an apparatus to perform an action and not merely for the action to occur some of the time. *See* RIB at 67-68 (citations omitted); RRB at 36-37. Specific to claim 15, Respondents contend:

In the context of claim 15, the claim language “a shader core *configured to* ...” means it is insufficient that the shader core is simply “capable of” executing a first

task and a second task simultaneously. Rather, the shader core must be specifically programmed or designed to cause a first task and a second task to be executed simultaneously—and not just any two tasks, but a first task and a second task from the same plurality of tasks received by the plurality of engines. If a first task and a second task from a plurality of tasks are occasionally executed at the same time, and it happens as a matter of probability or chance, this is nothing more than a shader core that is capable of executing a first task and a second task simultaneously. Claim 15 clearly requires something more—that the shader core be designed or programmed to cause the simultaneous execution of two tasks from the same plurality of tasks.

RIB at 67-68. AMD views this as an untimely argument, but regardless, states, “there is not much dispute between the parties [but] to the extent Respondents are suggesting that the claimed shader core must always execute two tasks at the same time at each and every instance, such a requirement is not imposed by a construction of the term ‘configured to’ to mean ‘designed to,’ ‘made to,’ or ‘programmed to.’” CRB at 35-36.

As it does not appear that Respondents contend that “configured to” means the claimed action must happen each and every time (without exception), there is no dispute between the parties. And it is correct that “configured to,” as used in claim 15, denotes a design or implementation that achieves the stated effect(s). No further construction is necessary. *See, e.g.*, RRB at 36-37 (immediately transitioning to discussion of non-infringement by Accused Products).

3. “each of the plurality of tasks”

Respondents next contend that “a plurality of engines configured to receive . . . a plurality of tasks and to load state data associated with each of the plurality of tasks,” means “configured to receive . . . [one or more pluralities] of tasks.” RIB at 73. More specifically, or perhaps additionally, Respondents assert that each individual engine “must be configured to load state data for each of the pluralities of tasks that are received.” *Id.* at 74. Respondents reason:

If the plurality of engines were intended to only receive a single “plurality of tasks” and load state data associated with that single plurality of tasks, the claim would have read “load state data associated with each of the tasks in the plurality of tasks.” But it is a well-settled principle that courts must “construe the claim as written, not

[REDACTED]

as the patentees wish they had written it. *Ecolab, Inc. v. FMC Corp.*, 569 F.3d 1335, 1344-45 (Fed. Cir. 2009). The claim as written requires that the plurality of engines be configured to load state data for each plurality of tasks. Thus, each engine must be configured to load state data for at least two tasks (*i.e.*, there must be two pluralities of tasks present in the plurality of engines). If each engine is only “configured to” or “designed to” receive and load state data for a single task and is not also designed to receive and load state data for a plurality of tasks, it does not satisfy the literal requirements of the claim.

RIB at 74. Respondents do not view their proposal as excluding embodiments where “the OS scheduler issuing a first task to one virtual engine and a second task to a second virtual engine.” *Id.* (“Nothing in Respondents’ proposed construction excludes such a scenario.”). Respondents do assert, however, that “[a]n engine that is not configured to receive and load state data for a plurality of tasks (as well as a single task, depending on what is issued to it by the scheduling module) does not meet the plain language of the claim.” *Id.* at 75. Respondents interpret a single question and answer from AMD’s expert, Dr. Reinman, as suggesting there is no dispute here. *Id.* (citing Hr’g Tr. (Reinman) at 398:21-24 (“Q. In addition to receiving a plurality of tasks, each engine has to also load state data for a plurality of tasks, correct? A. Yes, sir.”)).

AMD disagrees and argues Respondents’ proposal is not a plain reading of the claim and otherwise conflicts with the specification’s embodiment “in which each engine receives only one task at a given point in time.” CIB at 79 (citing 381 patent at 9:34-36)); CRB at 37-38. AMD clarifies, “the term ‘each’ in ‘load state data associated with each of the plurality of tasks’ modifies the term ‘tasks,’ rather than imposing a requirement that there be multiple pluralities.” CRB at 37-38. AMD also views its expert’s testimony as limited to the context of the operation of the Accused Products, and not claim construction which was discussed earlier in the hearing. *Id.* at 38 (citing Hr’g Tr. (Reinman) at 131:8-25, 398:12-24).

Respondents’ construction is not adopted. To “load state data associated with each of the plurality of tasks,” means, in the context of the specification, to load state data associated with

[REDACTED]

each of the tasks that make up the plurality. The patent is clear, and its purported benefits derive from the fact that “state data” belong to tasks on a task-by-task basis. *See, e.g.*, 381 patent at 2:2-56 (discussing the problem in switching out state data in memory when individual tasks are switched out of the shader core). There is no disclosure, on the other hand, of “state data” pertaining to a plurality of tasks, which greatly undercuts Respondents’ view that “each engine must be configured to load state data for *each* of the pluralities of tasks that are received.” RIB at 74 (emphasis in original). In other words, “each” in the claim modifies “task,” not “plurality of tasks.” And any extrinsic expert testimony does not override this clear intrinsic reading. *Phillips*, 415 F.3d at 1318 (“We have viewed extrinsic evidence in general as less reliable than the patent and its prosecution history in determining how to read claim terms, for several reasons.”).

D. Infringement

According to AMD’s post-hearing briefing, the use, manufacture, or sale of Accused Products infringe claims 15-17, 19, and 20. *See* CIB at 12, 13. Of these, claim 15 is independent, and claims 16, 17, 19, and 20 depend from claim 15. For the reasons discussed below, AMD has shown infringement of these claims by the Accused Products.

1. Claim 15

For reference, claim 15 of the 381 patent requires:

15. A computing system comprising:

[15a] a memory;

[15b] a first processing unit; and

[15c] a bus coupled to the memory and the first processing unit,

[15d] wherein the first processing unit comprises: a plurality of engines configured to receive, from a scheduling module associated with a second processing unit, a plurality of tasks and to load state data associated with each of the plurality of tasks; and

[REDACTED]

[15e] a shader core configured to receive the plurality of tasks from at least one of the plurality of engines and to execute a first task from the plurality of tasks while executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks.

381 patent at cl. 15.

The primary disputes between the parties concern the “memory,” “plurality of engines,” and “shader core” in limitations [15a], [15d], and [15e], as well as whether indirect infringement has been shown or is even necessary given the makeup of certain Accused Products. These are discussed below. For those remaining limitations which are not in dispute ([15b], [15c]), the Accused Products have been shown to meet them as alleged in view of the testimony of Dr. Reinman. *See* CIB at 86-88 (citing Hr’g Tr. (Reinman) at 127:10-129:10); CRB at 42-43 (citing Hr’g Tr. (Reinman) at 125:25-126:11, 127:10-19, 128:12-129:1).

a. [15a] “a memory”

Limitation [15a] reads, “a memory.” 381 patent at cl. 15. AMD contends the limitation is met in two ways. CIB at 85. AMD’s first theory is that for all Accused Products the “memory” comprises [REDACTED]. *Id.* (citing, *inter alia*, Hr’g Tr. (Reinman) at 125:25-126:11; CX-0773C.21). AMD’s second theory is that for [REDACTED]
[REDACTED]
[REDACTED] *Id.* (citing, *inter alia*, Hr’g Tr. (Reinman) at 126:15-22; CX-0889C-CX-0893C; JX-0276C (Zeng) at 10:4-15)).

AMD identifies the only dispute as [REDACTED]
[REDACTED]. CIB at 85 (citing Hr’g Tr. (Pfister) at 784:8-15). As relevant to indirect infringement, discussed below, AMD claims, “[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]

[REDACTED]’ *Id.* (citing, *inter alia*, Hr’g Tr. (Reinman) at 126:23-127:4; JX-0251C (wang) at 47:7-19; JX-0143C (Wei) at 111:8-13, 114:20-115:9).

There is no dispute from Respondents that [REDACTED] constitutes a “memory” or [REDACTED] is a “memory.” RIB at 77. Respondents also do not dispute that [REDACTED]. *Id.* But they do contend that [REDACTED] and as the claimed “computing system” has been identified as [REDACTED], there can be no direct infringement—even for [REDACTED]. *Id.* ([REDACTED])

[REDACTED]

[REDACTED]; RRB at 42-43 (citing CPB at 127).

There being no dispute, and in light of the testimony from Dr. Reinman, AMD has shown that all Accused Products include the claimed “memory” through [REDACTED]. AMD has also shown that [REDACTED] include the claimed “memory” through [REDACTED]. And while AMD has perhaps committed an oversight in identifying [REDACTED] as the claimed “computing system” for claim 15’s preamble because [REDACTED]

[REDACTED]

[REDACTED],

AMD’s pre and post-hearing briefing clearly identifies [REDACTED], as they are comprised in the ordinary course of business, as infringing apparatuses. CPB at 7 (citing Appendix A); CIB at 7 (citing EDIS Doc. ID 772055). And further, Respondents’ failure to contest infringement of claim 15’s preamble by any Accused Product in their own pre-hearing brief, or even arguing it is limiting, demonstrates a lack of prejudice. RPB at 128; CRB

at 54 (citing RIB at 112); *Am. Med. Sys. V. Biolitec, Inc.*, 618 F.3d 1354, 1358 (Fed. Cir. 2010) (holding that, generally, preambles are not limiting) (citations omitted). Thus, AMD has met its burden for limitation [15a] on all accused products. To be clear, however, infringement theories which depend on [REDACTED] as the “memory” can only succeed as indirect infringement against those Accused Products which [REDACTED]. CIB at 86 ([REDACTED]); CRB at 43.

- b. [15d] “a plurality of engines configured to receive, from a scheduling module associated with a second processing unit, a plurality of tasks and to load state data associated with each of the plurality of tasks”

Limitation [15d] reads, “a plurality of engines configured to receive, from a scheduling module associated with a second processing unit, a plurality of tasks and to load state data associated with each of the plurality of tasks.” 381 patent at cl. 15. Again, AMD contends the limitation is met in two ways. CIB at 88. AMD’s first theory is that [REDACTED] are the “plurality of engines.” *Id.* AMD’s second theory is that [REDACTED] are the “plurality of engines.” *Id.* As for the other elements, AMD contends: [REDACTED] is the claimed “scheduling module,” [REDACTED] is the claimed “second processing unit,” [REDACTED] are the claimed “tasks,” and [REDACTED] are the claimed “state data.” *Id.* (citing Hr’g Tr. (Reinman) at 133:4-21).

AMD, through Dr. Reinman, explains more specifically how [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

....

[REDACTED]

[REDACTED]

....

[REDACTED]

[REDACTED]

....

[REDACTED]

CIB at 88-92.

AMD also contends [REDACTED] in accordance with its proposed construction of “engine,” and “processes data” in accordance with Respondents’ proposed construction. *See* CIB at 93-95. But as determined above, neither is a proper construction for the term and thus need not be considered for claim 15.

[REDACTED]

Based on the evidence it cites, as excerpted above, AMD has shown it is more likely than not the limitation is met within the Accused Products. *See* CIB at 88-92. Respondents’ multiple arguments to the contrary are not persuasive.

First, Respondents contend [REDACTED] do not “load state data” because [REDACTED]

[REDACTED] RIB at 78 (citing Hr’g Tr. (Reinman) at 399:18-400:2); RRB at 43-44. However, Respondents’ own expert, Mr. de la Iglesia, opined that [REDACTED]

[REDACTED] Hr’g Tr. (de la Iglesia) at 864:17-25; *see* Hr’g Tr. (Reinman) at 143:24-144:19. Moreover, the plain and ordinary reading of “load state data” is broad. *See* Order No. 31 at 41. It seemingly encompasses the issuance of a command by [REDACTED]

[REDACTED]. There is nothing in the 381 patent specification, and Respondents cite to none, to show this is a distinction with a difference. *See* Order No. 31 at 41; RIB at 78; RRB at 43-44. Indeed, in just one example, the patent uses “provides” to cover direct as well as indirect connections between elements: “[t]o provide the priority type to GPU 110, application 302 provides bits to an API indicating the priority type of each task. The API, in turn, provides this information to the driver of GPU 110.” 381 patent at 9:54-57.

Respondents next challenge whether the data elements AMD identifies as “state data” are such. RIB at 79-80; RRB at 44-45. Essentially, Respondents contend that “state data,” as the 381 patent uses the term, “only resides in the thread context” and is not “data received from the CPU.” RIB at 79. Respondents acknowledge this “may include texture handles, shader constants, transform matrices, program counters, and other data that is stored locally in the GPU context and used by the processing engines to execute the thread,” but caution that it “only exists when a thread

[REDACTED]

context is created and while a thread is being executed.” *Id.* (citing 381 patent at 1:53-57).

Respondents reason:

The ‘381 patent describes how the patent was attempting to solve the prior art stall problem in which one thread context (its “state data”) would need to be switched out for another thread context in order for the processing engine to continue execution. JX-0003.0009 at 2:20-46. “After the more-important (*e.g.*, low-latency) task finishes executing, the state data associated with the threads of the more-important (*e.g.*, low-latency) task are flushed from the data-storage units, and the state data from the threads of the first task are swapped back into the data-storage units of the shader core. Then, the shader core can resume executing the threads of the first task.” *Id.* (emphasis added). Thus, the ‘381 patent is describing how “state data” is the context for a thread and the only reason that “state data” is associated with a “task” is because a “task” is eventually broken down into multiple threads. This is intuitive because the ‘381 patent and its claims are not related to anything that happens prior to thread execution, because prior to the execution of any threads, there is no stall problem yet and there is no “state data.” *Id.*

Id. Respondents add, “[n]othing that the ‘381 patent describes as ‘state data’ that is stored in a thread context in local GPU memory is even remotely similar to [REDACTED]

[REDACTED]” *Id.* at 80 (no citation). Rather, according to Respondents, [REDACTED]

[REDACTED]

[REDACTED] RIB at 80 (citing Hr’g Tr. (Pfister) at 709:22-710:18; RX-0001C ([REDACTED]) at 56:2-6:17).

Respondents intimate in their reply brief that state data only exists within [REDACTED]

[REDACTED]

[REDACTED]. See RRB at 44-45 ([REDACTED])

[REDACTED]).

There are several problems with these views. First is that Respondents provide no citation to evidence—in particular, no expert testimony—to support their claim that nothing in the 381 patent resembles [REDACTED]. To the contrary, the patent would seem to assign broad meaning to “state data” and gives examples of “texture handles,” as

well as “shader constants, transform matrices, *or the like*.” 381 patent at 1:53-56 (emphasis added).

And Dr. Reinman offered undisputed testimony that [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] Hr’g Tr. (Reinman) at 153:2-24.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

JX-0192 ([REDACTED]) at 129:23-131:23. Dr. Reinman further explained how [REDACTED]

[REDACTED] *See id.* at 165:13-167:21.

Altogether, there is more evidence than not that [REDACTED] load “state data.”

Second, even assuming that all of the data elements Dr. Reinman identifies as “state data” are [REDACTED]

[REDACTED], the intrinsic evidence of the 381 patent shows that it makes no difference. Admittedly, the patent uses the exact term “pointers” rarely; only twice and in the context of a sub-component located between the engines and arbiter which operates exclusive on vertex shader tasks. *See* 381 patent at 7:42-54 (discussing vertex analyzer 208), Fig. 2. But other disclosures strongly imply pointers are in use even when not explicitly recited. For example, claim 15 recites “engines” which “receive” “tasks” from a “scheduling module.” 381 patent at cl. 15. And in at least one embodiment, the specification describes this step as the engines “retriev[ing]” their tasks for later issuance to a shader core. 381 patent at 5:3-13. It is understood that retrieving data from a memory is done through pointers (Hr’g Tr. (Reinman) at 238:1-7; 382:12-18), indicating that what was “received” by the engine from the scheduling module was likely a pointer (*see* 381 patent at 8:59-9:9 (disclosing that any time the engines of command processor 230 need data, a request is made to the memory system)). And if receiving a pointer to data is treated the same as receiving the data, it would follow that loading a pointer is treated the same as loading the pointer’s data.

Respondents further challenge the propriety of AMD’s second “engines” theory, where the “engines” are [REDACTED]. RIB at 82-84. Respondents emphasize that according to the claim language, the engines must be outside of the shader core in order for that shader core to “receive” tasks from those engines. *Id.* at 82. Yet, they argue, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]. *Id.*
at 83-84 (citing CDX-0005C.20; Hr’g Tr. (Reinman) at 396:1-18, 402:4-10; Hr’g Tr. (Pfister) at
706:23-707:20); RRB at 46. Respondents summarize, [REDACTED]
[REDACTED] RIB at 84.

