

In the Matter of

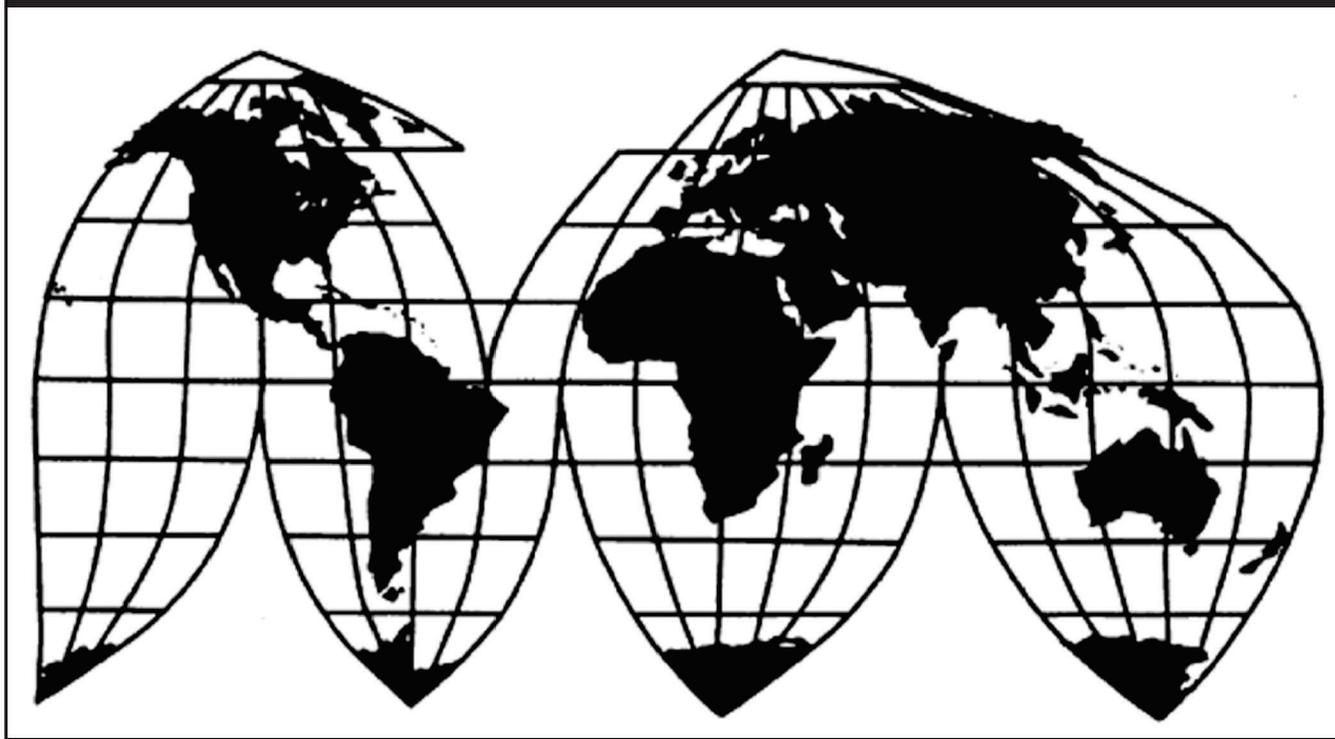
**CERTAIN CONSUMER ELECTRONICS AND
DISPLAY DEVICES WITH GRAPHICS
PROCESSING AND GRAPHICS
PROCESSING UNITS THEREIN**

337-TA-932

Publication 4903

June 2019

U.S. International Trade Commission



Washington, DC 20436

U.S. International Trade Commission

COMMISSIONERS

Meredith Broadbent, Chairman
Dean Pinkert, Vice Chairman
Irving Williamson, Commissioner
David Johanson, Commissioner
Scott Kieff, Commissioner
Rhonda Schmidlein, Commissioner

**Address all communications to
Secretary to the Commission
United States International Trade Commission
Washington, DC 20436**

U.S. International Trade Commission

Washington, DC 20436
www.usitc.gov

In the Matter of

CERTAIN CONSUMER ELECTRONICS AND DISPLAY DEVICES WITH GRAPHICS PROCESSING AND GRAPHICS PROCESSING UNITS THEREIN

337-TA-932



UNITED STATES INTERNATIONAL TRADE COMMISSION
Washington, D.C.

In the Matter of

CERTAIN CONSUMER ELECTRONICS
AND DISPLAY DEVICES WITH
GRAPHICS PROCESSING AND
GRAPHICS PROCESSING UNITS
THEREIN

Investigation No. 337-TA-932

**NOTICE OF COMMISSION DECISION NOT TO REVIEW THE ALJ'S FINAL INITIAL
DETERMINATION FINDING NO VIOLATION OF SECTION 337; TERMINATION OF
INVESTIGATION**

AGENCY: U.S. International Trade Commission.

ACTION: Notice.

SUMMARY: Notice is hereby given that the U.S. International Trade Commission has determined not to review the final initial determination (ID) issued on October 9, 2015, which found no violation of section 337 of the Tariff Act of 1930, 19 U.S.C. § 1337, in this investigation.

FOR FURTHER INFORMATION CONTACT: Ron Traud, Office of the General Counsel, U.S. International Trade Commission, 500 E Street, S.W., Washington, D.C. 20436, telephone (202) 205-3427. Copies of non-confidential documents filed in connection with this investigation are or will be available for inspection during official business hours (8:45 a.m. to 5:15 p.m.) in the Office of the Secretary, U.S. International Trade Commission, 500 E Street, S.W., Washington, D.C. 20436, telephone (202) 205-2000. General information concerning the Commission may also be obtained by accessing its Internet server at <http://www.usitc.gov>. The public record for this investigation may be viewed on the Commission's electronic docket (EDIS) at <http://edis.usitc.gov>. Hearing-impaired persons are advised that information on this matter can be obtained by contacting the Commission's TDD terminal on (202) 205-1810.

SUPPLEMENTARY INFORMATION: The Commission instituted this investigation based on a complaint filed by NVIDIA Corporation of Santa Clara, California (NVIDIA). The investigation was 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 certain consumer electronics and display devices with graphics processing and graphics processing units therein by reason of infringement of one or more of

The authority for the Commission's determination is contained in section 337 of the Tariff Act of 1930, as amended (19 U.S.C. § 1337), and in sections 210.42-46 of the Commission's Rules of Practice and Procedure (19 C.F.R. §§ 210.42-46).

By order of the Commission.

A handwritten signature in black ink, appearing to read 'Lisa R. Barton', written in a cursive style.

Lisa R. Barton
Secretary to the Commission

Issued: December 14, 2015

**CERTAIN CONSUMER ELECTRONICS AND DISPLAY
DEVICES WITH GRAPHICS PROCESSING AND
GRAPHICS PROCESSING UNITS THEREIN**

Inv. No. 337-TA-932

PUBLIC CERTIFICATE OF SERVICE

I, Lisa R. Barton, hereby certify that the attached **NOTICE** has been served by hand upon the Commission Investigative Attorney, R. Whitney Winston, Esq., and the following parties as indicated, on **December 14, 2015**.



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

On Behalf of Complainant NVIDIA Corporation:

I. Neel Chatterjee
ORRICK, HERRINGTON & SUTCLIFFE LLP
1000 Marsh Road
Menlo Park, California 94025

- Via Hand Delivery
 Via Express Delivery
 Via First Class Mail
 Other: _____

On Behalf of Respondent Qualcomm, Inc.:

Deanna Tanner Okun, Esq.
ADDUCI, MASTRIANI & SCHAUMBERG LLP
1133 Connecticut Ave., NW, 12th Floor
Washington, DC 20036

- Via Hand Delivery
 Via Express Delivery
 Via First Class Mail
 Other: _____

**On Behalf of Samsung Electronics Co., Ltd., Samsung
Telecommunications America, LLC, and Samsung
Semiconductor, Inc.:**

D. Sean Trainor, Esq.
KIRKLAND & ELLIS LLP
655 Fifteenth Street, NW
Washington, DC 20005

- Via Hand Delivery
 Via Express Delivery
 Via First Class Mail
 Other: _____

PUBLIC VERSION

**UNITED STATES INTERNATIONAL TRADE COMMISSION
Washington, D.C.**

In the Matter of

**CERTAIN CONSUMER ELECTRONICS AND DISPLAY
DEVICES WITH GRAPHICS PROCESSING AND
GRAPHICS PROCESSING UNITS THEREIN**

Inv. No. 337-TA-932

INITIAL DETERMINATION ON VIOLATION OF SECTION 337

Administrative Law Judge Thomas B. Pender

(October 09, 2015)

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 Consumer Electronics And Display Devices With Graphics Processing And Graphics Processing Units Therein*, Investigation No. 337-TA-932.