There is some merit to Respondents’ position, but it does not prevent a finding of
infringement. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

CX-0773C.30. [REDACTED]

[REDACTED] *See generally Plastic Omnium Advanced Innovation and Research v. Donghee Am., Inc.*, 943 F.3d 929, 935-6 (Fed. Cir. 2019) (emphasizing infringement analysis is between claim terms and accused product, not feature names used in defendant’s product literature). AMD’s theory is that the claimed “shader core” is [REDACTED]

[REDACTED]

[REDACTED]. CIB at 95. While

Contains Confidential Business Information Subject to Protective Order

AMD's pre-hearing brief is less clear, in that it seems to identify [REDACTED]
[REDACTED] as the claimed
"shader core," it does not indicate that [REDACTED]

[REDACTED]

CPB at 135-136.

Indeed, Respondents' argument seems to be based on the particular phrasing of their
question to Dr. Reinman, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Hr'g Tr. (Reinman) at 402:3-10; RIB at 83 (citing same). That Respondents switched "shader
core" for "execution engines" is not reason to alter AMD's infringement theory. See RRB at 47
(inaccurately describing CPB at 135). So there is no impermissible overlap in considering [REDACTED]
[REDACTED] the accused "engines."

Lastly, Respondents assert the limitation is not met because: [REDACTED]
[REDACTED]
[REDACTED]. None of these have merit as
they are based on claim constructions rejected above and so need not be addressed further.

Accordingly, limitation [15d] has been shown in the Accused Products.

- c. [15e] “a shader core configured to receive the plurality of tasks from at least one of the plurality of engines and to execute a first task from the plurality of tasks while executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks”

Limitation [15e] reads, “a shader core configured to receive the plurality of tasks from at least one of the plurality of engines and to execute a first task from the plurality of tasks while executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks.” 381 patent at cl. 15. AMD contends the limitation is met via

[REDACTED]
[REDACTED]. CIB at 95, 97. AMD explains [REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

CIB at 97-97.

As for “execute a first task from the plurality of tasks while executing a second task,” AMD points to testimony from [REDACTED]

[REDACTED]

[REDACTED] *See id.* at 98-99 (citing JX-0192C ([REDACTED]) at 114:7-18, 128:14-129:14). AMD also references the testimony of Dr. Reinman and Respondents’ expert Mr. de la Iglesia, who explained [REDACTED]

[REDACTED]. *Id.* at 102 (citing Hr’g Tr. (Reinman) at 406:17-407:13; Hr’g Tr. (de la Iglesia) at 866:18-868:15). Finally, as for “execute . . . based on respective state data associated with each of the first and second tasks,” AMD explains:

[REDACTED]

Id. at 103. Based on this cited evidence, AMD has shown it is more likely than not the limitation is met within the Accused Products. Respondents’ multiple arguments to the contrary are not persuasive.

[REDACTED]

Respondents first contend that the claim language “configured to” execute a first task while executing a second task means the products must not only have the capability to do both actions at the same time but must be “specifically ‘arranged or set-up to perform a specified function,’ ‘set up for operation especially in a particular way,[’] or ‘actually programmed or equipped with software to’ perform the recited function.” RIB at 76. Respondents view Dr. Reinman as only testifying [REDACTED]. *Id.* at 76-77 (citing Hr’g Tr. (Reinman) at 163:3-11, 405:24-406:5), 90-91 (citing Hr’g Tr. (Reinman) at 405:24-406:16); *see* RRB at 51-52.

The overall weight of expert testimony weighs against Respondents’ position. Dr. Reinman did use the word [REDACTED] to describe his evaluation of the Accused Products (Hr’g Tr. (Reinman) at 163:3-11) and answered in the affirmative that [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] (*id.* at 405:24-406:1). But he also testified that

[REDACTED]


[REDACTED] Hr’g Tr. (Reinman)

at 163:12-22.

[REDACTED]

[REDACTED]

[REDACTED]



CX-0773C.30. The evidence shows that [REDACTED]
[REDACTED]
[REDACTED]. See, e.g., Hr’g Tr. (de la Iglesia) at 865:20-867:20 ([REDACTED]
[REDACTED]); Hr’g Tr. (Reinman) at 406:17-407:23; JX-0192C ([REDACTED]) at 128:14-23; CX-1724C.9-10.
Assuming for present purposes (but confirmed below) that [REDACTED] qualifies as
“execute a first task . . . while executing a second task,” then [REDACTED]
[REDACTED]
will satisfy the “configured to” claim limitation. Again, this [REDACTED] is
confirmed by Dr. Reinman, Mr. de la Iglesia, Dr. Pfister, and [REDACTED]:
[REDACTED]
[REDACTED]

[REDACTED]

Hr'g Tr. (Reinman) at 162:16-23 (emphasis added);

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Hr'g Tr. (de la Iglesia) at 824:15-826:5 (emphasis added);

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Hr'g Tr. (Pfister) at 800:11-801:16;

[REDACTED]

CX-1730C.39; *see* CX-1730C.10 ([REDACTED]).

[REDACTED]

[REDACTED]

[REDACTED]. *Id.* at 11. Thus, the “configured to” limitation is met.

Next, Respondents take issue with the fact that [REDACTED]

[REDACTED]

[REDACTED]. *See* RIB 85-87, 88-90. They contend the

claim requires that “what is received from the scheduling module must be what the shader core executes” (*id.* at 88) and “[t]he claim does not state that merely a portion of a ‘first task’ must execute while a portion of the “second task” is executing” (*id.* at 90). Put another way by

Contains Confidential Business Information Subject to Protective Order

Respondents, [REDACTED]

[REDACTED] *Id.* at 90; *see* RRB at 48-50.

But executing an entire job or task is unrealistic given the ordinary understanding that GPUs break down high-level user application requests into small pieces for faster parallel processing—and that those small pieces often take the form of a *sequence* of instructions known as threads (and discussed at length above in connection with the 053 patent). This basic understanding is disclosed in the 381 patent itself. 381 patent at 1:33-2:7; *see generally* Hr’g Tr. (Pfister) at 569:23-570:7, 573:17-23 (stating a program counter keeping track of which instruction in the sequence is necessary and a GPU cannot function without it). Executing a *portion* of a first task while executing a *portion* of a second task thus falls comfortably within the scope of the claim. Order No. 31 at 38 (finding a “task” is, broadly, “a unit of processing work or portion of an application to be executed.”). Respondents’ argument is therefore rejected.

Moving on, Respondents assert that there is a lack of evidence that any Accused Product actually includes [REDACTED]

[REDACTED]

RIB at 89 n.9. This is not persuasive. [REDACTED]

[REDACTED]

[REDACTED]. CIB at 95; CRB at 47. So it is irrelevant

that [REDACTED]

Contains Confidential Business Information Subject to Protective Order

[REDACTED]
[REDACTED] JX-0192C ([REDACTED]) at 33:9-23; *see*
CIB at 7; CDX-0005C.15; RDX-0005C.5; CX-0773C.30.

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]² Relatedly, Respondents argue [REDACTED]

[REDACTED] do not meet the construction of “while.” RIB at 87-88. This is based on an overly narrow claim construction rejected above, however.

Lastly, Respondents argue that [REDACTED]
[REDACTED]
[REDACTED] RIB at 91 (emphasis in original).

They explain:

[REDACTED]

² This settles the issue described in the parties’ joint motion (1318-071) for leave to file an additional post-hearing statement.

[REDACTED]

Id. at 91-92 *see* RRB at 52-53. AMD claims this is a new argument appearing for the first time, with an implication that it should be struck under Ground Rule 9.2. CRB at 52. But there was fair notice here. RPB at 145.

Nevertheless, it is not persuasive. There are two possible interpretations for “based on” within “execute a first task from the plurality of tasks while executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks.” 381 patent at cl. 15. The first is Respondents’, where the “while” condition is “based on respective state data.” The second, favored by AMD, is that the execution of the tasks simply needs to be “based on respective state data.” CRB at 52. The intrinsic evidence supports the second. The patent describes how, in conventional GPUs, state data are associated with particular tasks, must be loaded for the task to execute, and to switch tasks, must be “flushed from the data-storage units, and the state data from the threads of the first task are swapped back into the data-storage units of the shader core.” 381 patent at 2:34-36. According to the patent, this is known as “context switching” and comes at a cost. *Id.* Thus, the purpose of the patent is to simply have state data for more than one task pre-loaded so that more than one task can be executed “without requiring a context switch.” *See, e.g.*, 381 patent at 8:27-32. This reflects AMD’s interpretation. Meanwhile, there is almost no discussion in the patent as to how state data would cause two tasks to be executed at the same time, or benefit from that circumstance, and Respondents do not point to any such disclosures. Thus, their interpretation and non-infringement argument are not accepted.

Accordingly, limitation [15e] has been shown in the Accused Products, and thus, infringement of claim 15 has been established.

2. Claim 16

For reference, claim 16 of the 381 patent requires:

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

[REDACTED]

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

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.

381 patent at cl. 16. AMD contends [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] CIB at 104 (citing, *inter alia*, Hr’g Tr. (Reinman) at 169:3-173:1; CX-1724C.7). Respondents oppose on various grounds, including that AMD has failed to show [REDACTED]

[REDACTED]

[REDACTED] RIB at 92.

AMD has met its burden. [REDACTED] constitute the claimed “plurality of processing elements,” and [REDACTED] there would be at least two pluralities. There is no dispute that [REDACTED]. *See, e.g.*, JX-0192 ([REDACTED]) at 56:17-21, 59:4-17; 381 patent at 1:47-57. And as determined above, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] CX-01730C.8, 10. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] 381 patent at 5:14-28, 8:8-17.

Accordingly, infringement of claim 16 has been shown.

[REDACTED]

3. Claim 17

For reference, claim 17 of the 381 patent requires:

17. The computing system of claim 15, 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.

381 patent at cl. 17. AMD contends that [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] CIB at 104 (citing Hr’g Tr. (Reinman) at 173:12-174:22; JX-0192C ([REDACTED]) at 1325:25-133:8). Respondents do not dispute infringement of claim 17 apart from its dependency on claim 15. RIB at 93; RRB at 53. Respondents’ reference to the impact of claim 17’s language on claim 15 is not understood to be an assertion that the additional limitations of claim 17 are not present in the Accused Products. RIB at 93. In light of the cited evidence and opinion, claim 17 is infringed by the Accused Products.

4. Claim 19

For reference, claim 19 of the 381 patent requires:

19. The computing system of claim 15, wherein the first task comprises a graphics-processing task and the second task comprises a general-compute task.

381 patent at cl. 19. AMD contends [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] CIB at 105 (citing Hr’g Tr. (Reinman) at 174:24-176:9; CX-0773C.26-27). Supporting AMD’s theory is [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

CX-0773C.27. AMD has met its burden. Respondents oppose, but their argument is based on the mistaken construction for “while executing” rejected above. *See* RIB at 93.

Accordingly, AMD has shown claim 19 is infringed by the Accused Products.

5. Claim 20

For reference, claim 20 of the 381 patent requires:

20. The computing system of claim 15, wherein the shader core is configured to receive the plurality of tasks based on a priority associated with each of the plurality of tasks.

381 patent at cl. 20. AMD contends [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

CIB at 105-106. AMD has met its burden to show, as testified to by Dr. Reinman, [REDACTED]

[REDACTED]

[REDACTED]. Respondents dispute that [REDACTED]

[REDACTED] but cite to no evidence in support. RIB at 93.

Accordingly, AMD has shown claim 20 is infringed by the Accused Products.

6. Indirect Infringement

AMD states that it pursues an indirect infringement theory under the 381 patent but gives no explanation as to what claim limitation might be missing in any Accused Product so as to necessitate such a theory. *See* CIB at 106; CRB at 52. For example, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] Instead, AMD tersely incorporates by reference its argument and evidence presented in connection with the 053 patent. CIB at 106; CRB at 52. This is insufficient, even setting aside AMD's failure under the 053 patent to show direct infringement by the end user. Accordingly, no indirect infringement of the 381 patent has been shown.

E. Domestic Industry – Technical Prong

According to AMD's post-hearing briefing, the use, manufacture, or sale of DI Products practice claims 15-17, 19, and 20 of the 381 patent. *See* CIB at 14. Of these, claim 15 is independent, and claims 16, 17, 19, and 20 depend from claim 15. For the reasons discussed below, AMD has shown practice of these claims by the DI Products.

[REDACTED]

1. Claim 15

Claim 15 of the 381 patent is presented above in connection with infringement. The primary disputes between the parties concern the “scheduling module associated with a second processing unit,” “plurality of engines,” and “shader core” in limitations [15d] and [15e]. These are discussed below. For those remaining limitations which are not in dispute ([15a], [15b], and [15c]), the DI Products have been shown to practice them as alleged in view of the testimony of Dr. Reinman. *See* CIB at 107-108 (citing Hr’g Tr. (Reinman) at 184:15-189:8).

a. [15d] “a scheduling module associated with a second processing unit”

Limitation [15d] reads, “wherein the first processing unit comprises: a plurality of engines configured to receive, from a scheduling module associated with a second processing unit, a plurality of tasks and to load state data associated with each of the plurality of tasks.” 381 patent at cl. 15. AMD contends the limitation is met through two components of [REDACTED]

[REDACTED]
which, together, AMD identifies as the “plurality of engines.” CIB at 108 (citing Hr’g Tr. (Reinman) at 189:10-194:25, 208:1-210:2, 212:10-215:21). According to AMD, [REDACTED]

[REDACTED] . *Id.* (citing, *inter alia*, CX-0080C.5-26, CX-0129C.5-56). Moreover, it alleges [REDACTED] communicate with [REDACTED]

[REDACTED]
[REDACTED], and load state data associated with those tasks:

[REDACTED]

[REDACTED]

Id. at 109-110.

Respondents oppose on two grounds. First, Respondents advance several interrelated arguments surrounding the limitation “from a scheduling module associated with a second processing unit”; namely, that it is limiting (RIB at 94-96; RRB at 55-57 (citing, *inter alia*, *INVT SPE LLC v. Int’l Trade Comm’n*, 46 F.4th 1361 (Fed. Cir. 2022))), that the “scheduling module” and “second processing unit” must be within the overall “computing system” as found in the preamble (RIB at 96-100; RRB at 55), and that AMD has failed to provide “*any* evidence at the hearing that *any* particular driver code was used either in or with any of the AMD DI Products” (RIB at 103 (emphasis in original); RRB at 57-61 (citing, *inter alia*, Hr’g Tr. (Reinman) at 191:12-

[REDACTED]

223:12)). On the second point, Respondents view AMD as having waived any position on whether the “scheduling module” and “second processing unit” are within the “computing system” or not—as well as any argument that AMD supplies a “scheduling module” to customers of the DI Products. *See* RRB at 54 (citing Order No. 4 at G.R. 13.1).

On the last point, Respondents emphasize that Order No. 54 struck AMD’s ability to rely on certain [REDACTED], leaving them only with a version

. RIB at 103-104. Consequently, Respondents argue “AMD has no evidence in the record of any [REDACTED] that

” and “there is a complete failure of proof that this claim limitation is met by the AMD Products.” *Id.* at 105; RRB at 58 (“no evidence was introduced at the hearing—and none exists—that [REDACTED]”).

Even if it is shown that the DI Products receive tasks from a scheduling module associated with a second processing unit, Respondents argue it has not been shown the DI Products are “configured” for that reception. RRB a 60 (“would receive the tasks regardless of what they were received from”).

In its reply brief, AMD acknowledges the first point, that “from a scheduling module associated with a second processing unit” is limiting, but disagrees that the two elements are part of the “computing system.” CRB at 53. AMD contends, “Federal Circuit precedent is clear that when a first component is configured to interact with a second component, the claim is met even in the absence of the second component.” *Id.* at 53-54 (citing *Revolution Eyewear, Inc. v. Aspex Eyewear, Inc.*, 563 F.3d 1358, 1370 (Fed. Cir. 2009); *Certain Vision-Based Driver Assistance System Cameras, Components Thereof*, Inv. No. 337-TA-907, Comm’n Op. at 12, 13 (Dec. 1,

[REDACTED]

2015)). Moreover, AMD argues there is no dispute between the parties that the preamble “computing system” is *not* limiting. *Id.* at 54 (citing RIB at 112, 124). And so, according to AMD,

[REDACTED]

[REDACTED] *Id.* at 55. AMD views the prosecution history excerpts cited by Respondents (JX-0009.175-187) as confirming that the “scheduling module” and “second processing unit” are “component[s] that [are] outside of the claimed system, whereas the ‘scheduling module’ claimed by [the prior art] was a component within the shader core, and thus would be part of the first processing unit and part of the claimed system.” *Id.* AMD concludes:

Respondents’ experts presented absolutely no testimony regarding this limitation of Claim 15 as it relates to the DI Products. Therefore, Respondents’ allegation that “it cannot be determined whether a product satisfies this limitation and has the required configuration without reference to a specific scheduling module and its associated second processing unit,” amounts to nothing more than unsubstantiated attorney argument.