PUBLIC VERSION

Table of Contents

I. INTRODUCTION 5

 A. The Parties 3

 1. Complainant 3

 2. Respondents 3

 B. Products at Issue 4

II. JURISDICTION 6

 A. Subject Matter Jurisdiction 6

 B. Personal Jurisdiction 6

 C. In Rem Jurisdiction 7

III. RELEVANT LAW 7

 A. Infringement 7

 1. Claim Construction 7

 2. Direct Infringement 8

 a. Literal Infringement 8

 b. Doctrine of Equivalents 9

 3. Indirect Infringement 9

 B. Invalidity 10

 1. Anticipation 11

 2. Obviousness 11

 3. Written Description and Enablement 13

 C. Domestic Industry – Technical Prong 14

IV. U.S. PATENT NO. 7,209,140 15

 A. Level of Ordinary Skill in the Art 17

 B. Claim Construction 17

 1. Order No. 20: Construing Terms of the Asserted Patents 17

 2. Agreed Construction - “programmable by the user” 17

 3. Disputed Constructions 18

 a. “graphics data”/“operations on the graphics data” 18

 b. “hardware graphics accelerator” 23

 c. “reciprocal instruction” 29

 C. Infringement 34

 D. Domestic Industry – Technical Prong 37

 E. Invalidity 38

PUBLIC VERSION

1.	Anticipation – 35 U.S.C. 102.....	38
2.	Obviousness – 35 U.S.C. 103	43
a.	Renderman on Horizon860	43
b.	“C” Language on TMS34082	45
3.	Inventorship – 35 U.S.C. 102(f).....	52
V.	U.S. PATENT NO. 6,690,372	59
A.	Level of Ordinary Skill in the Art.....	60
B.	Claim Construction	60
1.	Order No. 20: Construing Terms of the Patent - “shading calculation” ...	60
2.	Disputed Term - “texture look-up module coupled to the shading module” (Claim 23)	61
C.	Infringement.....	71
1.	Claim 23.....	71
a.	Sufficiency of NVIDIA’s infringement evidence.....	72
b.	Do the Accused Products have a “shading module” and a “texture look-up module” coupled to a “shading module”?	74
c.	Conclusion	81
2.	Claim 24.....	82
D.	Domestic Industry – Technical Prong.....	86
E.	Invalidity	87
1.	Anticipation.....	87
a.	“McCool”	87
b.	U.S. Patent No. 6,236,413.....	101
c.	“Ackerman”	108
2.	Obviousness	114
VI.	U.S. PATENT NO. 7,038,685	114
A.	Level of Ordinary Skill in the Art.....	115
B.	Claim Construction	117
1.	Agreed construction - “dynamically controlling...”	117
2.	Disputed construction - “graphics processor”	117
C.	Infringement.....	123
1.	Claim 1	123
a.	Mali GPUs	123
b.	Adreno GPUs	132
2.	Claim 15.....	137

PUBLIC VERSION

a.	Mali GPUs	137
b.	Adreno GPUs	137
D.	Domestic Industry - Technical Prong	138
E.	Invalidity	138
1.	Anticipation (35 U.S.C. § 102) – The ATI Unified Shader	138
a.	Conception	140
b.	Reduction to Practice	145
c.	Conclusion	149
2.	Obviousness – 35 U.S.C. 103	149
a.	Selzer in view of Van Hook	157
b.	Amanatides in view of Van Hook	160
VII.	DOMESTIC INDUSTRY	162
VIII.	CONCLUSIONS OF LAW	164
IX.	INITIAL DETERMINATION AND ORDER	165

PUBLIC VERSION

Table of Abbreviations

CDX	Complainant's Demonstrative Exhibit
CIB	Complainant's Initial Post-Hearing Brief
CRB	Complainant's Reply Post-Hearing Brief
CX	Complainant's Exhibit
Dep.	Deposition
JX	Joint Exhibit
RDX	Respondent's Demonstrative Exhibit
RIB	Respondent's Initial Post-Hearing Brief
RRB	Respondent's Reply Post-Hearing Brief
RX	Respondent's Exhibit
Tr.	Transcript
DWS	Direct Witness Statement (Including Revised Direct Witness Statements)
RWS	Rebuttal Witness Statement

I. INTRODUCTION

Complainant NVIDIA Corporation ("NVIDIA") filed a complaint on September 4, 2014 alleging a violation of Section 337 based on infringement of seven patents: (i) U.S. Patent No. 6,198,488 ("the '488 patent"); (ii) U.S. Patent No. 6,992,667 ("the '667 patent"); (iii) U.S. Patent No. 7,209,140 ("the '140 patent"); (iv) U.S. Patent No. 6,690,372 ("the '372 patent"); (v) U.S. Patent No. 7,038,685 ("the '685 patent"); (vi) U.S. Patent No. 7,015,913 ("the '913 patent"); and (vii) U.S. Patent No. 6,697,063 ("the '063 patent"). (Compl. ¶ 5). On October 10, 2014, the Commission instituted an investigation, by publication of a notice in the *Federal Register*, to determine:

PUBLIC VERSION

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 certain consumer electronics and display devices with graphics processing and graphics processing units therein by reason of infringement of one or more of claims 1, 19, and 20 of the '488 patent; claims 1-29 of the '667 patent; claims 1-5, 7-19, 21-23, 25-30, 34-36, 38, 41-43 of the '685 patent; claims 5-8, 10, 12-20 and 24-27 of the '913 patent; claims 7, 8, 11-13, 16-21, 23, 24, 28, and 29 of the '063 patent; claims 1-10, 12, and 14 of the '140 patent; and claims 1-6, 9-16, and 19-25 of the '372 patent, and whether an industry in the United States exists as required by subsection (a)(2) of section 337....

79 Fed. Reg. 61338 (Oct. 10, 2014).

A *Markman* hearing was held on February 2 and 3, 2015, and an order construing the disputed terms issued on April 2, 2015. (Order No. 20 (Apr. 2, 2015) ("*Markman* Order")). In addition, the Commission has determined that NVIDIA satisfies the economic prong of the domestic industry requirement with respect to each of the asserted patents. (Notice of Commission Determination not to Review an Initial Determination Granting Complainants' Unopposed Motion for Summary Determination that the Economic Prong of the Domestic Industry Requirement is Satisfied (Apr. 22, 2015).)

NVIDIA filed four motions for partial termination based on its withdrawal of certain allegations set forth in its complaint. On December 31, 2014, NVIDIA filed a motion seeking termination with respect to claim 19 of the '488 patent. (Mot. Docket Nos. 932-013). The motion was granted, and the investigation was terminated with respect to that claim on February 18, 2015. (*See* Order No. 10 (Jan. 15, 2015); Notice (Feb. 18, 2015)). On June 8, 2015, June 22, 2015, and June 25, 2015, NVIDIA filed motions seeking termination with respect to the '488 patent, the '667 patent, the '913 patent, the '063 patent, and claims 4, 16, 19, and 21 of the '685 patent. (Mot. Docket No. 932-047, 932-051, 932-060). These motions were granted and the investigation was terminated with respect to these claims and patents. (*See* Order No. 35.)

PUBLIC VERSION

The following patent assertions remain: claim 14 of the '140 patent; claims 23 and 24 of the '372 patent; and claims 1 and 15 of the '685 patent (respectively, the “Asserted Claims” and the “Asserted Patents”).