Id. at 57.

Respondents are not persuasive. AMD acknowledges that “from a scheduling module associated with a second processing unit” is limiting (CRB at 53), but Respondents acknowledge that the preamble, “a computer system comprising” is not limiting (RIB at 112). Thus, whether the “scheduling module” and “second processing unit” are inside or outside the “computing system” need not be resolved.

To Respondents, the only alternative is to view the “scheduling module” and “second processing unit” as part of the “plurality of engines.” RIB at 100. This is far from correct. As a matter of grammar, the “scheduling module” and “second processing unit” would only be part of the “first processing unit”—and not necessarily the subcomponent “plurality of engines”:

15. A computing system comprising:
a memory;

[REDACTED]

a first processing unit; and

a bus coupled to the memory and the first processing unit,

wherein *the first processing unit comprises:*

a plurality of engines configured to receive, from *a scheduling module associated with a second processing unit*, a plurality of tasks and to load state data associated with each of the plurality of tasks; and . . .

381 patent at cl. 15 (emphasis added). Even then, the specification never places the “scheduling module” and “second processing unit” (*i.e.*, CPU) as within the GPU’s engines, or within the first processing unit. *See generally* 381 patent; RIB at 100-103 (citing no intrinsic evidence of such an embodiment). So Respondents’ alternative reading is rejected.

There is no reasonable confusion here that would cause indefiniteness either, as Respondents suggest. *See id.* at 101-102 n.12. A person of ordinary skill in the art would understand the “scheduling module” and “second processing unit” are outside of the GPU, as every relevant 381 patent figure shows and the prosecution history records—a point Respondents effectively concede. *See, e.g.*, 381 patent at Figs. 1, 4; JX-0009.175-187; RIB at 96-100 (“The specification again discloses that in this embodiment, ‘there are two different types of processing units, CPU 102 and GPU 110,’ and the ‘OS scheduler’ is what provides tasks to the GPU.”).

AMD has further shown it is more likely than not that the “plurality of engines” within the DI Products are “configured to” “receive . . . a plurality of tasks” from “a scheduling module associated with a second processing unit.” Nowhere within Respondents’ briefing, or their expert’s testimony, is it disputed that the DI Product GPUs are intended to and do operate with a CPU (claimed “second processing unit”). Nor is there any assertion that they can operate *without* a CPU under any circumstances. And while AMD failed to enter into the record appropriate [REDACTED]

[REDACTED] definitively proving this relationship, other reliable evidence shows that *some*

[REDACTED] is used, including testimony from Dr. Reinman, AMD representatives, and

Contains Confidential Business Information Subject to Protective Order

AMD documentation. Hr'g Tr. (Reinman) at 193:1-194:25; Hr'g Tr. (McCrary) at 37:19-22, 39:1-40:18; JX-0131C (Ramsey) at 95:15-21, 97:10-13, 66:8-68:7, 86:12-87:6; CX-0079C.13. For example, CX-0129C, authored by witness and 381 patent inventor Mr. Rex McCrary, describes the interaction between [REDACTED]

[REDACTED]:

[REDACTED]

[REDACTED]

CX-0129C.5; *id.* at 51 (“[REDACTED]

[REDACTED]


[REDACTED]

[REDACTED]

[REDACTED]”). Another document shows the “[REDACTED]” is [REDACTED]
[REDACTED]. CX-2007.4. Respondents critique AMD’s evidence as individually not showing the limitation is met (RRB at 58-60), but as AMD observes, this is entirely attorney argument (*see id.* (without citation to Dr. Pfister)). Regardless, the question is whether a preponderance of all evidence shows that the “plurality of engines” within the DI Products are “configured to” “receive . . . a plurality of tasks” from “a scheduling module associated with a second processing unit.” The evidence cited by AMD comfortably answers this question in the affirmative. *See* CIB at 109-10.

b. [15d] “a plurality of engines . . . to load state data associated with each of the plurality of tasks”

Respondents’ second ground of opposition to limitation [15d] is that the “plurality of engines” identified as the [REDACTED] in the DI Products, do not “load state data associated with each of the plurality of tasks.” RIB at 107. Rather, according to Respondents and Dr. Pfister, a [REDACTED] called [REDACTED] loads the state data, and it receives that state data from [REDACTED]
[REDACTED], neither of which are the [REDACTED] alleged “engines.” *See id.* at 107-109 (citing Hr’g Tr. (Pfister) at 746:7-747:8; RX-0435; RX-0468). Respondents cite a diagram showing this arrangement:



RX-0468.8; *see* RIB at 108-9; RRB at 61. Respondents acknowledge that *some* [REDACTED] passes from the [REDACTED], but dismiss it as “[REDACTED],” which are distinct from vertex/pixel shaders and the type of “state data” relied on by Dr. Reinman. *See* RIB at 109, 109 n.16 (citing JX-0131C (Ramsey) at 79:12-89:5).


As discussed above with respect to infringement, it has already been determined that “[t]he claim does not require that the engines load the state data into their own memories (as opposed to some other component or components) . . . and it does not specify how or from which component

Contains Confidential Business Information Subject to Protective Order

the shader core receives the state data.” Order No. 31 and 41. And so, the CPC and CPG can satisfy the claim if they cause the state data to be loaded *somewhere*.

The record evidence shows this is the case. The [REDACTED] document, relied on by AMD and Respondents, describes the “[REDACTED]” as “[REDACTED]”

RX-0435C.28. Dr. Reinman testified that [REDACTED] flows to the [REDACTED] through the [REDACTED], whether it be from [REDACTED] (together making the [REDACTED]), and cited source code in support. Hr’g Tr. (Reinman) at 219:11-220:7, 221:3-222:4. AMD technical documentation shows the [REDACTED] connecting all of the [REDACTED] [REDACTED], which AMD witness Mr. Randy Ramsey testified occurs [REDACTED] [REDACTED]:



[REDACTED]

CX-0129C.5; JX-0131C (Ramsey) at 65:14-66:7 ([REDACTED])
[REDACTED], 77:15-78:12 ([REDACTED]), 82:2-83:5
([REDACTED]); *see* CX-0129C.6-7 (listing CPG
Interfaces with SPI, VGT, GRBM, etc.); CX-1029C.5 (AMD documentation explaining the [REDACTED]
[REDACTED]).

Seen in this context, Respondents' chosen figure from RX-0468C.8 is inadequate as it
omits all blocks to the [REDACTED] that might have a role in [REDACTED]. RIB at 108-109.
And Dr. Pfister not only cited no code but also seemed to base his strict reading on what RX-
0435C.28 "suggests." Hr'g Tr. (Pfister) at 746:4-747:8 (opinion based on what RX-0435C.28
"suggests"), 804:5-805:14 ("I don't recall the specifics."). His opinion is less persuasive for these
reasons, and AMD has met its burden.

Respondents also argue that not all DI Products necessarily support [REDACTED], such
that [REDACTED]
[REDACTED] *See* RIB at 109 n.16 (citing JX-0131C (Ramsey) at 79:12-89:5); RRB at 62
n.9 (same). This is entirely based on an off-hand comment from Mr. Ramsey, and when read in
context, does not cause as much doubt as suggested:

[REDACTED]

See JX-0131C (Ramsey) at 80:10-23; *see also id.* at 79:12-80:9 (no relevant discussion), 80:24-
89:5 (same). Respondents also argue that [REDACTED] are exclusively [REDACTED] and thus

[REDACTED]

should be excluded from the scope of “engines” (RIB at 106-107), but this is based on a rejected claim construction as discussed above. And to the extent [REDACTED] are used in the DI Products, Respondents do not make any challenge based on their [REDACTED]. *See id.* at 94-110; RRB at 53-62.

Accordingly, this portion, and all of limitation [15d], has been shown in the DI Products.

- c. **[15e] “a shader core configured to receive the plurality of tasks from at least one of the plurality of engines and to execute a first task from the plurality of tasks while executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks”**

Limitation [15e] reads, “a shader core configured to receive the plurality of tasks from at least one of the plurality of engines and to execute a first task from the plurality of tasks while executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks.” 381 patent at cl. 15. AMD contends the limitation is met, in part, through arrays of [REDACTED] wherein each array is a “shader core,” and that each [REDACTED] further includes [REDACTED] CIB at 112-113 (citations omitted). AMD also contends “the array of [REDACTED] is coupled to the [REDACTED] such that each [REDACTED] can receive a plurality of tasks from any of the [REDACTED] via the [REDACTED], and the array of [REDACTED] is configured to execute the first and second tasks.” *Id.* at 113 (citations omitted). This much is undisputed.

Respondents only take issue that the “shader core” is not “configured to” execute two tasks together based upon respective state data for each task. RIB at 109-110. Respondents reiterate their position that executing a first task while executing a second task requires more than

[REDACTED]. *Id.* at 110. Respondents argue that merely

[REDACTED] only shows “capability.” RRB at 62; Hr’g Tr. (Pfister) at 747:9-24.

Respondents are only modestly persuasive. As determined above, some overlap in time is sufficient under the claim so long as the article is “configured” to have that overlap.

[REDACTED]

And the evidence shows the DI Products are so configured. Specifically, tasks are received by the [REDACTED] which [REDACTED]. Hr’g Tr. (Reinman) at 217:3-218:19. Dr. Reinman testified that, from his source code review, the products are configured to [REDACTED]
[REDACTED]
[REDACTED]” *Id.* at 218:20-219:6. Mr. Ramsey also described a [REDACTED]
[REDACTED]. JX-0131C at 49:16-52:3. Dr. Pfister does not dispute these operations other than to generally argue an “ability” does not mean “configured” (*see* Hr’g Tr. (Pfister) at 747:9-24), but there is a configuration here to send [REDACTED] to [REDACTED]. On balance, therefore, AMD has the stronger position.

Accordingly, this portion, and all of limitation [15e], has been shown in the DI Products, and the DI Products are found to practice claim 15.

2. Claim 16

Claim 16 of the 381 patent is presented above in connection with infringement. AMD contends, as testified to by Dr. Reinman, that “each AMD GPU includes an [REDACTED]
[REDACTED]” and “[e]ach [REDACTED]
[REDACTED]
[REDACTED].” CIB at 116; Hr’g Tr. (Reinman) at 223:14-225:6. Respondents do not dispute practice of claim 16 apart from its dependency on claim 15. RIB at 94-110; RRB at 53-62. In light of the cited evidence and opinion, the additional elements of claim 16 are met in the DI Products.

3. Claim 17

Claim 17 of the 381 patent is presented above in connection with infringement. AMD contends, as testified to by Dr. Reinman, that “the [REDACTED]

[REDACTED]

receive tasks for the [REDACTED]

[REDACTED].” CIB at 117; Hr’g Tr. (Reinman) at 225:8-226:5. Respondents do not dispute practice of claim 17 apart from its dependency on claim 15. RIB at 940-110; RRB at 53-62. In light of the cited evidence and opinion, the additional elements of claim 17 are met in the DI Products.

4. Claim 19

Claim 19 of the 381 patent is presented above in connection with infringement. AMD contends, as testified to by Dr. Reinman, that “the graphics-processing task would include graphics applications which would be received and handled by [REDACTED].” CIB at 118; Hr’g Tr. (Reinman) at 226:6-227:6. Respondents do not dispute practice of claim 19 apart from its dependency on claim 15. RIB at 94-110; RRB at 53-62. In light of the cited evidence and opinion, the additional elements of claim 19 are met in the DI Products.

5. Claim 20

Claim 20 of the 381 patent is presented above in connection with infringement. AMD contends, as testified to by Dr. Reinman, that, among other things, “[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] CIB at 118-119; Hr’g Tr. (Reinman) at 227:7-228:15. Respondents do not dispute practice of claim practice of claim 20 apart from its dependency on claim 15. RIB at 94-110; RRB at 53-62. In light of the cited evidence and opinion, the additional elements of claim 20 are met in the DI Products.

F. Validity

Respondents identify the following invalidity theories for the 381 patent:

Claims	Theory
15-17, 20	Anticipated under 35 U.S.C. § 102 by Jiao (RX-0423)
15-17, 20	Anticipated under 35 U.S.C. § 102 by Stuttard (RX-0416)
15-17, 19, 20	Rendered obvious under 35 U.S.C. § 103 by Jiao (RX-0423)
15-17, 19, 20	Rendered obvious under 35 U.S.C. § 103 by Stuttard (RX-0416)
15-17, 19, 20	Invalid under 35 U.S.C. § 101 for lack of patentable subject matter

See generally RIB at 110-144. Pub. No. WO 00/62182 (“Stuttard”) published on October 19, 2000. RX-0416. U.S. Patent No. 8,319,774 (“Jiao” or “Jiao ’774”) issued on November 27, 2012, was published on March 22, 2012, and is a continuation of an application filed on April 4, 2008. RX-0423. Respondents make no effort to show that Stuttard is prior art to the 381 patent, and assert, erroneously, that Jiao is prior art under pre-AIA 35 U.S.C. § 102(b). *See* RIB at 110, 124. AMD nevertheless does not dispute their prior art status, and demonstrating that status is trivial: Stuttard published years before the priority date of the 381 patent, and is therefore prior art under Section 102(b); and Jiao claims priority to an application filed before the 2009 priority date of the 381 patent, and so is prior art under pre-AIA Section 102(e). 35 U.S.C. § 102(b), (e) (pre-AIA); *see generally* CIB at 119-139; CRB at 59-73. Thus, Stuttard and Jiao are prior art to the 381 patent.

1. Jiao

As noted, Respondents contend claims 15-17, 19, and 20 are, collectively, anticipated or rendered obvious by prior art reference Jiao, which describes methods for managing the buffer of

[REDACTED]

constant data values used by shader programs. RIB at 110-124, 133-135; RX-0423 at 2:24-53. As discussed below, Respondents have shown invalidity of claims 15-17 under 35 U.S.C. § 102 as anticipated by Jiao. Respondents have not shown invalidity of claims 19 or 20 under 35 U.S.C. § 102 as anticipated by Jiao or under 35 U.S.C. § 103 in view of Jiao.

a. Claim 15

With respect to claim 15, Respondents contend Jiao anticipates and renders obvious all of its limitations. *See* RIB at 110, 112-122. Indeed, limitations [15d] and [15e] are the only challenged limitations in AMD’s post-hearing briefing. *See* CIB at 120-127; CRB at 59-65. In light of the evidence and expert testimony provided by Respondents’ expert, Dr. Pedram, limitations [15a]-[15c] are taught in Jiao as alleged. *See* RIB at 113-114 (citing Hr’g Tr. (Pedram) at 931:20-932:24).

Limitation [15d] states “wherein the first processing unit comprises: a plurality of engines configured to receive, from a scheduling module associated with a second processing unit, a plurality of tasks and to load state data associated with each of the plurality of tasks.” 381 patent at cl. 15. Respondents contend the “plurality of engines” is met in two ways: first, by Jiao’s vertex shader (VS) 304, geometry shader (GS) 306, and pixel shader (PS) 310 that lie within a “computational core 204”; and second, assuming AMD’s infringement theory regarding “frontend units” is accepted, by the same three shaders “plus the ‘thread slots’ within the execution units 420 of the EU pool 414.” RIB at 114-115 (citing RX-0423 at 5:47-51, 6:1-64, 9:11-24; Hr’g Tr. (Pedram) at 934:1-8, 949:8-11)). The following figure from Jiao is useful to see the arrangement of GPU blocks:




FIG. 4

RX-0423 at Fig. 4. Respondents continue:

Jiao discloses that the plurality of vertex, geometry, and pixel shaders (“engines”) are configured to receive a plurality of tasks from the EU pool control unit 206. Tr. (Pedram) at 935:14-936:16.; RDX-0004C.97-99; *see* RX-0423.00003, 6, 14 at 5:17-23 and 6:65-7:27 and Fig. 1 and 4. The shaders load constants, rendering context information, and graphics data before commencing execution of any tasks. *See, e.g.*, RX-0423.00003, 13-15 at 6:13- 44; 65-7:27, 8:35-41, 9:15-24, 9:39-54, Fig. 3 and 4; Tr. (Pedram) at 936:17-938:1; RDX-0004C.100-101.

....

Similar to how AMD is reading the claim for infringement, Jiao’s vertex shader 304, geometry shader 306, and pixel shader 310, receive graphics (high level) tasks from the driver software 128 residing in the system memory 124 via the EU pool control unit 206 and produce (low level) shader tasks that are then mapped to threads of the various execution units 420 by the EU pool control unit 206. *Id.* at 9:43-49, at 5:47-51. A POSITA would have understood that Jiao discloses a high-level task that has given rise to the low-level task that is mapped to a thread. *See* Tr. (Pedram) at 936:17-938:1; RDX-0004C.100-101.

[REDACTED]

Based on the ALJ's construction, Jiao discloses the claimed tasks. *See, e.g.*, RX-0423.00016 at 9:42-47 ("...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, and also performs an associated 'bookkeeping' of the tasks and threads."). *See, e.g.*, RX-0423.00015-16 at 8:65-67 and 9:10-13 (sometimes using task and thread interchangeably). Thus, a POSITA would have understood Jiao's inherent disclosure of a high-level task to be the construed "task." To the extent not considered an express disclosure, for the same reasons explained above, Jiao's disclosure would have also rendered this claim limitation obvious. *See* Tr. (Pedram) at 936:17-938:1; RDX-0004C.100-101.

Id. at 115.

As for "scheduling module," Respondents assert this is also met in two ways: first, by Jiao's "control unit 206"; and second, by Jiao's "driver software 128." RIB at 117-118. More specifically, Respondents argue:

Jiao's EU pool control unit 206 is the "scheduling module." Tr. (Pedram) at 934:9-21; RDX-0004C.93-94. The EU pool control unit 206 sends to the vertex, geometry, and pixel shaders ("engines") various tasks. For example, the EU pool control unit 206 "functions as a global scheduler for the tasks" (while also being responsible for assigning vertex shader, geometry shader, and pixel shader tasks to threads of the various execution units 420"). *See, e.g.*, RX-0423.00005-6, 14-16 at 6:65-7:27, 9:15-24, 9:39-54, Fig. 3 and 4. Further, Jiao's EU pool control unit 206 is associated with the command stream processor 406, which is the claimed "second processing unit." Tr. (Pedram) at 934:22-935:13; RDX.0004C.95-96; *see, e.g.*, RX-0423.00015 at 7:19-27 ("The command stream processor 406 provides triangle vertex indices to the EU pool control unit 206....The EU pool control unit 206 assembles vertex shader inputs from the stream cache and sends data to the computational core 204 (input E)...[and] assembles geometry shader inputs and provides those inputs to the computational core 204 (input F)... control[ling] the respective inflow and outflow to the computational core 204."). The command stream processor 406 is a processing unit associated with EU pool control unit 206 at least because it performs "new constant update[s]" for use by the EU pool control logic, which assembles inputs for shader tasks. RX-0423.00005-6, 14-16 at 3:48-60, 4:10-23, 10 55:11-46, 14:13-23; RDX-0004C.101.

Alternatively, the driver software 128 constitutes the claimed "scheduling module." Tr. (Pedram) at 935:15-936:4; RDX-0004C.93-94. Jiao discloses that the driver software "communicates instruction sets or commands through the use of the CPU 126 to registers in the GPU 114." RX-0423.00014 at 5:17-23. Responsive to a command from the driver software 128, vertex information is passed to the vertex shader 304 in the computational core 204 to implement vertex transformations, and after processing in the fixed-function graphics engine 208, are passed to the pixel

shader 310 in the computational core 204. RX-0423.00015 at 7:57-8:14. And, Jiao's driver software 128 is associated with the CPU 126, which is the claimed "second processing unit." *See, e.g.*, RX-0423.00014 at 5:8-13; Tr. (Pedram) at 934:24-935:13; RDX-0004C.95-96.

Id. at 117-118. Respondents reject a meaning for "associated with" as requiring the scheduling module to be a part of the second processing unit as opposed to a more general relationship, and further reject any prosecution history disclaimer occurred to limit its meaning. *See id.* at 118-120 (citing, *inter alia*, JX-0009.185).

As for the "state data" in "a plurality of engines configured to . . . load state data associated with each of the plurality of tasks," Respondents point to Jiao's "constant data pool" which are "context-specific constants stored in a constant buffer." RIB at 120 (citing RX-0423 at 1:50-58).

Respondents continue:

Jiao discloses that its shaders load constants, rendering context information, and graphics data before commencing execution of a task. Tr. (Pedram) at 936:17-938:17; RDX-0004C.100-101. Further, constants for performing one or more of the shader computations are provided by a constant buffer of a constant buffer system 200. RX-0423.00014-16 at RX-0423.0005-6, 16-18 at 6:1- 7:55, 8:65-9:14; 9:15-24, 12:33-52, 12:63-13:6, Figs. 3 and 4; Tr. (Pedram) at 936:17-938:17; RDX-0004C.100-101.

Id.

AMD argues limitation [15d] is not disclosed for at least nine reasons. *See* CIB at 120-125; CRB at 59-63. None are persuasive.

First, AMD contends, "Dr. Pedram, in rendering his opinion, failed to link the tasks that are received by the engines (*i.e.*, high-level tasks from an alleged scheduling module) and sub-tasks (*i.e.*, low-level tasks) that are generated to carry out the work associated with the high-level tasks issued by the scheduling module." CIB at 120. Put another way, AMD asserts, "[f]atal to Dr. Pedram's opinion is that he does not and cannot link any of these alleged 'lower level' sub-

tasks that he identifies as satisfying the claimed tasks to the higher level tasks that are actually issued by the alleged scheduling module.” *Id.* at 121; *see* CRB at 59-60.

There are two problems here. First, the premises behind the contention seem to come from Dr. Pedram’s deposition and not his hearing testimony, which is improper. *See* CIB at 120-121 (citing CX-0672); RRB at 64 n.10. Second, it is difficult to understand, as a general matter, how the tasks which are received by Jiao’s VS 304, GS 306, and PS 310, or the “sub-tasks” leaving those same units for the Execution Units 420, would not be “linked” or associated with what had been previously issued by Jiao’s EU Pool Control 206 (first possible “scheduling module”) or driver software 128 (second possible “scheduling module”). There is no suggestion in Jiao or any other apparatus at issue in this investigation that GPUs conduct processing activities untethered to end-user software applications. Indeed, Jiao’s Figure 3 shows the “graphics pipeline 208” involving each of VS 304, GS 306, and PS 310 and beginning with EU Pool Control 206. RX-0423 at Fig. 3, 7:56-60. And the 381 patent teaches in the background of the art, “[t]here are several layers of software between the end-user application and the GPU” and “[t]he API communicates with a driver. The driver translates standard code received from the API into a native format of instructions understood by the GPU. The driver is typically written by the manufacturer of the GPU. The GPU then executes the instructions from the driver.” 381 patent at 1:33-34, 1:41-46. So there is no reason to find that the assignments issued by the driver 128, eventually received by VS 304, GS 306, and PS 310, and then received by EUs 420 are not “linked” or related to one another. Hr’g Tr. (Reinman) at 108:8-111:7 (“Jobs consist of tasks, and so individual tasks comprise this particular job, and performing these jobs would include performing the tasks in them, and performing the tasks would be part of performing the job.”).

[REDACTED]

Second, AMD challenges Respondents’ use of the “same component of [Jiao] as satisfying the shader core limitation and the plurality of engines limitation.” CIB at 122. Specifically, AMD explains, “Dr. Pedram identifies the same component, Execution Pool Unit 414 as satisfying the claimed plurality of engines (*i.e.*, when executing the shader programs) and the claimed shader core. That is improper. In order to satisfy multiple elements, the prior art must disclose separate components.” *Id.* Respondents, however, do not contend the *entire* Execution Pool Unit (EUP) 414 is the “plurality of engines.” Under the first theory described above, the “engines” are only VS 304, GS 306, and PS 310, which as shown in Jiao’s Figure 4 are not within the EUP. RX-0423 at Fig. 4. Under their second theory, the “engines” are the same VS 304, GS 306, and PS 310 along with just the thread slots of the execution units 420, that are within the EUP 414. The claimed “shader core” is then the leftover portions of Jiao’s EUP 414. This second theory is akin to AMD’s own infringement theory which permitted the “plurality of engines” in the Accused Products to be not only the “Job Slots” but also “Frontends,” even though the latter are very much inside the products’ shader core and the former are very much outside. *See* CX-0773C.24, 26, 30; RIB at 82-84; RRB at 46. And, to AMD’s benefit, that infringement theory has been found persuasive above. Respondents’ rebuttal—that AMD did not present this dispute in its pre-hearing brief, and thus it is waived (RRB at 62)—is also well taken (*see* CPB at 165-168), but moot.

Third, AMD argues Jiao’s vertex shader 304, graphics shader 306, and pixel shader 310 cannot be the “plurality of engines” because “a POSITA would understand that shader programs are actually the constituent parts of a task that would be received by the engines and executed by the shader core.” CIB at 122-123 (citing, *inter alia*, 381 patent at 1:58-61). Respondents and Dr. Pedram persuasively respond, however, that Jiao uses “shader” more broadly than the 381 patent:

As Dr. Pedram testified, Jiao's shaders are components of Jiao's computational core that feed their output, which Jiao calls "shader tasks" and which meet the definition of tasks as construed by ALJ, to Jiao's execution unit pool (the "shader core") within Jiao's computational core. More precisely, Jiao's shader tasks are assigned an execution unit and a slot number in Jiao's execution unit pool and are sent to the specified slot in the specified execution unit to be executed by the corresponding execution unit. Tr. (Pedram) at 930:4-931:1 and 1014:13-18. Dr. Pedram's explanation is supported by express disclosures in the Jiao reference. RX-0423 at 9:15-21 and 9:35-37.

RRB at 64. This perspective is supported in no small part by the figure reproduced above and its accompanying disclosures, which unambiguously disclose vertex shader 304, graphics shader 306, and pixel shader 310 as block components of the "graphics processor 202." See, e.g., RX-0423 at 6:1-8:14, Figs. 3, 4.

Fourth, and relatedly, AMD contends, "as described by Dr. Hart and Jiao, the vertex shader, geometry shader, and pixel shader generate the alleged shader tasks." CIB at 123 (citing Hr'g Tr. (Hart) at 1270:1-16; Jiao at 9:35-41). AMD reasons, "[t]he shader programs cannot receive something that they create or generate." *Id.*³ While it is true that Jiao states, "various shader tasks are generated by the vertex shader 304, geometry shader 306, and pixel shader 310" (RX-0423 at 9:35-38), it cannot be seriously argued that vertex shader 304, geometry shader 306, and pixel shader 310 do not receive anything to prompt their generation of tasks. Jiao—and the 381 patent—are clear that for all practical purposes, graphics rendering starts with a host processor (*i.e.*, CPU) and its driver. RX-0423 at 5:7-23; 381 patent at 1:23-46. Dr. Pedram has credibly testified that a POSITA would understand this. Hr'g Tr. (Pedram) at 930:3-21, 1014:1-6. AMD's expert, Dr. Reinman, testified to the same basic principle:

Q. Can you explain to us how a CPU and a GPU generally work together to perform graphics processing?

³ Again, AMD cites to expert discovery materials not admitted as evidence, and so have not been considered. See, e.g., CIB at 123-124 (citing CX-0672C); CRB at 62 (citing CX-0672C); RRB at 64 n.10.

[REDACTED]

A. Yes, sir. The CPU will be processing certain set of applications; and then when one of those applications has a graphics processing task that it wants to run, it will make a request through the operating system or some driver to be able to make use of the GPU. The GPU will then execute that particular portion of the application on behalf, and perform the graphics processing, and then return back to the CPU to say it's done. It's basically being used as a proxy for graphics processing.

Q. Thank you. Are you familiar with the concept of a graphics pipeline?

A. Yes, sir.

Q. Can you explain to us how work typically flows through a graphics pipeline?

A. Sure. So, Your Honor, typically, there are work units that are passed. We talked before about how these applications are providing work for the GPU, offloading that work from the CPU to the GPU, and there would be units of work that would be provided from the CPU to the GPU itself. These work units have different names depending on who you talk to. Some people will call these tasks. Some will call them jobs. One ARM terminology term for this is a job. A job is provided from the CPU to the GPU.

Hr'g Tr. (Reinman) at 107:16-108:21. In addition, AMD itself describes Jiao as "a conventional graphics pipeline with a single pathway" (CIB at 120) and that pipeline starts with EU Pool Control 206 supplying assembled primitives to the VS 304 for processing (RX-0423 at Fig. 3, 6:17-27). Jiao continues, "EU pool control unit 206 also assembles geometry shader inputs and provides those inputs to the computational core 204 (input F). In general, the EU pool control unit 206 controls the respective inflow and outflow to the computational core 204." *Id.* at 7:23-27. So there is no merit to AMD's position here.

Fifth, AMD argues the "alleged engines are not independent" and, more specifically, "[t]he shader programs are not configured to be independent of one another and do not independently service a command buffer." CIB at 124. But this position seems to be derived from a claim construction for "engine" already considered and rejected above. *See* CIB at 77-78 (arguing as a

[REDACTED]

matter of claim construction that “a POSITA would understand that each engine must service its own command buffer”).

Sixth, AMD argues, “the evidence has shown that the identified engines do not load state data.” *See* CIB at 124-125. AMD explains:

Jiao ’774 describes sending state data (*i.e.*, rendering context information) to each execution unit 420. Hrg. Tr. (Day 5) (Hart) at 1275:3-8; CX-1869 (Jiao ’774) at 8:65-9:7. State data is not passed to the identified engines (*i.e.*, vertex, geometry, and pixel shaders). *Id.* at 1274:24-1275:15. Dr. Pedram relies on passing the rendering context, including constants, to the alleged shader core (computational core 204) as teaching the engines loading state data. Hrg. Tr. (Day 3) (Pedram) at 936:18-938:1. The logic identified by Dr. Pedram as loading state data (*i.e.*, constants) (Jiao, Fig. 3) is executed by the shader core, not the alleged engines (shader programs). CX-1869 (Jiao ’774) at 5:42-51.

Id. As determined, and applied above, however, Order No. 31 set a broad construction for “load state data” which “does not require that the engines load the state data into their own memories (as opposed to some other component or components), it does not require the engines (as opposed to some other component) to first extract the state data, and it does not specify how or from which component the shader core receives the state data.” Order No. 31 at 41. In view of this breadth, Jiao’s disclosure that “[c]onstants for performing one or more of the aforementioned computations are provided by a constant buffer of the constant buffer system 200” to vertex shader 304, geometry shader 306, and pixel shader 310 is sufficient. RX-0423 at 6:22-26 (vertex shader), 6:32-34 (geometry shader), 6:40-42 (pixel shader). Additionally, it should be appreciated that Jiao shares the same overall objective as the 381 patent—the elimination of latency from flushing contextual state data across threads:

One challenge to existing systems occurs in regard to the frequent updating of at least a portion of the constants in the constant buffer due to rendering of the sequence of objects with different visual properties defined in an appropriate rendering context. In general, an update of constants requires “flushing” the entire constant buffer having the “old” context related data and loading “new” constants corresponding to a new context.

Further, the execution unit corresponding to the constant buffer, which may be running a plurality of threads, may be stalled during the flushing of the constant buffer. That is, the execution unit may flush the threads (*e.g.*, to allow completion of processing using the previous constants) before execution using the set of new constants can commence. Thus, such rendering context switching (*e.g.*, from a first context to a second context) in the shader running on a multi-threaded computational core may cause significant performance penalties.

....

Disclosed herein are various embodiments of constant buffer (CB) systems and methods (herein, also collectively constant buffer system(s) or CB system(s)) implemented in each execution unit of a multi-threaded, parallel computational core of a graphics processing unit (GPU). A CB system enables the updating of a constant buffer with new context data related to a new rendered object while current or existing context data in the buffer is being used by shader program threads rendering a previous object. Such updating of the constant buffer occurs without stalling the computational core execution units for reload of constants. That is, as compared to existing systems that flush the constant buffer with each rendering context change (flush the “old” constants entirely from the buffer to make room to load new constants), the CB system enables the simultaneous processing of constants using the same or different shader from more than one rendering context.

Thus, unlike conventional systems, changes to constants in the preferred embodiments do not impose a penalty in speed and efficiency. In particular, the below described CB systems enable updates or changes to constants to occur while enabling execution of previous constants. That is, execution of shader programs with constants of a different context occurs in simultaneous paths, with one path corresponding to the previous context constants and at least one other path corresponding to the new context constants. The updating or replacement of constants through implementation of the various CB system embodiments enables a smooth procedure of constant replacement in shader execution blocks, which may result in an improvement in computational efficiency.

RX-0423 at 2:3-20, 3:31-60;

A graphics-processing task performed by a GPU typically involves complex mathematical computations, such as matrix and vector operations. To perform a single graphics-processing task, a GPU may execute a plurality of different threads (sequence of instructions). Each thread may comprise a shader program, such as a geometry shader, a pixel shader, a vertex shader, or the like. Each thread (*e.g.*, shader program) is typically associated with a set of state data (such as texture handles, shader constants, transform matrices, or the like) that is locally stored in data-storage units of the GPU. The locally stored state data is called a context.

....

After the more-important (*e.g.*, low-latency) task finishes executing, the state data associated with the threads of the more-important (*e.g.*, low-latency) task are flushed from the data-storage units, and the state data from the threads of the first task are swapped back into the data-storage units of the shader core. Then, the shader core can resume executing the threads of the first task.