A. The Parties

1. Complainant

Complainant NVIDIA Corporation (“NVIDIA”) is a Delaware corporation having a principal place of business in Santa Clara, California. (Complaint at ¶ 12).

2. Respondents

Respondent Samsung Electronics Co., Ltd. is a corporation organized under the laws of South Korea having a principal place of business at Samsung Main Building, 250 Taepyung-ro 2-ka, Chung-ku, Seoul 100-742, Korea. (Samsung Answer at ¶ 25). Samsung Electronics Co., Ltd. manufactures the accused mobile phones and tablet computers at issue in this investigation. Respondent Samsung Telecommunications America, LLC is a Delaware limited liability company having a principal place of business at 1301 Lookout Drive, Richardson, Texas 75802. (*Id.* at ¶ 28.) Samsung Telecommunications America, LLC is involved in the importation and sale in the United States after importation of the accused mobile phones and tablet computers at issue in this investigation. (*Id.*) Respondent Samsung Semiconductor, Inc. is a California corporation having a principal place of business at 3655 North First Street, San Jose, California 95134. (*Id.* at ¶ 29.) Respondent Samsung Electronics America, Inc. is a New York corporation having a principal place of business at 85 Challenger Road, Ridgefield Park, New Jersey 07660. (*Id.* at ¶ 27.) Samsung Electronics America, Inc. is a wholly-owned subsidiary of Samsung Electronics Co., Ltd. (*Id.* at ¶ 27.) Samsung Telecommunications America, LLC and Samsung Semiconductor, Inc. are subsidiaries of Samsung Electronics America, Inc. (*Id.* at ¶¶ 28-29.) Samsung Electronics Co., Ltd., Samsung Telecommunications America, LLC, Samsung

PUBLIC VERSION

Semiconductor, Inc., and Samsung Electronics America, Inc. are referred to collectively herein as “Samsung.”

Respondent Qualcomm, Inc. (“Qualcomm”) is a Delaware corporation having a principal place of business at 5775 Morehouse Drive, San Diego, California 92121. (Qualcomm Answer at ¶ 23). Qualcomm designs, has-manufactured, and sells mobile processors with Adreno graphics processing units (“GPU”) that are used in certain accused Samsung devices. (*Id.*)

B. Products at Issue

The parties still disagree about what products are actually at issue in this investigation. (CRB at 3; RRB at 4-6.) It is black letter law that the scope of a Section 337 investigation is determined by the Commission’s Notice of Investigation (“NOI”). Thus, in determining what products are at issue in this investigation, I must turn to the NOI to see how the Commission defined the scope of this investigation.

The scope of this investigation is defined in the NOI as “certain consumer electronics and display devices with graphics processing and graphics processing units therein.” 79 Fed. Reg. 61338 (Oct. 10, 2014). NVIDIA argues that the NOI defines the scope as: (1) certain consumer electronics and display devices with graphics processing; and (2) graphics processing units therein. Respondents and the Staff maintains the scope of the investigation is certain consumer electronics and display devices with graphics processing and graphics processing units therein (*i.e.*, “with graphics processing and graphics processing units therein” is a single prepositional phrase that modifies “certain consumer electronics and display devices.”).

The syntax of the NOI only supports the position of the Respondents and the Staff. If the phrase “graphics processing units therein” was meant to be a separate clause from “certain consumer electronics and display devices with graphics processing” there would be a comma

PUBLIC VERSION

after "...with graphics processing." Without such a comma, the phrase "with graphics processing and graphics processing units therein" reads as a single prepositional phrase.

Moreover, NVIDIA's argument is belied by the following table from its own opening post-hearing brief.

Patent	Claim	Accused Products
6,690,372	23	Samsung consumer products with Adreno, Mali, or PowerVR GPUs
	24	Samsung consumer products with Mali GPUs
7,209,140	14	Samsung consumer products with Adreno, Mali, or PowerVR GPUs
7,038,685	1	Samsung consumer products with Adreno or Mali GPUs
	15	Samsung consumer products with Adreno or Mali GPUs

(CIB at 2.) The above table makes clear the Accused Products are "Samsung consumer products ..." NVIDIA does not identify any GPU, by itself, as an Accused Product and certainly Qualcomm's GPUs are not Samsung consumer products. Accordingly, for at least the reasons above, I find the scope of this investigation is limited to consumer electronics and display devices that include graphics processing capabilities and that have graphics processing units therein.

Consistent with the NOI, the Accused Products in this investigation are Samsung mobile devices, including phones and tablets, that employ one of the following five GPU architectures: (i) Qualcomm Adreno A3X (302, 305, 306, 320, 330); (ii) Qualcomm Adreno A4X (420, 430); (iii) ARM Mali T-6X (604, 624, 628); (iv) ARM Mali T-7X (720, 760); and (v) Imagination PowerVR SGX 5X (540, 544). (CIB at 2.)

For purposes of the domestic industry requirement, NVIDIA relies upon its GPUs with Kepler, Fermi, and Maxwell architectures (*i.e.*, GeForce Titan Z, GeForce GTX750 TI, and Tegra K1 processors) and products incorporating those GPUs. (CIB at 23, 63, 129.)

II. JURISDICTION

In order to have the power to decide a case, a court or agency must have both subject matter jurisdiction and jurisdiction over either the parties or the property involved. 19 U.S.C. § 1337; *Certain Steel Rod Treating Apparatus and Components Thereof*, Inv. No. 337-TA-97, Commission Memorandum Opinion, 215 U.S.P.Q. 229, 231 (1981).

A. Subject Matter Jurisdiction

Section 337 confers subject matter jurisdiction on the International Trade Commission to investigate, and if appropriate, to provide a remedy for, unfair acts and unfair methods of competition in the importation, the sale for importation, or the sale after importation of articles into the United States. (*See* 19 U.S.C. §§ 1337(a)(1)(B) and (a)(2).) NVIDIA alleges in the Complaint a violation of Subsection 337(a)(1)(B) in the importation and sale of products that infringe the Asserted Patents. (*See* Complaint.) NVIDIA has alleged sufficient facts that, if proven, would show Samsung and Qualcomm imported articles that infringe NVIDIA's patents. *See Certain Elec. Devices with Image Processing Sys., Components Thereof, & Assoc. Software*, Inv. No. 337-TA-724, Comm'n Op. at 8-10 (Dec. 2, 2011) (citing *Amgen, Inc. v. ITC*, 902 F.2d 1532, 1536 (Fed. Cir. 1990)). Accordingly, I find the Commission has subject matter jurisdiction over this Investigation under Section 337 of the Tariff Act of 1930

B. Personal Jurisdiction

Respondents Samsung and Qualcomm have appeared and fully participated in this investigation by, among other things, participating in discovery, participating in the evidentiary hearing, and filing pre-hearing and post-hearing briefs. Accordingly, I find Respondents Samsung and Qualcomm have submitted to the jurisdiction of the Commission. *See Certain Miniature Hacksaws*, Inv. No. 337-TA-237, Pub. No. 1948, Initial Determination at 4, 1986 WL 379287 (U.S.I.T.C., October 15, 1986) (unreviewed by Commission in relevant part).

PUBLIC VERSION

C. In Rem Jurisdiction

Samsung has stipulated to importation of the Accused Products and further stipulated it does not and will not dispute the importation requirement within the meaning of 19 USC § 1337(a)(1)(B) is satisfied in this Investigation. (See JX-020C.) See *Certain Kinesiotherapy Devices*, Inv. No. 337-TA-823, ID, at 11-12 (Jan. 8, 2013) (unreviewed in relevant part) citing *Amgen, Inc. v. U.S. Int'l Trade Comm'n*, 902 F.2d 1532, 1536 (Fed. Cir. 1990).

Qualcomm denies engaging in any importation-related activities with respect to the Accused Products, but does not contest that the Accused Products (*i.e.*, the Samsung mobile devices at issue in this investigation) have been imported, as stipulated to by Samsung. Thus, I find the Accused Products have been imported into the United States. Accordingly, the Commission has *in rem* jurisdiction over the Accused Products. See *Sealed Air Corp. v. United States Int'l Trade Comm'n*, 645 F.2d 976, 985 (C.C.P.A. 1981).