Although context switching allows a GPU to process tasks out of the order in which they were placed in the command buffer, context switching is problematic for several reasons. As an initial matter, a substantial amount of time is required to perform a context switch, thereby limiting the performance of the GPU. Also, context switching requires additional local memory (*e.g.*, back-up storage units) to store the context that is being switched. The additional local memory takes up precious chip area, resulting in a larger GPU.

In addition to requiring a substantial amount of time and area, context switching makes the GPU ineffective at processing low-latency, high-priority tasks. To prepare the shader core for executing a low-latency, high-priority task, a conventional GPU must perform a context switch. The time associated with the context switch (*e.g.*, hundreds of clock cycles) makes the effective time for executing the low-latency, high-priority task relatively long, even though the actual time for executing the low-latency, high-priority task may be relatively short (*e.g.*, tens of clock cycles).

Given the foregoing, what is needed is a processing unit that can efficiently process important (*e.g.*, low-latency) tasks without context switching.

381 patent at 1:47-2:59.

Seventh, as to the claimed “scheduling module,” AMD observes that Dr. Pedram offered no testimony on Jiao’s “driver software 128” as meeting this limitation. CRB at 61-62. AMD reasons, “[a]s such, Respondents’ allegations with respect to the driver software are bare attorney argument and should be afforded no weight.” *Id.* (citing *Invitrogen Corp. v. Clontech Labs., Inc.*, 429 F.3d 1052, 1068 (Fed. Cir. 2005); *Estee Lauder, Inc. v. L’Oreal, S.A.*, 129 F.3d 588, 595 (Fed. Cir. 1997)). This is not persuasive. Dr. Pedram clearly testified that “on column 5, lines 17 to 23 [of Jiao], that does explain that there is driver software 128 residing in the system memory space 124, which communicates instruction sets or commands, tasks, through the use of the CPU 126 to

registers in the GPU 114.” Hr’g Tr. (Pedram) at 936:8-12. And Jiao does disclose this (RX-0423 at 5:7-23)—and more (*id.* at 7:56-60 (“The graphics pipeline 208 comprises fixed-function graphics processing functionality. Responsive to a command from the driver software 128, such as to draw a triangle, vertex information is passed to vertex shader 304 in the computational core 204 to implement vertex transformations.”)).

Eighth, assuming the “driver software 128” as “scheduling module” theory is entertained (it is), AMD argues it has not been shown that Jiao’s shaders “receive tasks from the driver software”:

The only strained evidence Respondents could provide was that Jiao ’774 describes the driver software communicating with “registers of the GPU 114” generally, and that Jiao ’774 suggests that there could be registers associated with the units shown in Fig. 3 of Jiao ’774. RIB at 116-117 (citing RX-0423 at 5:17-23 and 7:57-62); *See also, e.g.*, RDX-0004C.99; Hrg. Tr. (Day 3) (Pedram) at 936:7-16. This ignores the claim requirements and does not address if the registers of the shaders receive tasks from the driver software. *See* CX-0672C (Pedram Tr.) at 340:17-19 (confirming that the shaders receive only “graphics primitives”).

CRB at 62. This is incorrect. As noted in preceding paragraphs, Jiao discloses that “vertex shader 304” within the “computational core 204” is responsive to commands from the driver software 128—and that graphics shader 306 and pixel shader 310 act on the data outputted by vertex shader 304 in a pipeline. RX-0423 at Fig. 3, 6:1-48, 7:56-60.

Lastly, AMD contends Jiao’s “EU Pool Control 206” cannot be the “scheduling module” because it is part of the GPU and therefore not “associated with” a “second processing unit.” CIB at 125; CRB at 62-63. AMD makes a related point in connection with Stuttard, as discussed below. *See* CIB at 129-30. This presents an issue of claim construction, and more precisely of prosecution history disclaimer.

As noted by AMD, prosecution history disclaimer occurred which limited the scope of “associated with.” See CIB at 125. Specifically, the applicant made, among others, the following amendment to the claims:

Reply to Final Office Action of January 14, 2014 - 5 - MANTOR et al.
Appl. No. 12/874,134

13. *(Currently Amended)* A computing system comprising:
a memory;
a first processing unit; and
a bus coupled to the memory and the first processing unit, wherein the first processing unit comprises:
a plurality of engines configured to receive, from ~~an operating system~~
~~(OS)~~ a scheduling module associated with a second processing unit, a plurality of tasks
and to load state data associated with each of the plurality of tasks; and
a shader core configured to receive the plurality of tasks from at least one
of the plurality of engines and to execute a first task from the plurality of tasks while
executing a second task from the plurality of tasks based on respective state data
associated with each of the first and second tasks.

JX-0009.178. This new “scheduling module associated with a second processing unit” language was an important update as the patentee’s subsequent remarks rely on it to avoid prior art (identified as Jiao, but not the reference applied in this investigation):



Jiao discloses a graphics processor with parallel scheduling and execution of threads. *See* Jiao, Title. In referring to FIG. 2 of Jiao (reproduced below), Jiao states that “[w]ithin shader core 130, a multiplexer (Mux) 208 receives threads from programs 110a through 110n and provides these threads to a thread scheduler 210. Thread scheduler 210 performs various functions to schedule and manage execution of the threads . . .” *Id.*, 2:58-62. In particular, “[t]hread scheduler 210 may determine which hardware unit to use for each instruction . . . Thread scheduler 210 may then issue different instructions to

Atty. Dkt. No. 1972.1540001

AMD1318_0008924

different hardware units for execution in parallel, which may significantly increase throughput.” *Id.*, 4:19-23. In other words, thread scheduler 210 is within shader core 130 and issues instructions to hardware units that are also within shader core 130. Thread scheduler 210 and the hardware units are not associated with different processing units; rather, they are associated with the same processing unit—namely, shader core 130. Jiao therefore does not teach or suggest “a shader core associated with the first processing unit and configured to receive the plurality of tasks from at least one of the plurality of engines . . .” in which the plurality of engines are “associated with a first processing unit and configured to receive, from a scheduling module associated with a second processing unit, a plurality of tasks and to load state data associated with each of the plurality of tasks,” as recited in claim 1 (emphasis added) and similarly-recited in claims 7, 13, 19 and 25.




FIG. 2 of Jiao

Since Jackson and Jiao do not teach or suggest at least the above-noted distinguishing features, these references cannot be used to establish a *prima facie* case of

JX-0009.184-185. The only substantive reason given for why the applied reference did not teach the limitation is that the reference's thread scheduler 210 is "within" the shader core, and therefore "associated with" that shader core, when it needed to be associated with a different "second processing unit." Thus, as a matter of claim construction, to meet the "associated with a second processing unit" limitation the scheduling module cannot be "within" or "embedded" (CIB at 130 (emphasis omitted)) in the "first processing unit."

Contrary to Respondents' position (RIB at 119 ("the patentee's statements are directed to a claim limitation not recited in the Asserted Claims")), the remarks plainly encompass any 381 patent claim—presently asserted or not—that recites a "scheduling module" "associated with" a processor. *Phillips*, 415 F.3d at 1314 ("Because claim terms are normally used consistently throughout the patent, the usage in one claim can often illuminate the meaning of the same term in other claims."). Respondents' other points in opposition (RIB at 118-120; RRB at 66, 69) do not overcome this intrinsic evidence.

This construction, however, does not detract from Respondents' invalidity case based on Jiao. Respondents did not identify Jiao's entire GPU 114 as the "first processing unit"—only its computational core 204. RIB at 113 (citing, *inter alia*, Hr'g Tr. (Pedram) at 932:4-8). And AMD has not pointed to any disclosure that EU Pool Control 206 is within the computational core 204. In fact it appears to be outside. RX-0423 at Fig. 4. AMD's contention is also irrelevant in light of Respondents' alternate theory that "driver software 128" constitutes the "scheduling module," because the driver software 128 is unequivocally outside the GPU 114. *See* RX-0423 at Fig. 1.

Accordingly, Respondents have met their burden to show Jiao discloses limitation [15d].

Limitation [15e] reads, "a shader core configured to receive the plurality of tasks from at least one of the plurality of engines and to execute a first task from the plurality of tasks while

executing a second task from the plurality of tasks based on respective state data associated with each of the first and second tasks.” 381 patent at cl. 15. Respondents contend that Jiao discloses this limitation through its Execution Unit Pool (EUP) 414:

Jiao’s computational core 204 comprises an execution unit pool (EUP) 414, which comprises one or more execution units (EUs) 420 a through 420 h (collectively referred to herein as EUs 420). Tr. (Pedram) at 939:22-942:11; RX-0423.00015 at 8:15-18. Each of the EUs 420 includes control logic and local memory, and each EU 420 is capable of processing multiple instructions within a single clock cycle. RX-0423.00015 at 8:18-21. The EUP is configured to receive the plurality of tasks from at least a vertex shader, a geometry shader, and/or a pixel shader (*i.e.*, the “engines”) in graphics pipeline processing. *See, e.g.*, RX-0423.00016 at 9:35-38 (“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.”). A POSITA would have understood Jiao’s EU pool (EUP) 414 to be a “shader core” as recited in claim 15. *See* Tr. (Pedram) at 938:21-939:3; RDX-0004C.102.

Jiao’s EUP is configured to execute multiple, simultaneous tasks based on respective state data including constants, rendering context information, and graphics data. *See, e.g.*, RX-0423.00013-15 at 3:48-60, 6:49-7:18. Further, the EUP is configured to execute a first task while executing a second task. *See, e.g.*, RX-0423.00015 at 8:21-23 (“Thus, the EU pool 414, at its peak, can process multiple threads simultaneously or substantially simultaneously.”). Each task requires its corresponding context-specific constants to be able to be mapped or executed on some execution unit of the computational core. *See, e.g.*, RX-0423.00016 at 1:43-58, 6:49-7:18, 8:15-9:14, 9:25-54, 9:65-10:8; *see also* Tr. (Pedram) at 936:17-938:1; RDX-0004C.104.

RIB at 120-121.

AMD’s opposition to finding this limitation in Jiao is not persuasive. First, AMD argues that the first and second tasks must be first and second as “determined by the scheduling module.” CIB at 125. AMD posits, “tasks, as defined by the ’381 Patent, are higher level tasks. . . . A POSITA would not assume any arbitrary collection of threads and/or shader programs is necessarily a ‘task.’” *Id.* at 125-126 (citing Hr’g Tr. (Hart) at 1275:20-1276:19). AMD continues:

The evidence has shown that the ability of a GPU to process multiple shader programs and/or threads in parallel does not necessarily mean that that GPU necessarily is capable of receiving or executing multiple tasks at the same time. *Id.*

[REDACTED]

Those multiple shader programs and/or threads can pertain to the same, singular, task. *Id.*; *See also, e.g.,* CX-0672C (Pedram Tr.) at 262:4-14 (agreeing that executing multiple threads at the same time is not necessarily the same as executing multiple tasks at the same time) (emphasis added). Absent evidence, among other things, linking those multiple threads / shader programs to different tasks received from the scheduling module – and there is none – Dr. Pedram cannot demonstrate that the GPU is capable of executing multiple tasks at the same time.

Id. at 126; *see* CRB at 63-64. But “task,” as used in claim 15, has been construed as “unit of processing work or portion of an application to be executed.” There is no further demarcation of high-level, low-level, or anything in-between. Thus, to the extent a single thread—a single sequence of instructions—is a “portion of an application to be executed,” then it qualifies as a “task.” And AMD has firmly taken the position, in connection with infringement, that working on low-level activities qualifies as working on the high-level tasks those low-level activities derive from:

Dr. Pfister’s testimony also ignores that a POSITA would understand that by receiving and executing the warps/quads or threads of a task/ARM Job, the Mali GPU is receiving and executing the task/ARM Job. Indeed, the ’381 Patent itself discloses that “[t]o perform a single graphics processing task, a GPU may execute a plurality of different threads (sequence of instructions).” JX-0003 (’381 Patent) at 1:49-51; Hrg. Tr. (Day 1) (Reinman) at 130:10-19, 228:21-230:8. Dr. Pfister also ignores the ’381 Patent’s disclosure in Fig. 2 and its related description, that the tasks may undergo pre-processing before being forwarded along to the shader core for execution. *See, e.g.,* JX-0003 (’381 Patent) at Fig. 2; *see also id.* at 7:25-8:43, 14:1-15. The ’381 Patent discloses, for example, a scan converter 212 (which is located between the engines 202D and 202E and shader core 214) that traverses a task’s primitives by converting them to pixels, and then forwarding along the pixels to the shader core for execution of the task. *See, e.g.,* JX-0003 (’381 Patent) at Fig. 2; *see also id.* at 7:25-8:43, 14:1-15. The shader core 214 receives pixels from the scan converter 212, which is not the same entity that is received and output by the 3D engines 202D and 202E. *Id.* Therefore, the ’381 Patent, and claim 15 in particular, allows the shader core to receive tasks, indirectly (*i.e.,* via other modules) and in a different form of expression, from the tasks received by the engines. *Id.* Thus, Dr. Pfister’s testimony is contrary to a POSITA’s understanding of what it means to execute a task, and should be rejected. *See Phillips*, 415 F.3d at 1312-13.

[REDACTED]

CIB at 100; *see id.* at 97-98. AMD has further taken the position, also in connection with infringement, that if a first task involves vertices and a second task involves pixels/fragments, then they are *de facto* different tasks in satisfaction of the claim. *See id.* at 89, 97-99 (citing, *inter alia*, Hr’g Tr. (Reinman) at 161:14-163:22 (“So the first task here could be the fragment shader requests, and the second task here could be the vertex shader request.”); JX-0192C (Larri) at 114:7-8)). In light of these positions, Jiao’s disclosure that its execution units 420 act as a common pool for VS 304, GS 306, and PS 306 to draw upon as each completes its assigned work is sufficient to meet the claim:

In general, an embodiment of the graphics processor 202 utilizes programmable vertex 304, geometry 306, and pixel shaders 310. Rather than implementing the functions or operations of these components as separate fixed function units with different designs and instruction sets, the operations are instead executed by the pool 414 of execution units 420 a, 420 b . . . 420 n with a unified instruction set. Each of these execution units 420 is identical in design and configurable for programmed operation. In one embodiment, each execution unit 420 is capable of multi-threaded operations simultaneously. *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.*

As individual tasks are generated, the EU pool control unit 206 handles the assigning of those tasks to available threads within the various execution units 420. As tasks are completed, the EU pool control unit 206 further manages the release of the relevant threads. In this regard, 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, and also, performs an associated “bookkeeping” of the tasks and threads. Specifically, the EU pool control unit 206 maintains a resource table of threads and memories for all execution units 420. The EU pool control unit 206 particularly knows which threads have been assigned tasks and are occupied, which threads have been released after thread termination, how many common register file memory registers are occupied, and how much free space is available for each execution unit.

RX-0423 at 9:25-54 (emphasis added). AMD’s additional argument that “the structure of Jiao ’774 prevents parallel execution of tasks” (CIB at 126-127 (citing Hr’g Tr. (Hart) at 1266:15-18)) is simply inaccurate in light of these disclosures.

Accordingly, Respondents have met their burden to show Jiao discloses limitation [15e], and thus all limitations of claim 15. Claim 15 is determined to be invalid under 35 U.S.C. § 102. There is no need to discuss obviousness which, in any event, has not been pled with any particularity as to the contested limitations. *See* RIB at 133-134; RRB at 70-71.

b. Claim 16

Claim 16 depends upon claim 15 and recites, “wherein the shader core comprises: 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 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.” 381 patent at cl. 16. Respondents contend this is met by Jiao’s EUP 414, comprised of multiple EUs 420, and the disclosure that “[e]ach of the EUs 420 includes control logic and local memory, as described further below, and each EU 420 is capable of processing multiple instructions within a single clock cycle.” RIB at 122 (citing RX-0423 at 8:18-23). Respondents continue, “Jiao’s EUP is configured to execute multiple threads simultaneously or substantially simultaneously based on respective state data including constants, rendering context information, and graphics data.” *Id.* (citing RX-04223 at 6:1-64). Respondents highlight Jiao’s disclosure that “[a]s 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.” *Id.* (citing RX-0423 at 9:25-28).

Respondents have shown the claim is disclosed in Jiao. In addition to those passages noted above, Jiao discloses that “[c]onstants for performing one or more of the aforementioned computations are provided by a constant buffer of the constant buffer system 200” to vertex shader 304, geometry shader 306, and pixel shader 310 for processing their respective threads. RX-0423 at 6:22-26 (vertex shader), 6:32-34 (geometry shader), 6:40-42 (pixel shader). Additionally, the constant buffer system holds the constants for multiple rendering contexts (*id.* at 4:12-14), and

[REDACTED]

there is no dispute that the rendering context includes state data (*id.* at 8:65-9:14; *see* CIB at 124-125). AMD offers no dispute for claim 16 specifically beyond those points discussed above for limitation [15e]. *See* CIB at 124-125; CRB at 64-65.