III. RELEVANT LAW

A. 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 v. Westview Instruments, Inc.*, 52 F.3d 967, 976 (Fed. Cir. 1995) (*en banc*) (internal citations omitted), *aff'd*, 517 U.S. 370 (1996).

1. Claim Construction

“Claim construction is generally a matter of law ... but it may have underlying factual determinations.” *Summit 6, LLC v. Samsung Electronics Co., Ltd.*, 2015 WL 5515331 (Fed. Cir. 2015) (citing *Teva Pharm. USA, Inc. v. Sandoz, Inc.*, —U.S.—, —, 135 S.Ct. 831, 837, — L.Ed.2d —, — (2015). “The construction of claims is simply a way of elaborating the

PUBLIC VERSION

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).

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*, 52 F.3d at 979. 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). 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.

2. Direct Infringement

A complainant must prove either literal infringement or infringement under the doctrine of equivalents. Infringement 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).

a. Literal Infringement

Literal infringement is a question of fact. *Finisar Corp. v. DirecTV Group, Inc.*, 523 F.3d 1323, 1332 (Fed. Cir. 2008). Literal infringement requires the patentee to prove that the accused device contains each and every limitation of the asserted claim(s). *Frank's Casing Crew &*

PUBLIC VERSION

Rental Tools, Inc. v. Weatherford Int'l, Inc., 389 F.3d 1370, 1378 (Fed. Cir. 2004). 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).

b. Doctrine of Equivalents

Where literal infringement is not found, infringement nevertheless can be found under the doctrine of equivalents. Determining infringement under the doctrine of equivalents “requires an intensely factual inquiry.” *Vehicular Techs. Corp. v. Titan Wheel Int'l, Inc.*, 212 F.3d 1377, 1381 (Fed. Cir. 2000). According to the Federal Circuit:

Infringement under the doctrine of equivalents may be found when the accused device contains an “insubstantial” change from the claimed invention. Whether equivalency exists may be determined based on the “insubstantial differences” test or based on the “triple identity” test, namely, whether the element of the accused device “performs substantially the same function in substantially the same way to obtain the same result.” The essential inquiry is whether “the accused product or process contain elements identical or equivalent to each claimed element of the patented invention[.]”

TIP Sys., LLC v. Phillips & Brooks/Gladwin, Inc., 529 F.3d 1364, 1376-77 (Fed. Cir. 2008)

(citations omitted). Thus, if an element is missing or not satisfied, infringement cannot be found under the doctrine of equivalents as a matter of law. *London v. Carson Pirie Scott & Co.*, 946 F.2d 1534, 1538-39 (Fed. Cir. 1991).

3. Indirect Infringement

Section 271(b) of the Patent Act prohibits inducement: “[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 aid[ed] and abett[ed] another’s direct infringement.”)

PUBLIC VERSION

Section 271(c) of the Patent Act prohibits contributory infringement: “Under 35 U.S.C. § 271(c), 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).

B. Invalidity

It is Respondents’ burden to prove invalidity, and the burden of proof never shifts to the patentee to prove validity. *Scanner Techs. Corp. v. ICOS Vision Sys. Corp. N.V.*, 528 F.3d 1365, 1380 (Fed. Cir. 2008). “Under the patent statutes, a patent enjoys a presumption of validity, *see* 35 U.S.C. § 282, which can be overcome only through facts supported by clear and convincing evidence[.]” *SRAM Corp. v. AD-II Eng’g, Inc.*, 465 F.3d 1351, 1357 (Fed. Cir. 2006).

The clear and convincing evidence standard placed on the party asserting the invalidity defense requires a level of proof beyond the preponderance of the evidence. Although not susceptible to precise definition, “clear and convincing” evidence has been described as evidence which produces in the mind of the trier of fact “an abiding conviction that the truth of a factual contention is ‘highly probable.’” *Price v. Symsek*, 988 F.2d 1187, 1191 (Fed. Cir. 1993) (citing *Buildex, Inc. v. Kason Indus., Inc.*, 849 F.2d 1461, 1463 (Fed. Cir. 1988).)

“When no prior art other than that which was considered by the PTO examiner is relied on by the attacker, he has the added burden of overcoming the deference that is due to a qualified government agency presumed to have properly done its job[.]” *Am. Hoist & Derrick Co. v. Sowa & Sons, Inc.*, 725 F.2d 1350, 1359 (Fed. Cir. 1984). Therefore, the challenger’s “burden is especially difficult when the prior art was before the PTO examiner during prosecution of the

PUBLIC VERSION

application.” *Hewlett-Packard Co. v. Bausch & Lomb Inc.*, 909 F.2d 1464, 1467 (Fed. Cir. 1990).

1. Anticipation

Under 35 U.S.C. § 102(a), a patent is invalid for anticipation if it was “patented, described in a printed publication, or in public use, on sale, or otherwise available to the public before the effective filing date of the claimed invention.”¹ 35 U.S.C. § 102(a). The Federal Circuit has held that “[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 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). “Inherency, however, may not be established by probabilities or possibilities. The mere fact that a certain thing may result from a given set of circumstances is not sufficient.” *Continental Can Company USA v. Monsanto Company*, 948 F.2d 1264, 1269 (Fed. Cir. 1991). To be considered anticipatory, a prior art reference must describe the applicant’s “claimed invention sufficiently to have placed it in possession of a person of ordinary skill in the field of the invention.” *Helifix Ltd. v. Blok-Lok, Ltd.*, 208 F.3d 1339, 1346 (Fed. Cir. 2000) (quoting *In re Paulsen*, 30 F.3d 1475, 1479 (Fed. Cir. 1994)). Anticipation is a question of fact. *Texas Instruments, Inc. v. U.S. Int’l Trade Comm’n*, 988 F.2d 1165, 1177 (Fed. Cir. 1993).

2. Obviousness

Under 35 U.S.C. § 103(a), a patent is valid unless “the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would

¹For patent applications filed before March 16, 2013, the relevant priority date is “before the invention thereof by the applicant for a patent.” See MPEP § 2131.

PUBLIC VERSION

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.” 35 U.S.C. § 103(a). The ultimate question of obviousness is a question of law, but “it is well understood that there are factual issues underlying the ultimate obviousness decision.” *Richardson-Vicks Inc. v. Upjohn Co.*, 122 F.3d 1476, 1479 (Fed. Cir. 1997); *Wang Lab., Inc. v. Toshiba Corp.*, 993 F.2d 858, 863 (Fed. Cir. 1993). 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. *Graham v. John Deere Co.*, 383 U.S. 1, 17 (1966).

Although the Federal Circuit has historically required that, in order to prove obviousness, the patent challenger must demonstrate, by clear and convincing evidence, that there is a “teaching, suggestion, or motivation to combine,” the Supreme Court has rejected this “rigid approach.” *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 417-418 (2007). In *KSR*, the Supreme Court 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. Since *KSR* was decided, 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

PUBLIC VERSION

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).

3. Written Description and Enablement

35 U.S.C. § 112 is the basis for the written description and enablement requirements:

The specification shall contain a written description of the invention, and the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same . . .

35 U.S.C. § 112, ¶ 1.

The hallmark of the written description requirement is the disclosure of the invention.

Ariad Pharm., Inc. v. Eli Lilly and Co., 598 F.3d 1336, 1351 (Fed. Cir. 2010) (en banc). The test for determining the sufficiency of the written description in a patent requires “an objective inquiry into the four corners of the specification from the perspective of a person of ordinary skill in the art. Based on that inquiry, the specification must describe an invention understandable to that skilled artisan and show that the inventor actually invented the invention claimed.” *Id.*

Compliance with the written description requirement is a question of fact and “the level of detail required to satisfy the written description requirement varies depending on the nature and scope of the claims and on the complexity and predictability of the relevant technology.” *Id.*

“To be enabling, the specification of a patent must teach those skilled in the art how to make and use the full scope of the claimed invention without ‘undue experimentation.’”

Genentech, Inc. v. Novo Nordisk, A/S, 108 F.3d 1361, 1365 (Fed.Cir.1997) (quoting *In re Wright*, 999 F.2d 1557, 1561 (Fed. Cir. 1993)). Enablement serves the dual function in the patent system of ensuring adequate disclosure of the claimed invention and of preventing claims broader than the disclosed invention. *MagSil Corp. v. Hitachi Global Storage Technologies, Inc.*, 687 F.3d 1377, 1380 -1381 (Fed. Cir. 2012). “The scope of the claims must be less than or equal to the

PUBLIC VERSION

scope of the enablement to ensure that the public knowledge is enriched by the patent specification to a degree at least commensurate with the scope of the claims.” *Sitrick v. Dreamworks, LLC*, 516 F.3d 993, 999 (Fed. Cir. 2008) (quoting *Nat’l Recovery Techs., Inc. v. Magnetic Separation Sys., Inc.*, 166 F.3d 1190, 1195-96 (Fed. Cir. 1999). The enablement determination proceeds as of the effective filing date of the patent. *Plant Genetic Sys., N.V. v. DeKalb Genetics Corp.*, 315 F.3d 1335, 1339 (Fed. Cir. 2003).

C. Domestic Industry – Technical 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. *Certain Stringed Musical Instruments and Components Thereof*, Inv. No. 337-TA-586, Comm’n Op. at 12-14 (May 16, 2008). 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) (unreviewed by Commission in relevant part).

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) and (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 (Jan. 16, 1996). “In order to satisfy the technical prong of the 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

PUBLIC VERSION

patent.” *Certain Ammonium Octamolybdate Isomers*, Inv. No. 337-TA-477, Comm’n Op. at 55 (August 28, 2003).

The test for claim coverage for the purposes of the technical prong of the domestic industry requirement is the same as that for infringement. *Certain Doxorubicin and Preparations Containing Same*, Inv. No. 337-TA-300, Initial Determination at 109, (May 21, 1990), *aff’d*, Views of the Commission at 22 (October 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.” Inv. No. 337-TA-300, Initial Determination at 109. To prevail, the patentee must establish by a preponderance of the evidence that the domestic product practices one or more claims of the patent. 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, Initial Determination at 44, Pub. No. 2575 (May 15, 1992).

IV. U.S. PATENT NO. 7,209,140

U.S. Patent No. 7,209,140 (“the ‘140 patent”) is titled “System, Method and Article of Manufacture for a Programmable Vertex Processing Model With Instruction Set.” (JX-006, ‘140 patent). The ‘140 patent issued on April 24, 2007, and lists John Erik Lindholm, David B. Kirk, Henry P. Moreton, and Simon Moy as inventors. (*Id.*). The ‘140 patent has five figures and 14 claims. (*Id.*). Independent claims 1, 5-7, 12, and 14, and dependent claims 2-4 and 8-10 were asserted in this investigation. *See* 79 Fed. Reg. 61338 (Oct. 10, 2014) (“Notice of Investigation”). However, only claim 14 of the ‘140 patent remains at issue in this investigation. (CIB at 2; Tr. at 14:17-20; CX-006C (Aliaga) at Q&A 21).

PUBLIC VERSION

The '140 patent generally relates to hardware accelerated computer graphics. (See '140 patent at 1:17-19, 51-55). According to the '140 patent:

Graphics application program interfaces (API's) have been instrumental in allowing applications to be written to a standard interface and to be run on multiple platforms, i.e. operating systems. Examples of such API's include Open Graphics Library (OpenGL®) and D3D™ transform and lighting pipelines. OpenGL® is the computer industry's standard graphics API for defining 2-D and 3-D graphic images. With OpenGL®, an application can create the same effects in any operating system using any OpenGL®-adhering graphics adapter. OpenGL® specifies a set of commands or immediately executed functions. Each command directs a drawing action or causes special effects.

(JX-006 at 1:23-34). One of the benefits of standardized APIs, such as OpenGL and D3D, is the ability to optimize the available commands using hardware graphics accelerators. (*Id.* at 1:43-47). However, standardized APIs were slow to change. (*Id.* at 1:43-50). Accordingly, the '140 patent asserts that there was “a need to provide a new computer graphics programming model and instruction set that allows convenient implementation of changes to the graphics API, while preserving the driver and hardware optimization afforded by currently established graphics API's.” (*Id.* at 1:57-61).

To that end, the '140 patent proposes a programmable hardware graphics accelerator that provides “instructions from a predetermined instruction set” to make various operations available to programmers. (*Id.* at 1:65-2:4).

NVIDIA alleges infringement of claim 14 of the '140 patent, which reads as follows with emphasis added to indicate disputed and construed terms:

Claim 14. A system, comprising:

a central processing unit; and

a hardware graphics accelerator for receiving graphics data, and performing programmable operations on the graphics data in order to generate output;

PUBLIC VERSION

wherein the operations are programmable by a user utilizing instructions from a predetermined instruction set capable of being executed by the hardware graphics accelerator, the predetermined instruction set including a reciprocal instruction, a reciprocal square root instruction, a three component dot product instruction, a four component dot product instruction, a distance instruction, a minimum instruction, a maximum instruction, an exponential instruction, and a logarithm instruction.

(JX-006 at 24:23-38.)

A. Level of Ordinary Skill in the Art

In Order No. 20, I found that a person of ordinary skill in the art of the '140 patent would have “at least a four-year degree in Electrical Engineering, Computer Engineering, Computer Science, or equivalent, as well as at least two years of experience in graphics processing including developing, designing or programming software or hardware for graphics processing units, hardware graphics accelerators or other graphics processing systems.” (Markman Order at 18 (April 2, 2015); CIB at 18, 58; RIB at 117.)

B. Claim Construction

1. Order No. 20: Construing Terms of the Asserted Patents

With respect to claim 14 of the '140 patent, I construed the following terms:

Claim Language	Construction (Order No. 20)
“operation”	“an action or process recognized by the hardware graphics accelerator”
“instructions from a predetermined instruction set”	“the complete set of instructions recognized by a given computer or provided by a given programming language”

(Markman Order at 25, 32.)

2. Agreed Construction - “programmable by the user”

The parties agree that the term “programmable by the user” should be construed to mean “an application writer can create graphics functionality by causing instructions from a predetermined instruction set to be executed.” (CIB at 45; RIB at 12-13; SIB at 14.)

PUBLIC VERSION

3. Disputed Constructions

a. “graphics data”/“operations on the graphics data”

The parties dispute the meaning of the phrases “graphics data” and “operations on the graphics data” in claim 14 of the '140 patent. (CIB at 51-54; RIB at 13-14, SIB at 14-15.)

Term	Proposed Constructions		
	Complainant	Respondents	Staff
“graphics data”	data used in vertex processing	data related to graphics	data related to graphics
“operations on the graphics data”	operations during vertex processing of graphics data	No construction necessary. See “operations” and “graphics data” above.	No construction necessary. See “operations” and “graphics data” above.

The Parties’ Positions

NVIDIA argues that its proposed constructions of “graphics data” and “operations on the graphics data” properly define the scope of the claimed invention. (CIB at 51.) NVIDIA asserts the '140 Patent is directed to the vertex processing portion of the graphics pipeline. (*Id.*)

NVIDIA argues its constructions reflect this and focus on operations performed (and data used) during the vertex processing. Specifically, NVIDIA argues its construction of “graphics data” is correct because it covers both types of data used during vertex processing (*e.g.*, constant and/or vertex data as disclosed in the '140 Patent), and excludes data used only in portions of the pipeline that are not subject of the '140 Patent, (*e.g.*, pixel data). (*Id.* at 51-52.)

NVIDIA argues Respondents’ proposal construction of “graphics data” as any “data related to graphics” would broaden the claims to cover operations performed in portions of the graphics pipeline not contemplated by the '140 Patent. (*Id.* at 52.)