Accordingly, Respondents have met their burden to show Jiao anticipates claim 16. There is no need to discuss obviousness which, in any event, has not been pled with any particularity as to claim 16. *See* RIB at 133-134; RRB at 70-71.

c. Claim 17

Claim 17 depends upon claim 15 and recites, “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.” 381 patent at cl. 17. Respondents offer two theories as to how this claim is met, but both depend upon Jiao’s “thread slots,” which exist in each of the execution units 420 and enable the storage of a plurality of tasks or threads, potentially up to 32 at the same time. *See* RIB at 123-124 (citing, *inter alia*, Hr’g Tr. (Pedram) at 948:4-949:11). Respondents remind that one of their theories as to “plurality of engines” includes all of vertex shader 304, geometry shader 306, and pixel shader 310 in addition to the “thread slots” which exist in the EUs 420. *Id.* at 122.

Respondents have shown the claim is disclosed in Jiao as there is no dispute that “thread slots” are queues. AMD argues that the theory fails because the thread slots belong to Jiao’s computational core 204 and so cannot be considered part of the alleged “plurality of engines.” CIB at 127; CRB at 64. But as discussed above, this is not persuasive and contrary to AMD’s approach to infringement, which combines the Accused Products’ Job Slots and Frontends.

Accordingly, Respondents have met their burden to show Jiao anticipates claim 17. There is no need to discuss obviousness which, in any event, has not been pled with any particularity as to claim 17. *See* RIB at 133-134; RRB at 70-71.

d. Claim 19

Claim 19 depends from claim 15 and reads, “wherein the first task comprises a graphics-processing task and the second task comprises a general-compute task.” 381 patent at cl. 19. Respondents do not contend this limitation is disclosed in Jiao, but do contend it would have been obvious. RIB at 134-135. Respondents explain:

[T]he system comprises components such as arithmetic logic units and other logic and performs functionalities suitable for graphics-processing tasks and general-compute tasks. RX-0423.0014 at 6:1-48, 11:4-13. Further, Jiao teaches that the graphics processor 202 utilizes programmable vertex 304, geometry 306, and pixel shaders 310 and that each of these execution units 420 is identical in design and configurable for programmed multi-threaded operations. RX-0423.00014, 17 at 9:25-38. A POSITA would have understood and found it obvious that such programmable execution units can be used to do both graphics-processing and general-compute tasks. Tr. (Pedram) at 949:12-950:12; RDX-0004C.117.

Id.

Respondents also contest that any indicia of non-obviousness apply because AMD has failed to show the required nexus. *See generally* RIB at 135-138; RRB at 72-73 (citing, *inter alia*, 381 patent at 2:8-19). Specifically, Respondents argue that AMD’s evidence relates to a feature called “asynchronous dispatch” but “the asserted claims do not require asynchronous dispatch [which] was already known in the art.” RIB at 135 (citing Hr’g Tr. (Pedram) at 978:15-980:5); *see id.* at 136 (citing Hr’g Tr. (Hart) at 1288:13-15, 1262:20-1263:1). Respondents add that, for commercial success, AMD has failed to show the successful products are co-extensive with the alleged invention of the 053 patent. *Id.* at 137-138 (citing, *inter alia*, *Teva Pharms. Int’l GmbH v. Eli Lilly & Co.*, 8 F.4th 1349, 1360 (Fed. Cir. 2021)).

In opposition, AMD argues:

[T]he evidence has shown that Jiao ’774’s system is expressly directed to processing of shader programs in a single graphics pipeline. What Dr. Pedram points to as processing compute tasks within EUPD 514 would be understood by a POSITA as relating to general capabilities of arithmetic logic units. *See, e.g.*, CX-1869 (Jiao ’774) 6:1-48, 9:25-38. This general capability would not teach or suggest

[REDACTED]

to a POSITA that Jiao '774's system could be modified to conform with the requirements and limitations of the shader core of claim 15(e).

CIB at 135. AMD also provides a lengthy discussion of indicia of non-obviousness tied to the idea that the 381 patent claims represent “asynchronous task dispatch.” *See generally* CIB at 74-75, 136-139.

Respondents have not met their burden. While it seems more than likely that at the time of the 381 patent invention it was known that GPUs could be used for general compute tasks in addition to graphics rendering tasks (*see, e.g.*, 381 patent at 10:14-19 (“During the course of running, applications 402 may generate a plurality of tasks 404A-D to be executed on GPU 110 rather than CPU 102. Tasks 404 may comprise data-parallel processing tasks (*e.g.*, graphics-processing tasks, general-compute tasks, or the like) that GPU 110 can likely perform faster than CPU 102 . . .”); RX-0416 at 11:16-25), Respondents have not made a clear and convincing case that the EUs 420 of Jiao already function in this way or could be made to with modification. At the hearing, Dr. Pedram seemed to present an opinion that compute tasks were actually disclosed in Jiao as opposed to an obvious use or alteration—which is at odds with Respondents’ briefing. *See Hr’g Tr. (Pedram)* at 950:1-12 (citing RX-0423 at 11:4-8). He further did not present any motivation that a POSITA would have had to use Jiao for compute tasks alongside the disclosed graphics rendering tasks, and it does not otherwise appear Jiao is set up to handle general compute tasks (*see, e.g.*, RX-0423 at Fig. 4 (shader programs are only vertex, geometry, and pixel)). This is a critical omission. *InTouch Techs.*, 751 F.3d at 1351.

Accordingly, Respondents have not shown by clear and convincing evidence that claim 19 is invalid as obvious in light of Jiao.

e. Claim 20

Claim 20 depends on claim 15 and reads, “wherein the shader core is configured to receive the plurality of tasks based on a priority associated with each of the plurality of tasks.” 381 patent at cl. 20. Respondents contend it is disclosed by Jiao. RIB at 124 (citing Hr’g Tr. (Pedram) at 950:4-952:3). For support, Respondents argue “Jiao explains that the constant blocks may be loaded to the CB system 200 using a fast access path 506 or a slow access path 508 based on the high or low priority level 504 associated with the respective rendering context.” *Id.* (citing RX-0423 at 10:13-54); *see* RRB at 67 (citing RX-0423 at 2:24-36). In reply, Respondents add, “Dr. Pedram testified that Jiao’s system will execute tasks on high and low priority levels (fast or slow access paths), where each must be associated with a respective rendering context. . . . It is beyond dispute that the priority is used to establish the timing and execution of the reception of tasks by the shader core.” RRB at 67 (citing RX-0423 at 10:13-54).

AMD opposes and alleges the priority is not used for reception of tasks but for loading of constants (*i.e.*, claimed “state data”):

[W]hat Respondents allege as meeting this claim limitation, *i.e.*, the loading of constants using different access paths and priorities (fast access path 506, slow access path 508, and the priority level 504 of constant blocks) are instead priorities associated with access to data by the execution units to load contexts and associated constant data, not the reception of tasks by the alleged shader core. *See, e.g.*, CX-1869 (Jiao ’774) at 10:28-32; Hrg. Tr. (Day 3) (Pedram) at 951:5-19, Hrg. Tr. (Day 5) (Hart) at 1277:17-1278:1.

CIB at 127; *see* CRB at 65.

AMD is persuasive here. The priority level disclosed in Jiao relates to the loading of constants from the top-level API space 130 to the constant buffer (CB) system 200. RX-0423 at 10:13-18, Fig. 5. There is no disclosure that the priority level is used for any other purpose, such as what claim 20 requires—*i.e.*, the transfer of threads (the claimed “tasks”) from the combined

[REDACTED]

shaders/thread slots (the claimed “engines”) to the remainder of the EUs (the claimed “shader core”). *See* RX-0423 at 2:24-36, 10:13-54.

Accordingly, Respondents have not shown that claim 20 is anticipated by Jiao. Nor have they shown it would have been obvious, which has not been pled with any particularity as to claim 20. *See* RIB at 133-134; RRB at 70-71.

2. Stuttard

As with Jiao, Respondents contend claims 15-17, 19, and 20 are, collectively, anticipated or rendered obvious by Stuttard, which, among other things, describes a data processing apparatus with multiple processing elements being operable to process respective groups of data items. RIB at 45, 110; RX-0416 at Abstract. As discussed below, Respondents have not shown invalidity of claims 15-17, 19, or 20 under 35 U.S.C. § 102 as anticipated by Stuttard or under 35 U.S.C. § 103 as obvious in view of Stuttard.

a. Claim 15

With respect to claim 15, Respondents contend Stuttard anticipates and renders obvious all of its limitations. *See* RIB at 124-132, 135. Indeed, limitations [15d] and [15e] are the only challenged limitations in AMD’s post-hearing briefing. *See* CIB at 128-131; CRB at 65-68. In light of the evidence and expert testimony provided by Respondents’ expert, Dr. Pedram, limitations [15a]-[15c] are determined to be taught in Stuttard as alleged. *See* RIB at 125 (citing Hr’g Tr. (Pedram) at 964:4-20).

For limitation [15d], Respondents contend the “plurality of engines” consists of Stuttard’s “array controller 104, channel controller 108, and components of the thread manager 102 including cache memory unit 1024, status block 1030, and thread processors 1026.” RIB at 125 (citing RX-0416 at Figs. 3, 4, 6:6-19; Hr’g Tr. (Pedram) at 965:19-25). As for the “scheduling module associated with a second processing unit,” Respondents look to Stuttard’s “instruction fetch unit

[REDACTED]

1023 and thread scheduler 1025, working together” and then also the “embedded processor unit EPU 8.” *Id.* (citing RX-0416 at Figs. 3, 4; Hr’g Tr. (Pedram) at 966:1-11). The relevant operation, according to Respondents, is as follows:

[T]he instructions for a thread are fetched from the local external memory 103 or from the EPU 8 by the fetch unit 1023. RDX-0004C.132; RX-0416.0126 at 23:13-28; Tr. (Pedram) at 966:13-25. These instructions are supplied to the cache memory 1024 (which is one of the identified engines in Stuttard) via connecting logic. RDX-0004C.132; RX-0416.0126 at 23:13-28. The thread scheduler 1025 then schedules the execution of each thread, ensuring that each processor block 106 is kept busy as much as possible. *Id.* Thus, each of the fetch unit 1023 and/or thread scheduler 1025, are associated with the EPU, which directly feeds the fetch unit 1023 the plurality of tasks. Tr. (Pedram) at 966:1-25. As Dr. Pedram showed at trial, a task may comprise one or more threads. RDX-0004C.107 For example, the “Multi Task” drawing illustrates a scenario where two tasks are active, a first task comprising 5 threads and a second task comprising 3 threads. As explained above, the threads for these two tasks are supplied to the cache memory 1024 by the fetch unit 1023. When all 8 threads associated with the two tasks are received by the cache memory 1024, the corresponding tasks will also have been received by the identified engines (which include the cache memory 1024.) Similarly, when the 8 threads are sent to the processing blocks 106 by the thread processors 1026, the corresponding tasks will also have been sent to the processing blocks 106. Finally, when the processing of these threads is completed, the processing of the corresponding tasks will also have been completed.

Id. at 125-126. For background, Stuttard’s Figure 3 shows a block diagram of a processing block 106, and Figure 4 shows the thread manager 102, both reproduced below:




FIG. 3




FIG. 4

RX-0416 at Figs. 3, 4. As for “load[ing] state data,” Respondents contend:

[REDACTED]

Stuttard expressly states that the status block 1030 receives status information 1036 from each of the threads associated with a task, which is then transferred to the thread scheduler 1025 (“by the status block” – one of the claimed engines), which uses the status information to schedule the threads. Tr. (Pedram) at 899:19-901:9, 965:19-966:25; RX-0416.0026 at 23:13-28. Dr. Pedram explained that the status information is state data. Tr. (Pedram) at 967:18-23.

Id. at 127-128.

AMD disputes that limitation [15d] is met on at least five distinct grounds, several of which have been discussed previously with respect to Jiao and to Stuttard as to the 053 patent. First, AMD argues that Stuttard fails to disclose “tasks,” relying on Stuttard’s statements that tasks are made of threads which are made of instructions, and that the threads which are processed have not been shown to be of different “tasks.” See CIB at 128-129 (citing, *inter alia*, RX-0416 at 10:7-17); CRB at 66-67. As explored above, however, the ordered construction for “task” in this investigation is quite broad: “unit of processing work or portion of an application to be executed.” Order No. 31 at 38. Thus, Stuttard’s threads, which are indisputably portions of an application to be executed (RX-0416 at 11:16-25) qualify as the claimed “tasks.” And Respondents aptly observe that “AMD’s infringement argument is that tasks are comprised of threads (*i.e.*, threads are associated with tasks) such that receiving threads means you are receiving the associated tasks.” RRB at 68 (citing CIB at 81, 92, 97, 100). Moreover, Stuttard expressly discloses, as reproduced above in connection with 053 patent validity, that:

The threads are assigned priorities relative to one another. Of course, although the example described here has eight threads, any number of threads can be controlled in this manner. At any particular moment in time, each thread may be assigned to any one of a number of tasks. For example, thread zero may be assigned for general system control, thread 1 assigned to execute 2D (two dimensional) activities, and threads 2 to 7 assigned to executing 3D activities (such as calculating vertices, primitives or rastering).

RX-0416 at 11:17-25. Thus, Stuttard discloses “a plurality of tasks” and AMD’s criticism has no merit.

[REDACTED]

AMD further argues “Dr. Pedram provided no testimony at trial that each of Stuttard’s alleged engines (i) receives a task from a scheduling module, (ii) loads state data of the task, or (iii) independently services a command buffer.” CIB at 129 (emphasis added). This is not persuasive based on the operative word “each,” because, again, as discussed above, it is not a requirement of the claims that each engine in the plurality receive the tasks from a scheduling module. Only the plurality, as a collective, needs to do this. 381 patent at cl. 15.

AMD argues none of the alleged engines “load state data.” CIB at 129. Specifically, AMD views Stuttard’s “status block 1030” as merely “monitor[ing] the status of threads or load[ing] information related to lower-level instructions.” *Id.* (citing Hr’g Tr. (Pedram) at 967:2-23; RX-0416 at 44:1-8; Hr’g Tr. (Hart) at 1281:14-23). This is wholly without merit. Stuttard explicitly discloses, in a section directly describing status block 1030, that “[t]he status information is used by the thread scheduler 1025 to determine which thread should be active at any one time.” RX-0416 at 12:15-18. Moreover, “determin[ing] which thread should be active” is a prime purpose of Stuttard, as it seeks to reduce problematic context switching during thread switches—*i.e.*, the aim of the 381 patent as well. *Id.* at 1-7; 381 patent at 2:20-56. Nor does AMD explain its argument that “[a] single subcomponent [like status block 1030] cannot serve the function of loading state data of the plurality of tasks on behalf of the plurality of engines,” which is plainly meritless in any event. CRB at 66. There is no reason not to consider the data stored in status block 1030 as “state data.”

Additionally, AMD contends “Stuttard’s alleged engines are not configured in parallel, and therefore do not operate independently as required by the ’381 Patent.” CIB at 129. The need for engines to operate independently, however, is a flawed claim construction already discussed and rejected above. *See, e.g.*, RRB at 68.

[REDACTED]

Next, AMD argues that Stuttard's fetch unit 1023 and thread scheduler 1025 cannot meet the claimed "scheduling module associated with a second processor" because they are part of the processing core of the claimed "first processor" (*i.e.*, GPU) and thus cannot be "associated with" a second processor. *See* CIB at 129-130. AMD also argues that the "embedded processing unit (EPU) 8," alleged to be the claimed "second processor," is likewise part of the GPU. *See id.* (citing Hr'g Tr. (Hart) at 1278:22-1283:14).

There is no dispute that Stuttard's fetch unit 1023 and thread scheduler 1025 are within the first processing unit and not the second. RX-0416 at 11:3-10 ("The thread manager also includes an instruction fetch unit 1023, a thread scheduler 1025, thread processors 1026, a semaphore controller 1028 and a status block 1030."); RIB at 127 (citing RX-0416 at 11:12-14). Given the construction accorded "associated with" above, the limitation is not disclosed. Unlike with Jiao, Respondents offer no alternative "scheduling modules."

Accordingly, Respondents have not shown limitation [15d] is disclosed in Stuttard. As Respondents have not put forward an obviousness theory for this element (RIB at 135; RRB at 71) the limitation is also not obvious, and so claim 15 has not been shown to be invalid for anticipation or obviousness.