NVIDIA argues that Respondents’ construction ignores the legal requirement that the term be construed in the context of the intrinsic evidence (*e.g.*, the claims, specification, and file history). (*Id.*) According to NVIDIA, “Where the specification makes clear at various points

PUBLIC VERSION

that the claimed invention is narrower than the claim language might imply, it is entirely permissible and proper to limit the claims.” (*Id.* (citing *Alloc, Inc. v. ITC*, 342 F.3d 1361, 1370 (Fed. Cir. 2003).) NVIDIA argues the specification consistently describes “the present invention” as enabling a user to “program a portion of the graphics pipeline that handles vertex processing.” (*Id.*) NVIDIA also asserts that the specification distinguishes between “the programmable vertex processing of the present invention” and the “remaining portions of the graphics pipeline” that are not part of the invention. (*Id.* at 53.) NVIDIA argues that every embodiment in the ‘140 patent is directed to the vertex processing portion of the graphics pipeline and performing operations on data used in vertex processing, specifically constant and/or vertex data. (*Id.*) NVIDIA argues nothing in the patent suggests the invention is directed to another portion of the graphics pipeline. (*Id.*) Thus, NVIDIA argues that because the ’140 Patent specification repeatedly and consistently describes the “present invention” as a whole as directed to programmable vertex processing, and only discloses embodiments directed to the same, its construction is proper and should be adopted. (*Id.* at 54.)

Respondents argue that NVIDIA improperly seeks to limit the scope of claim 14 by construing the term “graphics data” to include only one type of graphics data—graphics data used in vertex processing. (RIB at 13.) Respondents argue the intrinsic evidence places no such limitation on the term. (*Id.*) Respondents argue the term “graphics data” should be construed in accordance with its plain meaning: “data related to graphics.” (*Id.* at 14.) Likewise, Respondents argue the plain meaning of “operations on the graphics data” is “operations” on that same data. (*Id.*) Respondents assert that NVIDIA’s expert, Dr. Aliaga, confirmed that the plain meaning of “graphics data” is not limited to vertex data but includes, for example, pixel data. (*Id.*) Respondents argue that the intrinsic evidence also supports its construction noting that the

PUBLIC VERSION

Field of Invention states “the invention relates to computer graphics” in general and dependent claims 2 and 8 separately limit graphics data to vertex data. (*Id.*)

The Staff argues that the plain and ordinary meaning of the term “graphics data” is simply “data related to graphics.” (SIB at 15.) The Staff argues this construction is consistent with the claim language and with the specification. (*Id.*) The Staff argues NVIDIA seeks a narrower construction that would import a limitation from the specification to narrowly construe “graphics data” as “data used in vertex processing.” (*Id.*) Such a construction, the Staff asserts, is improper in the absence of an expression of intent by the patentee. (*Id.* (citing *Phillips*, 415 F.3d at 1323; *Vitronics*, 90 F.3d at 1582-83; *Markman*, 52 F.3d at 979-80; *Intel Corp. v. U.S. International Trade Commission*, 946 F.2d at 836 (“Where a specification does not require a limitation, that limitation should not be read from the specification into the claims.”).) The Staff argues neither the plain language of the claim nor the specification of the '140 patent evidence an intent by the patentee to limit the term “graphics data” beyond its plain and ordinary meaning. (*Id.*) Thus, according to the Staff, NVIDIA’s proposed construction should be rejected and the term “graphics data” should be construed as “data related to graphics.” (*Id.*) With regard to the phrase “operations on the graphics data,” the Staff argues this phrase need not be construed in light of the constructions of “operations” and “graphics data.” (*Id.*)

Discussion

The term “graphics data” is used only in the claims of the '140 patent. There is no recitation of the term in the specification. The Federal Circuit has stated that “[i]n some cases, the ordinary meaning of claim language as understood by a person of skill in the art may be readily apparent even to lay judges.” *Phillips v. AWH Corp.*, 415 F.3d 1303, 1314 (Fed. Cir. 2005). Here, the term “graphics data” is broad and its meaning self-explanatory. “Graphics

PUBLIC VERSION

data” is data related to graphics. This interpretation is consistent with the term’s plain and ordinary meaning as understood by one of ordinary skill in the art. (RX-005C (Diefendorff) at Q&A 92, 96.) This interpretation is also consistent with the claims and specification of the ‘140 patent. (*Id.* at Q&A 94.)

NVIDIA seeks to narrowly construe “graphics data” as “data used in vertex processing.” NVIDIA argues that term “graphics data” should be read to exclude certain types of graphics data, such as pixel data, based on its disavowal of claim scope. (*See* CIB at 51-54.) Specifically, NVIDIA contends that claim 14 of the '140 patent should be limited to programmable vertex processing for two reasons. First, NVIDIA argues that the specification of the '140 patent repeatedly and consistently describes the invention as a whole as requiring programmable vertex processing. (*Id.* at 54.) Second, NVIDIA argues that the '140 patent only discloses embodiments that provide programmable vertex processing. (*Id.*) For the reasons set forth below, the evidence does not support NVIDIA’s position.

NVIDIA relies on portions of the specification that refer to “vertex processing” or “programmable vertex processing” for the proposition that the claimed invention should be limited to only programmable vertex processing. (*See* CIB at 52-53.) The cited evidence, however, does not demonstrate clear and unmistakable disclaimer, as is necessary to depart from the usual and customary meaning. *Thorner*, 669 F.3d at 1366-67 (“To constitute disclaimer, there must be a clear and unmistakable disclaimer.”). Although the preferred embodiments describe programmable vertex processing, not programmable primitive processing, nothing in the patent precludes programmable primitive processing. Indeed, the '140 patent even contemplates alternative embodiments that enable programmable primitive processing. (JX-006 at 4:1-14 (“the programmable graphics mode may also supersede the standard graphics API

PUBLIC VERSION

during primitive processing”).) Moreover, the language of the claims suggests NVIDIA’s construction cannot be correct. Independent claims 1 and 7 both recite “graphics data,” while dependent claims 2 and 8 state that graphics data “includes vertex data.” This claim language implies the term graphics data is broader than just vertex data.

Also, contrary to NVIDIA’s argument, the '140 patent does not repeatedly and consistently require programmable vertex processing. For example, the “Field of Invention” states broadly that the “present invention relates to computer graphics, and more particularly to providing programmability in a computer graphics processing pipeline.” (JX-006 at 1:15-19.) The evidence shows vertex processing is just one part of the graphics pipeline. (RX-005C (Diefendorff) at Q&A 94.) Likewise, the “Description of the Invention” describes the invention without reference to vertex processing:

A system, method and article of manufacture are provided for programmable processing in a computer graphics pipe line. Initially, data is received from a source buffer. Thereafter, programmable operations are performed on the data in order to generate output. The operations are programmable in that a user may utilize instructions from a predetermined instruction set for generating the same. Such output is stored in a register. During operation, the output stored in the register is used in performing the programmable operations on the data.

(JX-006 at 1:65-2:7.) The Abstract of the '140 patent similarly describes the invention broadly without limiting the claimed invention to vertex processing. (*Id.*, Abstract).

Nothing in the intrinsic record indicates the applicant acted as his own lexicographer by assigning a special meaning to the term “graphics data” and nothing in the specification or prosecution history shows the applicant clearly and unmistakably disavowed claim scope. Thus, I find no reason to depart from the term’s plain and ordinary meaning. Such meaning, as discussed above, is also consistent with the language of the claims and specification.

PUBLIC VERSION

Accordingly, I find for at least the reasons above, that one of ordinary skill in the art at the time of the invention would have construed the term “graphics data” as “data related to graphics.”

In light of my construction of “graphics data” as “data related to graphics” I find it unnecessary to further construe the phrase “operations on graphics data.”

b. “hardware graphics accelerator”

The parties dispute the meaning of the phrase “hardware graphics accelerator” in claim 14 of the '140 patent. (CIB at 47-51; RIB at 14-17.)