For the sake of completeness, limitation [15e] is now considered. Respondents contend it is disclosed in Stuttard. Specifically, according to Respondents, the "plurality of processing blocks 106" are the claimed "shader core" and receive tasks from the thread manager 102. RIB at 128 (citing RX-0416 at 11:27-12:7; Hr'g Tr. (Pedram) at 968:11-25). Respondents remind that the thread processors 1026 within the thread manager 102 are part of the claimed "plurality of engines" and assert that they "control the issuance of core instructions from the thread manager so that the system can maintain the simultaneous active program threads within the processing blocks." *Id.*

[REDACTED]

As to whether Stuttard discloses executing a first task while executing a second task, they argue Stuttard is so configured “based on its disclosure of multithreading and multitasking.” *Id.* (citing RX-0416 at 10:7-17, 11:16-25; Hr’g Tr. (Pedram) at 969:2-971:14, 972:1-7, 973:12-974:20.

Respondents continue:

Stuttard discloses that multithreading refers to the use of several threads to perform a single task, whereas multitasking refers to the use of several threads to perform multiple tasks simultaneously. *Id.* Stuttard expressly discloses using at least eight threads. RX-0416.0014 at 11:16-25 (“Of course, although the example described here has eight threads, any number of threads can be controlled in this manner.”); RDX-0004C.142. With at least eight threads, Stuttard can perform both multithreading and multitasking. Tr. (Pedram) at 969:16-18.

Id. at 128-129. Respondents then point to that passage in Stuttard’s use of multitasking, and argue activities concerning vertices would be a first task and activities involving “primitives or rastering” would be a second task:

The threads are assigned priorities relative to one another. Of course, although the example described here has eight threads, any number of threads can be controlled in this manner. At any particular moment in time, each thread may be assigned to any one of a number of tasks. For example, thread zero may be assigned for general system control, thread 1 assigned to execute 2D (two dimensional) activities, and threads 2 to 7 assigned to executing 3D activities (such as calculating vertices, primitives or rastering).

RX-0416 at 11:17-25; *see* RIB at 130.

The limitation is taught as alleged. Although it was previously determined in connection with the 053 patent that Stuttard does not actually disclose the simultaneous execution of vertex threads and pixel threads, claim 15 of the 381 patent does not require this specificity. Rather, the shader core need only execute threads pertaining to a “first task” and “second task” of any identity so long as they are distinct. 381 patent at cl. 15. As discussed for Jiao, it is only at dependent claims 18 and 19 where the tasks’ identities becomes relevant. And Stuttard is clear that, for example, a first general system control task may be executing while a 3D activity task is executing.

[REDACTED]

This is sufficient for claim 15. AMD's opposition is largely based on points rejected above; namely, that Stuttard only discloses the processing of one instruction across all processing blocks 106 (CIB at 130-131; CRB at 66-68) and that "task" as used in the claims can only mean a sort of high-level task that comes out of the scheduler (CIB at 131; CRB at 68).

Accordingly, Respondents have shown the limitation [15e] is disclosed in Stuttard, but claim 15 is not rendered invalid in light of limitation [15d].

b. Claim 16

Respondents contend claim 16 is disclosed in Stuttard. For the "first plurality of processing elements" and "second plurality of processing elements," each configured to execute the first and second tasks based on state data, Respondents look to Stuttard's processing blocks 106 and disclosures of multitasking and that "[a]t any particular moment in time, each thread may be assigned to any one of a number of tasks." RIB at 132 (citing RX-0416 at 11:16-25 Fig. 3; Hr'g Tr. (Pedram) at 973:12-22). They reason, "[o]ne task could use a plurality of processing blocks, while the remaining processing blocks perform the second task." *Id.*

Respondents have met their burden to show the limitation is disclosed. As noted above, the construction for "task" is quite broad; and between Stuttard's acknowledgment of multi-tasking (RX-0416 at 10:10-17), that each processing block 106 operates as a SIMD (*id.* at 17:28-32), and that at any moment different threads can be used, for example, for 2D activities and 3D activities (*id.* at 11:16-25), the limitation is met. AMD's opposition is bundled with its opposition to limitation [15e] (CIB at 130-131; CRB at 66-68) and found not persuasive, above.

Accordingly, Respondents have shown the additional limitation of claim 16 is disclosed in Stuttard, but claim 16 is not rendered invalid in light of its dependence on claim 15.

c. Claim 17

Respondents contend claim 17 is disclosed in Stuttard. Respondents allege “Stuttard discloses that its cache memory can be replaced by a series of first-in/first-out buffers (FIFO), one per thread.” RIB at 132 (citing RX-0416 at 23:3-11; Hr’g Tr. (Pedram) at 975:3-14). They continue, “[w]ith different tasks running on the Stuttard system, each task would be using queues, *i.e.*, the FIFO buffers, for storage. . . . FIFO is just another term for a queue, as known by a person of ordinary skill in the art.” *Id.*

The limitation is taught. Given the broad construction of “task,” a FIFO buffer dedicated to each of eight threads, and the corresponding processing block 106 according to Figure 4, qualifies as “a first queue configured to receive the first task” and “a second queue configured to receive the second task.” 381 patent at cl. 17. AMD’s opposition is limited to the false assumption that threads cannot qualify as the claimed “tasks” (*see* CIB at 131; RRB at 70) discussed above.

Accordingly, Respondents have shown the additional limitation of claim 17 is disclosed in Stuttard, but claim 17 is not rendered invalid in light of its dependence on claim 15.

d. Claim 19

As with Jiao, Respondents do not explicitly contend claim 19 is disclosed in Stuttard, but instead assert it would have been “considered” to be taught through the “[a]t any particular moment in time” disclosure:

The threads are assigned priorities relative to one another. Of course, although the example described here has eight threads, any number of threads can be controlled in this manner. At any particular moment in time, each thread may be assigned to any one of a number of tasks. For example, thread zero may be assigned for general system control, thread 1 assigned to execute 2D (two dimensional) activities, and threads 2 to 7 assigned to executing 3D activities (such as calculating vertices, primitives or rastering).

RX-0416 at 11:16-25; *see* RIB at 135. Respondents explain, “3D activities would correspond to graphics processing tasks, and at least the general system control would be considered a general

[REDACTED]

compute task[.]” *Id.* (citing, Hr’g Tr. (Pedram) at 976:4-15, 902:2-903:6, 907:8-908:2, 969:19-970:8).

No party asserts that threads “executing 3D activities” would not be a “graphics-processing task” and Dr. Pedram’s testimony—that Stuttard’s disclosed “general system control” would be understood as the claimed “general-compute task”—is persuasive. Despite AMD’s characterizations, its expert, Dr. Hart, fails to meaningfully dispute this association:

Q. Let’s shift to talking about Claim 19, slide 125. What does Stuttard describe about general compute tasks?

A. It really doesn’t discuss general compute tasks. It describes a graphics processor. So it focuses on graphics processing tasks.

Q. All right. What is your opinion regarding the validity of Claims 15 through 20 of the ’381 Patent? . . .

CIB at 132 (citing Hr’g Tr. (Hart) at 1287:14-18).

Accordingly, a skilled artisan would have recognized that the additional limitation of claim 19 is disclosed in Stuttard. Nevertheless, claim 19 is not rendered invalid in light of its dependence on claim 15.

e. Claim 20

For claim 20, Respondents contend, “Stuttard expressly discloses selecting and assigning (*i.e.*, receiving) threads associated with tasks based on ‘priorities relative to one another.’” RIB at 133 (citing Hr’g Tr. (Pedram) at 976:25-977:15; RX-0416 at 11:16-25). They further reason:

[U]nder the Court’s construction of the term “task” as “a unit of processing work or portion of an application to be executed” and based on the disclosure of the ’381 Patent, a POSITA would understand that a task would comprise one or more threads and that a high-priority task would give rise to high-priority threads. Conversely, once all high-priority threads associated with a task are completed based on their assigned priority levels, the task has received a high priority processing by the computing system.

[REDACTED]

Id. AMD counters in that, “Dr. Pedram has not explained how such priorities associated with threads would relate to the reception of tasks based on a priority associated with the tasks by the alleged shader core.” CIB at 132 (citing RX-0416 at 11:16-25).

Respondents have met their burden. Stuttard teaches “[t]he threads are assigned priorities relative to one another” in a designated discussion of thread manager 102. *See* RX-0416 at 10:5-12:22. There is no dispute that threads are assembled and stored in thread manager 102 before being passed to the processing blocks 106, which are Respondents’ alleged “shader core.” RIB at 128 (citing Hr’g Tr. (Pedram) at 967:25-968:9). This is a convincing implication that threads, being given a certain relative priority to one another within the thread manager 102, will be passed to the processing blocks 106 based on that priority. Indeed, additional supporting details of this priority-based decision-making are disclosed elsewhere in the reference. RX-0416 at 23:3-26:10 (section header “Determining which thread should be active”).

Accordingly, Respondents have shown the additional limitation of claim 20 is disclosed in Stuttard, but claim 20 is not rendered invalid in light of its dependence on claim 15.

3. Patentable Subject Matter

Respondents contend “[t]he Asserted Claims of the ’381 Patent are invalid under 35 U.S.C. § 101 because they are directed to the abstract idea of task scheduling and impart no particular technological improvement to a computer or other technology.” RIB at 138. Respondents present their theory under the two-step analysis created by *Alice*: 1) “whether the patent claims are directed toward a patent-ineligible concept,” and 2) whether the claim’s elements, considered both individually and as an ordered combination, transform the nature of the claim into a patent-eligible application.” RIB at 138 (citing *Alice*, 573 U.S. at 217-18).

As for step one, Respondents state, “[s]cheduling tasks has been done by humans since humans first existed and is an abstract idea. The Asserted Claims of the ’381 Patent merely recite

[REDACTED]

the receipt of several ‘tasks’ and then sending two or more of these tasks to a computing unit to be executed.” RIB at 138-139. Respondents view this as “inherently abstract.” *Id.* at 139 (citing *Braemar Mfg., LLC v. ScottCare Corp.*, 816 F. App’x 465, 468 (Fed. Cir. 2020); *Interval Licensing LLC v. AOL, Inc.*, 896 F.3d 1335, 1338 (Fed. Cir. 2018)). Respondents find support in the specification of the 381 patent:

The specification of the ’381 Patent confirms the abstract nature of the claims. The specification identifies task scheduling as conventional and long-employed by processors. Tr. (Pedram) at 927:2-931:1; RDX-0004C.73-76; RX-0149.00010 at 1:15-2:59. In describing the history of the task scheduling, the ’381 Patent described the use of processors in the prior art including CPUs and GPUs, which may be used separately or together. *See* JX-0003.0009 at 1:23-32. Prior to the ’381 Patent, GPUs could already execute different threads or shader programs, “typically associated with a set of state data” that is locally stored in the GPU. *Id.* at 1:33-57. The GPUs had “an array of processing elements, called a shader core,” which could be organized into multiple SIMD devices capable of executing different threads, shader programs, or “tasks” at the same time, thereby performing operations in parallel. *Id.* at 1:58-2:7. Additionally, at that time, the CPUs could hand off “tasks” to the GPU, and “an operating-system (OS) scheduler [that] stores the tasks,” which might be processed by the GPU in the order they were received or out of order. *Id.* at 2:4-19. GPUs in the prior art were also able to recognize priorities of tasks, and could “perform the more-important (*e.g.*, low-latency) task before the first task has finished completing in the shader core.” *Id.* at 2:20-37. The fact that tasks would be scheduled based on their associated “state data” was routine and conventional – any tasks executed in a processor must have state associated with them. *Id.* at 1:53-57, 2:20-37. Nothing in the ’381 patent changes the fundamental fact that scheduling tasks is just determining which task to pass on followed by another. Few claimed inventions are as abstract as that, which is why the Federal Circuit routinely finds such “scheduling” claims abstract. *See, e.g., PersonalWeb Techs. LLC v. Google LLC*, 8 F.4th 1310, 1317 (Fed. Cir. 2021) (reiterating that processes which “can be performed in the human mind, or by a human using a pen and paper, ... [has been] repeatedly found unpatentable”) (citation omitted); *Solutran, Inc. v. Elavon, Inc.*, 931 F.3d 1161, 1166 (Fed. Cir. 2019) (“claims performed in a particular order or sequence are directed to abstract ideas”).

RIB at 139-140; *see* RRB at 73-74 (citing *Customedia Tech., LLC v. Dish Network Corp.*, 951 F.3d 1359, 1364-65 (Fed. Cir. 2020)).

Respondents are not persuasive. Independent claim 15 is not directed to abstract, ineligible subject matter. Respondents oversimplify when they describe the aim of the claim as “scheduling

tasks . . . since humans first existed.” Rather, like the claim at issue in *Enfish*, claim 15 is directed to improving a computer’s ability to render graphics. Specifically, claim 15 recites the reception of multiple “tasks” from at least one “engine,” which was provided with tasks from an external processor, and loads “state data” associated with each of those tasks for the purpose of executing at least two tasks at the same time. 381 patent at cl. 15. According to the patent, and not questioned by Respondents, there are downsides to the act of “context switching” which the invention of claim 15 allegedly ameliorates. *See id.* at 2:8-59, 4:46-61 (“Executing two or more different streams of processing tasks provides the low-latency benefits of context switching without the overhead associated with stopping and draining the processing unit of data.”). Indeed, the prior art references asserted by Respondents tout similar computer-improving benefits, as does ARM documentation for the Accused Products. *See, e.g.*, RX-0416 at 1:24-2:13, 12:1-7 (discussing the reduction in need of context switching); RX-0423 at 1:59-2:53, 3:31-47 (“The updating or replacement of constants through implementation of the various CB system embodiments enables a smooth procedure of constant replacement in shader execution blocks, which may result in an improvement in computational efficiency.”); CX-0773C.26 (describing ARM job slots, “Job slots can be assigned to run in different address spaces if necessary. This flexibility enables GPU sharing between different processes on the application processor”), 31 (describing execution core, “[a] quad can continue execution until its operation depends on the result from one of the accelerators that is attached to the message fabric. This ability allows memory latency to be hidden.”).

These circumstances are markedly similar to those in *Enfish* and dissimilar to those in Respondents’ cited cases. In *Enfish*, the technology was a “self-referential” database and involved a claim reciting “means for configuring said memory to a logical table” with the table including

“logical rows” and “logical columns,” and using “object identification number[s]” to identify each row/column. *See* 822 F.3d at 1329. In discussing precedent, the court held:

We do not read *Alice* to broadly hold that all improvements in computer-related technology are inherently abstract and, therefore, must be considered at step two. Indeed, some improvements in computer-related technology when appropriately claimed are undoubtedly not abstract, such as a chip architecture, an LED display, and the like. Nor do we think that claims directed to software, as opposed to hardware, are inherently abstract and therefore only properly analyzed at the second step of the *Alice* analysis. Software can make non-abstract improvements to computer technology just as hardware improvements can, and sometimes the improvements can be accomplished through either route. We thus see no reason to conclude that all claims directed to improvements in computer-related technology, including those directed to software, are abstract and necessarily analyzed at the second step of *Alice*, nor do we believe that *Alice* so directs. Therefore, we find it relevant to ask whether the claims are directed to an improvement to computer functionality versus being directed to an abstract idea, even at the first step of the *Alice* analysis.

Id. at 1335. The court viewed the claims at issue as “directed to a specific improvement to computer functionality” and remarked “[m]uch of the advancement made in computer technology consists of improvements to software that, by their very nature, may not be defined by particular physical features but rather by logical structures and processes.” *Id.* at 1339. The court concluded:

In sum, the self-referential table recited in the claims on appeal is a specific type of data structure designed to improve the way a computer stores and retrieves data in memory. The specification's disparagement of conventional data structures, combined with language describing the “present invention” as including the features that make up a self-referential table, confirm that our characterization of the “invention” for purposes of the § 101 analysis has not been deceived by the “draftsman's art.” Cf. *Alice*, 134 S.Ct. at 2360. In other words, we are not faced with a situation where general-purpose computer components are added post-hoc to a fundamental economic practice or mathematical equation. Rather, the claims are directed to a specific implementation of a solution to a problem in the software arts. Accordingly, we find the claims at issue are not directed to an abstract idea.

Id. Claim 15 of the 381 patent is most certainly a “solution to a problem in the software arts” that “improve[s] the way a computer stores and retrieves data in memory” for computer graphics rendering.

[REDACTED]

On the other hand, *Braemar* involved claims which invoked generic computer functions for the purpose of cardiac biological signal monitoring. 816 Fed. Appx. at 467-68; *see* RIB at 139. The court held that this was “the abstract idea of classification and filtering data, not an improvement in the functioning of computer capabilities.” *Id.* at 470. In *Interval Licensing*, the claim again invoked computer functions, but for an “attention manager” which was an automated manner of using unused screen space for advertisement. *See* 896 F.3d at 1338-39; RIB at 139. The court summarized, “Claims 15-18 are directed to an abstract idea: the presentation of two sets of information, in a non-overlapping way, on a display screen” (*Interval Licensing*, 896 F.3d at 1338) and held, “[r]ecitation, as in this case, of the collection, organization, and display of two sets of information on a generic display device is abstract absent a ‘specific improvement to the way computers [or other technologies] operate’” (*id.* at 1345 (citing *Enfish*, 822 F.3d at 1336)). The claim in *PersonalWeb* also had no connection to improving computer functionality but instead used a computer for data access control (8 F.4th at 1313; *see* RIB at 140), nor did the claim in *Solutran* which was a method for electronically processing paper checks (931 F.3d at 1164; *see* RIB at 140). In fact, the court in *Solutran* explicitly distinguished its claim from *Enfish*, stating, “[t]his is not a situation where the claims ‘are directed to a specific improvement to the way computers operate’ and therefore not directed to an abstract idea, as in cases such as [*Enfish*].” *Id.* at 1167. Lastly, *Customedia*’s claims were also not directed to improving computers but using them for targeted advertising. 951 F.3d at 1361; RRB at 74. There is no parity between any of Respondents’ cases and claim 15 of the 381 patent.

Accordingly, Respondents have not shown claim 15, or any of its dependent claims, of the 381 patent is directed to ineligible subject matter under *Alice*. There is no need to consider *Alice* step two. Invalidity under 35 U.S.C. § 101 has not been shown.

VI. DOMESTIC INDUSTRY - ECONOMIC PRONG

In a patent-based complaint, a violation of Section 337 can be found “only if an industry in the United States, relating to the articles protected by the patent ... concerned, exists or is in the process of being established.” 19 U.S.C. § 1337(a)(2). Under Commission precedent, this “domestic industry requirement” of Section 337 consists of an economic prong and a technical prong. *Stringed Instruments*, Inv. No. 337-TA-586, Comm’n Op. at 12-14. The complainant bears the burden of establishing that the domestic industry requirement is satisfied. *See Certain Set-Top Boxes and Components Thereof*, Inv. No. 337-TA-454, Initial Determination at 294 (June 21, 2002) (not reviewed in relevant part).

The economic prong of the domestic industry requirement is defined in subsection (a)(3) of Section 337 as follows:

(3) For purposes of paragraph (2), an industry in the United States shall be considered to exist if there is in the United States, with respect to the articles protected by the patent, copyright, trademark or mask work concerned --

(A) Significant investment in plant and equipment;

(B) Significant employment of labor or capital; or

(C) Substantial investment in its exploitation, including engineering, research and development, or licensing.

19 U.S.C. § 1337(a)(3). The economic prong of the domestic industry requirement is satisfied by meeting the criteria of any one of the three factors listed above. Importantly, the Commission has clarified that investments in plant and equipment, labor, and capital that may fairly be considered investments in research and development are eligible for consideration under subsections (A) and (B), in addition to subsection (C). *See Certain Solid State Storage Drives, Stacked Electronics Components, and Products Containing Same*, Inv. No. 337-TA-1097, Comm’n Op. at 14 (June 29, 2018) (“*Solid State Storage*”).

[REDACTED]

AMD had previously moved for summary determination that the economic prong was met, without opposition from Respondents. As noted above, the motion was granted-in-part on February 6, 2023 via initial determination. Order No. 62. On March 15, 2023, the Commission determined to review and, upon review, affirmed Order No. 62. EDIS Doc. ID 792553. No further economic prong argument was presented by the parties. *See* CIB at 142.

Accordingly, for the purposes of this investigation, AMD has satisfied the economic prong of the domestic industry requirement for the 053 and 381 patents under subsections (A) and (B).

VII. CONCLUSIONS OF LAW

1. AMD has not proven that any Accused Product infringes any asserted claim of U.S. Patent No. 7,742,053.
2. AMD has proven that the Accused Products infringe claims 15-17, 19, and 20 of U.S. Patent No. 8,854,381.
3. Respondents have not proven the invalidity of any claim of U.S. Patent No. 7,742,053.
4. Respondents have proven the invalidity of claims 15-17 of U.S. Patent No. 8,854,381.
5. AMD has proven the existence of articles protected by claims 5, 6, and 9 of U.S. Patent No. 7,742,053.
6. AMD has proven the existence of articles protected by claims 15-17, 19, and 20 of U.S. Patent No. 8,854,381, except that claims 15-17 are invalid.
7. AMD has proven the existence of an economic domestic industry as required by 19 U.S.C. § 1337(a)(2) for U.S. Patent No. 7,742,053 and U.S. Patent No. 8,854,381.
8. There is no violation of section 337 with respect to U.S. Patent No. 7,742,053.
9. There is a violation of section 337 with respect to U.S. Patent No. 8,854,381.

VIII. RECOMMENDED DETERMINATION ON REMEDY AND BOND

The Commission's Rules provide that subsequent to an initial determination on the question of violation of section 337 of the Tariff Act of 1930, as amended, 19 U.S.C. § 1337, the administrative law judge shall issue a recommended determination concerning the appropriate remedy in the event that the Commission finds a violation of section 337, and the amount of bond to be posted by respondent during Presidential review of the Commission action under section 337(j). *See* 19 C.F.R. § 210.42(a)(1)(ii).

The Commission has broad discretion in selecting the form, scope, and extent of the remedy in a section 337 proceeding. *Viscofan, S.A. v. Int'l Trade Comm'n*, 787 F.2d 544, 548 (Fed. Cir. 1986). Under Section 337(d)(1), if the Commission determines as a result of an investigation that there is a violation of section 337, the Commission is authorized to enter either a limited or a general exclusion order. 19 U.S.C. § 1337(d)(1). A limited exclusion order instructs the U.S. Customs and Border Protection ("CBP") to exclude from entry all articles that are covered by the patent at issue and that originate from a named respondent in the investigation. A general exclusion order instructs the CBP to exclude from entry all articles that are covered by the patent at issue, without regard to source. *Certain Purple Protective Gloves*, Inv. No. 337-TA-500, Comm'n Op. at 5 (Dec. 22, 2004). Under section 337(f)(1), the Commission may issue a cease and desist order in addition to, or instead of, an exclusion order. 19 U.S.C. § 1337(f)(1). The Commission generally issues a cease and desist order directed to a domestic respondent when there is a "commercially significant" amount of infringing, imported product in the United States that could be sold, thereby undercutting the remedy provided by an exclusion order. *See Certain Crystalline Cefadroxil Monohydrate*, Inv. No. 337-TA-293, USITC Pub. 2391, Comm'n Op. on Remedy, the Public Interest and Bonding at 37-42 (June 1991); *Certain Condensers, Parts Thereof*

[REDACTED]

and Prods. Containing Same, Including Air Conditioners for Automobiles, Inv. No. 337-TA-334 (Remand), Comm’n Op. at 26-28 (Sept. 10, 1997).

Additionally, during the 60-day period of Presidential review under 19 U.S.C. § 1337(j), “articles directed to be excluded from entry under subsection (d) . . . shall . . . be entitled to entry under bond prescribed by the Secretary in an amount determined by the Commission to be sufficient to protect the complainant from any injury.” *See* 19 U.S.C. § 1337(j)(3). “The Commission typically sets the bond based on the price differential between the imported infringing product and the domestic industry article or based on a reasonable royalty. However, where the available pricing or royalty information is inadequate, the bond may be set at one hundred (100) percent of the entered value of the infringing product.” *Certain Industrial Automation Systems and Components Thereof Including Control Systems, Controllers, Visualization Hardware, Motion and Motor Control Systems, Networking Equipment, Safety Devices, and Power Supplies*, Inv. No. 337-TA-1074, Comm’n Op. at 13 (Apr. 23, 2019) (“*Automation Systems*”) (public version) (citation omitted).

A. Limited Exclusion Order

Should a violation be found, there is no dispute that a limited exclusion order (“LEO”) should issue against each respondent. *See* CIB at 142 (“Realtek and TCL do not dispute that an LEO is appropriate if a violation of Section 337 is found”); RIB at 145; RRB at 74. And while there is some dispute whether the LEOs should include a certification provision (*see* CIB at 143; RIB at 145; CRB at 73-74; RRB at 74), certification provisions are known to “aid U.S. Customs and Border Protection (‘CBP’) in enforcing Commission orders but ‘do not mandate that CBP accept certification as proof that the articles in question are not covered’ by the limited exclusion order” (*Certain Robotic Vacuum Cleaning Devices and Components Thereof Such as Spare Parts*,

[REDACTED]

Inv. No. 337-TA-1057, Comm’n Op. at 55 (Feb. 1, 2019)). As “it has been Commission practice for the past several years to include certification provisions in its exclusion orders to aid CBP” (see *Certain Road Milling Machines and Components Thereof*, Inv. No. 337-TA-1067, Comm’n Op. at 15, 15 n. 5 (Aug. 7, 2019) (citations omitted)), it is my recommendation that limited exclusion orders should issue against each respondent and include the Commission’s standard certification provision.

Respondents further request, in the event of a violation, suspension of any remedial enforcement orders “until the conclusion of any IPRs regarding [the] patents.” RIB at 144. This is not recommended. Respondents have not identified which IPRs, if any, have been filed against the Asserted Patents, let alone the state of their progress towards a final written decision. For its part, AMD contends any IPR final decision would not issue before this investigation’s target date (CIB at 142), and so, overall, there is very little reason to suspend any remedial order on this basis.

Respondents also seek “exceptions for the importation of products or components necessary to service and repair any infringing products that are already in the possession of customers or to otherwise comply with respondents’ warranty obligations.” RIB at 145. The request is noted but turns on public interest concerns which were not instituted for this stage of the investigation. See 87 Fed. Reg. 34719 (June 7, 2022)); RRB at 74 (discussing harm to “innocent[]” prior purchasers of Accused Products).

B. Cease and Desist Order

Should a violation be found, AMD argues cease and desist orders (“CDOs”) should issue against all Respondents. CIB at 143-144. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Complainants bear the burden on the issue of cease and desist orders. *Certain Microfluidic Devices*, Inv. No. 337-TA-1068, Comm’n Op. at 23 (Jan. 10, 2020). Such orders “are generally issued when, with respect to the imported infringing products, respondents maintain commercially significant inventories in the United States or have significant domestic operations that could undercut the remedy provided by an exclusion order.” *Id.* at 22-23 (citations omitted). [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] Overall, AMD has not shown a CDO against Realtek is justified, and so one is not recommended to issue in the event the Commission finds a violation.

C. Bond

The Commission has held that “[t]he complainant bears the burden of establishing the need for a bond” during the Presidential Review period. *See Robotic Vacuums*, Inv. No. 337-TA-1057, Comm’n Op. at 68. The amount of the bond is, generally, set “to be sufficient to protect the complainant from any injury.” 19 U.S.C. § 1337(e)(1), (j3). AMD explains that it and the TCL Respondents stipulated to a zero bond, and it has otherwise “withdrawn its request for a bond amount on Realtek Products found in violation.” CIB at 150 (citing JX-0013C; Hr’g Tr. at 17:7-25). It is therefore recommended no bond should issue in the event of a violation.

IX. INITIAL DETERMINATION AND ORDER

Based on the foregoing,⁴ it is my Initial Determination that there is a violation of Section 337 of the Tariff Act of 1930, as amended, 19 U.S.C. § 1337, in the importation into the United States, the sale for importation, or the sale within the United States after importation of certain

⁴ The failure to discuss any matter raised by the parties or any portion of the Record herein does not indicate that said matter was not considered. Rather, any such matter(s) or portion(s) of the Record has/have been determined to be irrelevant, immaterial or meritless. Arguments made on brief which were otherwise unsupported by Record evidence or legal precedent have been accorded no weight.

[REDACTED]

graphics systems, components thereof, and digital televisions containing the same in connection with the asserted claims of U.S. Patent No. 8,854,381. There is no violation in connection with the asserted claims of U.S. Patent No. 7,742,053.

The undersigned hereby certifies to the Commission this Initial Determination, together with the Record of the hearing in this investigation consisting of the following: the transcript of the evidentiary hearing, with appropriate corrections as may hereafter be ordered; and the exhibits accepted into evidence in this investigation.⁵

Pursuant to 19 C.F.R. § 210.42(h), this Initial Determination shall become the determination of the Commission sixty (60) days after the date of service of the Initial Determination, unless a party files a petition for review of the Initial Determination within twelve (12) business days after service of the Initial Determination pursuant to 19 C.F.R. § 210.43(a) or the Commission, pursuant to 19 C.F.R. § 210.44, orders on its own motion, a review of the Initial Determination or certain issues therein. Any issue or argument not raised in a petition for review, or response thereto, will be deemed to have been abandoned and may be disregarded by the Commission in reviewing the Initial Determination pursuant to 19 C.F.R. § 210.43(b) and (c).

⁵ The pleadings of the parties filed with the Secretary need not be certified as they are already in the Commission's possession in accordance with Commission rules.

[REDACTED]

Confidentiality Notice:

This Initial Determination is being issued as confidential, and a public version will be issued pursuant to Commission Rule 210.5(f). Within seven (7) days of the date of this Initial Determination, the parties shall jointly submit: (1) a proposed public version of this opinion with any proposed redactions bracketed in red; and (2) a written justification for any proposed redactions specifically explaining why the piece of information sought to be redacted is confidential and why disclosure of the information would be likely to cause substantial harm or likely to have the effect of impairing the Commission's ability to obtain such information as is necessary to perform its statutory functions.⁶

SO ORDERED.



Cameron Elliot
Administrative Law Judge

⁶ Under Commission Rules 210.5 and 201.6(a), confidential business information includes: information which concerns or relates to the trade secrets, processes, operations, style of works, or apparatus, or to the production, sales, shipments, purchases, transfers, identification of customers, inventories, or amount or source of any income, profits, losses, or expenditures of any person, firm, partnership, corporation, or other organization, or other information of commercial value, the disclosure of which is likely to have the effect of either impairing the Commission's ability to obtain such information as is necessary to perform its statutory functions, or causing substantial harm to the competitive position of the person, firm, partnership, corporation, or other organization from which the information was obtained, unless the Commission is required by law to disclose such information. *See* 19 C.F.R. § 201.6(a). Thus, to constitute confidential business information the disclosure of the information sought to be designated confidential must likely have the effect of either: (1) impairing the Commission's ability to obtain such information as is necessary to perform its statutory functions; or (2) causing substantial harm to the competitive position of the person, firm, partnership, corporation, or other organization from which the information was obtained.

I, Lisa R. Barton, hereby certify that the attached **INITIAL DETERMINATION** has been served upon the following parties as indicated, on **July 7, 2023**.



Lisa R. Barton, Secretary
U.S. International Trade Commission
500 E Street, SW, Room 112
Washington, DC 20436

**On Behalf of Complainants Advanced Micro Devices, Inc.
and ATI Technologies ULC:**

Adam S. Rizk, Esq.
**MINTZ, LEVIN, COHN, FERRIS, GLOVSKY, AND
POPEO PC**
One Financial Center
Boston, MA 02111
Email: ARizk@mintz.com

- ☐ Via Hand Delivery
- ☐ Via Express Delivery
- ☐ Via First Class Mail
- ☒ Other: Email Notification of
Availability for Download

On Behalf of Respondent Realtek Semiconductor Corp.:

Robert J. Benson, Esq.
ORRICK, HERRINGTON & SUTCLIFFE LLP
2050 Main Street
Suite 1100
Irvine, CA 92614
Email: rbenson@orrick.com

- ☐ Via Hand Delivery
- ☐ Via Express Delivery
- ☐ Via First Class Mail
- ☒ Other: Email Notification of
Availability for Download

**On Behalf of Respondents TCL Industries Holdings Co.,
Ltd., TCL Industries Holdings (H.K.) Limited, TCL
Electronics Holdings Limited f/k/a TCL Multimedia
Technology Holdings, Ltd., TCL Technology Group
Corporation, TTE Corporation, TCL Holdings (BVI)
Limited, TCL King Electrical Appliances (Huizhou) Co.,
Ltd., Shenzhen TCL New Technologies Co., Ltd.,
TCL MOKA International Limited, TCL Smart Device
(Vietnam) Co., Ltd., Manufacturas Avanzadas SA de CV,**

**TCL Electronics Mexico, S de RL de CV, and TCL
Overseas Marketing
Ltd.:**

John P. Schnurer, Esq.
PERKINS COIE LLP
11452 El Camino Real, Suite 300
San Diego, CA 92130-2080
Email: JSchnurer@perkinscoie.com

- ☐ Via Hand Delivery
- ☐ Via Express Delivery
- ☐ Via First Class Mail
- ☒ Other: Email Notification of
Availability for Download

On Behalf of Intervenor TTE Technology, Inc.

John P. Schnurer, Esq.
PERKINS COIE LLP
11452 El Camino Real, Suite 300
San Diego, CA 92130-2080
Email: JSchnurer@perkinscoie.com

- ☐ Via Hand Delivery
- ☐ Via Express Delivery
- ☐ Via First Class Mail
- ☒ Other: Email Notification of
Availability for Download