Term	Proposed Constructions		
	Complainant	Respondents	Staff
“hardware graphics accelerator”	“graphics processing unit (GPU)”	plain and ordinary meaning – “hardware for processing graphics”	plain and ordinary meaning – “hardware for processing graphics”

The Parties’ Positions

NVIDIA asserts that the dispute regarding this term centers on whether a “hardware graphics accelerator” requires a hardware implementation of the graphics pipeline, *i.e.*, a GPU, or whether it can be construed to encompass general-purpose processors programmed to process graphics data. (CIB at 47-48.) NVIDIA argues the intrinsic record of the patent makes clear the term does not include general-purpose processors. (*Id.* at 48.) NVIDIA asserts that one of ordinary skill understands the '140 Patent uses “hardware graphics accelerator” to refer to a hardware implementation of a graphics pipeline, *i.e.*, a GPU. (*Id.*) NVIDIA argues that the claims distinguish between hardware graphics accelerators and CPUs and that claim 14 separately recites both. (*Id.*) NVIDIA argues that the patent specification also distinguishes between hardware graphic accelerators and CPUs. (*Id.*) NVIDIA argues the “Background of the Invention” makes clear that specialized chips having a hardware implementation of the graphics

PUBLIC VERSION

pipeline are central to the invention. (*Id.*) NVIDIA asserts that the '140 patent explains how the prior art approach of implementing the graphics pipeline in specialized hardware afforded benefits in terms of acceleration but suffered from a lack of flexibility due to its fixed-function nature. (*Id.*) NVIDIA asserts that the '140 patent solved this problem by adding programmability to the vertex processing portion of a specialized graphics chip, which is the claimed invention of the '140 patent. (*Id.* at 49.) NVIDIA argues that the inventors confirmed the '140 patent was directed at specialized hardware and not general-purpose processors such as CPUs. (*Id.*)

NVIDIA maintains the discussion in the '140 patent about adding programmability to graphics chips only makes sense as a description of hardware graphics accelerators as NVIDIA construes the term (*i.e.*, a hardware implementation of a graphics pipeline), because general-purpose processors are inherently programmable. (*Id.*) NVIDIA alleges that during prosecution, the applicants and the Examiner agreed that general-purpose processor art was not analogous to the claimed hardware graphics accelerators. (*Id.* at 49-50.) NVIDIA asserts that to overcome a prior art general-purpose processor, applicants amended the claims to repeatedly recite a "hardware graphics accelerator." (*Id.* at 50.)

NVIDIA contends Respondents ignore the specification, file history and other evidence to offer a construction that captures general purpose processors. (*Id.*) NVIDIA argues that according to Respondents, a general-purpose processor is a "hardware graphics accelerator" if used to process graphics, but not if used for a different purpose. (*Id.* at 51.) NVIDIA argues one of ordinary skill would understand that whether something is a "hardware graphics accelerator" is determined by its structure and not its use. (*Id.*) NVIDIA argues that Respondents reliance on the conclusory assertion of their expert, Mr. Diefendorff, to show that hardware graphics

PUBLIC VERSION

accelerators encompass general purpose processors must be rejected as inconsistent with the intrinsic evidence. (*Id.*)

Respondents argue that both the intrinsic evidence and their expert, Mr. Diefendorff, confirm that the term “hardware graphics accelerator” (“HGA”) should be given its plain and ordinary meaning as “hardware for processing graphics data.” (RIB at 14.) Respondents argue that the ’140 patent does not ascribe any special meaning to the term, instead stating that it can take on “various configuration[s].” (*Id.* at 15.) Respondents argue that despite conceding that the term “HGA” should be given its plain and ordinary meaning, NVIDIA improperly limits the term to a “graphics processing unit” with specialized, fixed-function circuitry for rasterization and texturing. (*Id.*) But, Respondents argue, the ’140 patent never mentions a “GPU,” a “rasterizer,” or “special” texture hardware, and NVIDIA cannot cite any intrinsic support. (*Id.*) Respondents assert that one of the named ’140 inventors conceded that such fixed-function hardware is not needed to perform the claimed instructions. (*Id.*)

Respondents argue NVIDIA’s reliance on the prosecution history in support of its proposed construction is misplaced. (*Id.* at 15-16.) Respondents assert the prosecution history never defines or explains the meaning of “hardware graphics accelerator” and that despite the claims being rejected three times, NVIDIA only mentioned HGA once, to contrast the instructions from an “archaic” general processor that never mentioned graphics with those of a hardware graphics accelerator. (*Id.* at 16.) Respondents argue that during prosecution, NVIDIA never mentioned a GPU, fixed-function hardware, a rasterizer, or a texture unit. (*Id.*) Respondents argue that while claim 14 mentions both a CPU and an HGA, it does not limit the meaning of HGA as NVIDIA contends nor does it support NVIDIA’s construction. (*Id.*) Respondents also argue that NVIDIA’s alleged plain and ordinary meaning of “hardware

PUBLIC VERSION

graphics accelerator” fails to account for the term’s actual usage in the field many years before NVIDIA’s coined the term “GPU.” (*Id.*)

The Staff argues the term should be given its plain and ordinary meaning, which is simply “hardware for processing graphics.” (SIB at 16.) The Staff argues this construction is consistent with the claim language and with the specification. (*Id.*) The Staff argues that NVIDIA improperly seeks a narrower construction that would limit the claimed “hardware graphics accelerator” to a GPU. (*Id.*)

Discussion

The evidence shows the plain and ordinary meaning of the term “hardware graphics accelerator” is simply “hardware for processing graphics.” (RX-005C (Diefendorff) at Q&A 86-91; RX-3497C (Diefendorff) at Q&A 23; *see also* RX-3494C (Yu) at Q&A 20-30.) This construction is consistent with the claim language, specification, and prosecution history. (RX-3497C (Diefendorff) at Q&A 25, 27, 28.) The patent specification does not use the term “hardware graphics accelerator,” and only once refers to a “hardware accelerator.” There is nothing in the intrinsic record to indicate that the patent applicant assigned this limitation a special meaning or disclaimed part of this limitation’s plain and ordinary meaning.

As the name describes, a “hardware graphics accelerator” accelerates the processing of graphics data, which the evidence shows in practice is circuitry that can be implemented in different ways depending on the design needs. (RX-005C (Diefendorff) at Q&A 91.) The evidence shows hardware graphics accelerators can take many forms, including “something as simple as additional hardware added to a CPU or something more complex such as a separate processor.” (*Id.* at Q&A 86, 91; RX-3497C (Diefendorff) at Q&A 23.) Contrary to NVIDIA’s argument, the evidence shows there were many known programmable processors described in

PUBLIC VERSION

the art as accelerating graphics processing that do not contain fixed, dedicated hardware. (RX-005C (Diefendorff) at Q&A 90, 91; RX-3497C (Diefendorff) at Q&A 23.)

NVIDIA's main argument in support of its proposed construction is that a hardware graphics accelerator must be different from a CPU because claim 14 separately recites both a hardware graphics accelerator and a CPU, the specification says a hardware graphics accelerator can be implemented in "hardware accelerators of various configuration" or a CPU, and the file history distinguishes a general purpose processor from an hardware graphics accelerator. (CIB at 48-50.) This argument, however, is irrelevant as no party has proposed construing "hardware graphics accelerator" to mean a CPU. (RX-3497C (Diefendorff) at Q&A 28 ("I have not said that a general-purpose CPU, regardless of how it is used in a system, is automatically a hardware graphics accelerator. On the contrary, I have contended only two things: first, that a CPU microprocessor can be employed to function in a system as a hardware graphics accelerator if it is assigned to offloading graphics work from the main CPU, or, second, that a general-purpose CPU microprocessor can be transformed into a hardware graphics accelerator through the addition of new instructions and hardware execution units for the purpose of processing graphics data effectively.")) I find no support in the claim language, specification, or prosecution history of the '140 patent for the limiting construction offered by NVIDIA. (RX-005C (Diefendorff) at Q&A 86-91; (RX-3497C (Diefendorff) at Q&A 25, 27, 28.)) Thus I find NVIDIA's argument in support of its proposed construction not persuasive.

NVIDIA seeks to narrow the plain and ordinary meaning of "hardware graphics accelerator" by construing the term as a separate processor (*i.e.*, a GPU). However, the portions of the intrinsic record relied on by NVIDIA in support of its construction neither define a hardware graphics accelerator nor limit a hardware graphics accelerator to a GPU. In fact, the

PUBLIC VERSION

terms “graphics processing unit” and “GPU” appear nowhere in the patent or file history of the ‘140 patent.

With regard to the prosecution history, NVIDIA argues that during prosecution “the applicants and the Examiner agreed that general-purpose processor art was not analogous to the claimed hardware graphics accelerators.” (CIB at 49-50.) Specifically, NVIDIA argues that to overcome a rejection based on a prior art general-purpose processor, the applicant stated:

In particular, only applicant teaches and claims “performing programmable operations on the graphics data utilizing the hardware graphics accelerator ...

In particular, applicant emphasizes that, while Deering teaches a graphics system, Struble discloses a general-purpose processor assembler language. ... To simply glean features from the art of general-purpose processor assembler languages and combine the same with the non-analogous art of graphics systems would be improper and frustrate the inventive concepts of applicant, especially in view of the fundamentally different problems which the two arts address.

(CIB at 50 (quoting JX-12.0438-.0448) (emphasis in original)). NVIDIA argues that following this response, the examiner no longer cited references disclosing programming of general-purpose processors. (*Id.*) Contrary to NVIDIA’s argument, the applicant never argued that general-purpose processors were non-analogous to the claimed graphics hardware accelerators. Rather, as the quoted text makes clear, the applicant was only arguing that “general-purpose processor assembler languages” were non-analogous to “graphics systems.” Moreover, contrary to NVIDIA’s assertion, the patent examiner never stated that general-purpose processors were “not analogous.” At most, the examiner was silent, and examiner “silence is not a proper basis on which to construe a patent claim.” *DeMarini Sports, Inc. v. Worth, Inc.*, 239 F.3d 1314, 1326 (Fed. Cir. 2001).

Accordingly, for at least the reasons above, I find one of ordinary skill in the art at the time of the invention would construe the term “hardware graphics accelerator” to have its plain and ordinary meaning of “hardware for processing graphics.”

PUBLIC VERSION

c. “reciprocal instruction”

Term	Proposed Constructions		
	Complainant	Respondents	Staff
“reciprocal instruction”	Plain and ordinary meaning once “instruction” is construed. However, if construed: “characters used to specify an action or process for calculating a reciprocal”	a particular instruction that inverts a single operand	a particular instruction that inverts a single operand

The Parties’ Positions

NVIDIA argues the term “reciprocal instruction” has a plain and ordinary meaning. (CIB at 46.) NVIDIA argues that a “reciprocal” is a well-known mathematical expression, it is just one divided by x, *i.e.*, 1/x or the “inverse” of x. (*Id.*) Thus, NVIDIA argues, a “reciprocal instruction” is a programming language statement specifying the calculation of a reciprocal. (*Id.*) NVIDIA argues this is consistent with the ’140 patent’s description of reciprocal. (*Id.*) NVIDIA argues nothing in the ’140 patent requires the claimed “reciprocal instruction” to be anything more than an instruction that performs a mathematical reciprocal. (*Id.*) Accordingly, NVIDIA contends “reciprocal instruction” should be construed to mean “characters used to specify an action or process for calculating a reciprocal.” (*Id.*)

NVIDIA argues the Respondents and Staff’s proposed construction – “a particular instruction that inverts a single operand” – inserts the extraneous and ambiguous words, “particular” and “single.” (*Id.*) NVIDIA argues the plain and ordinary meaning does not require such and neither does the ’140 patent. (*Id.*) NVIDIA argues the ’140 patent teaches broadly that the instruction may be in “any type of programming language” and thus no “particular” format is required. (*Id.*)

PUBLIC VERSION

NVIDIA contends that Respondents' non-infringement position demonstrates they actually seek an even narrower construction. (*Id.* at 47.) NVIDIA asserts that Respondents' expert, Mr. Diefendorff, bases his non-infringement opinion on a belief that reciprocal instructions must be completely distinct from divide instructions and thus Respondents are really seeking a construction of "reciprocal instruction" that expressly excludes any overlap with a division instruction. (*Id.*) NVIDIA argues the '140 patent contains no such disclaimer. (*Id.*) According to NVIDIA, the '140 Patent never distinguishes reciprocal instructions from divide instructions. (*Id.*) NVIDIA asserts that in fact, divide instructions are never mentioned in the patent or its prosecution history. (*Id.*) Instead, NVIDIA argues the evidence shows that reciprocal is a type of division. (*Id.*) NVIDIA argues Respondents' narrow construction finds no support in the intrinsic evidence or the plain and ordinary meaning of a reciprocal instruction and should be rejected. (*Id.*)

Respondents argue that "reciprocal instruction" means "a particular instruction that inverts a single operand." (RIB at 13.) Respondents argue "reciprocal" modifies "instruction," such that the term is limited to one type of instruction, not a mathematical calculation in the abstract. (*Id.*) Respondents argue a PHOSITA would have understood a "reciprocal instruction" to be monadic, *i.e.*, it receives and operates on a single input or "operand." (*Id.*) Respondents argue this is consistent with the patent's only explanation of the claimed reciprocal instruction. (*Id.*)

Respondents argue that NVIDIA incorrectly rewrites the claimed "reciprocal instruction" to cover a "division" instruction. (*Id.*) However, Respondents argue a PHOSITA would know that a division instruction is different. (*Id.*) Respondents argue that among other things, a division instruction is dyadic, *i.e.*, it receives and operates on two inputs, a numerator and

PUBLIC VERSION

denominator. (*Id.*) Respondents argue the difference is significant as reciprocal can be easily calculated at significantly greater speed than division. (*Id.*) Respondents also contend there are differences in latency and accuracy, and the prior art consistently distinguished between the two instructions. (*Id.*) Respondents note the '140 patent never mentions a division instruction (Diefendorff, Tr. 960:10-19), and the inventors distinguished between the claimed reciprocal instruction and a division instruction. (*Id.*)

The Staff argues that a division instruction is not the same as a reciprocal instruction. (SIB at 17.) The Staff asserts that division is a dyadic operation, while reciprocal is a monadic operation. (*Id.*) In addition, the Staff argues reciprocal instructions are more efficient to pipeline because they exhibit constant latency, while division instructions exhibit variable latency and increase complexity. (*Id.*) The Staff argues one of ordinary skill in the art would not understand a reciprocal instruction to mean a division instruction. (*Id.*) Thus, the Staff contends the term “reciprocal instruction” should be construed to mean “a particular instruction that inverts a single operand.”

Discussion

NVIDIA seeks a construction of “reciprocal instruction” that would cover division. While it is true that a reciprocal of a number can be calculated by taking that number and dividing by 1, NVIDIA’s proposed construction improperly conflates reciprocal with division focusing on the meaning of “reciprocal” in the abstract as a mathematical expression, while ignoring that the actual limitation in question is “*reciprocal instruction*” not just reciprocal. Here, the evidence shows the limitation “reciprocal instruction” is a term that is well known in the art and one that would have been known to one of ordinary skill in the art at the time of the invention. *Phillips*, 415 F.3d 1303, 1318 (Fed. Cir. 2005) (“extrinsic evidence in the form of

PUBLIC VERSION

expert testimony can be useful to a court for a variety of purposes, such as ... to establish that a particular term in the patent or the prior art has a particular meaning in the pertinent field.”

The evidence shows a division instruction is not the same as a reciprocal instruction. (RX-005C (Diefendorff) at Q&A 102-103; RX-3497C (Diefendorff) at Q&A 43-46, 73-77). A person having ordinary skill in the art would have understood a “division instruction” to be a dyadic operation (*i.e.*, it receives and operates on two inputs or “operands”) and a “reciprocal instruction” to be monadic (*i.e.*, it receives and operates on a single input or “operand”). (RX-3497C (Diefendorff) at Q&A 71; RX-005C (Diefendorff) at Q&A 102; Tr. 1006:11-18.) The evidence shows this difference to be significant as a reciprocal can be easily calculated at significantly greater speed than division. (RX-3497C (Diefendorff) at Q&A 73-77; RX-005C (Diefendorff) at Q&A 103. The evidence also shows there are differences in latency and accuracy, and that the prior art consistently distinguished between the two types of instructions. (RX-005C (Diefendorff) at Q&A 103; RX-3497(Diefendorff) at Q&A 43-46, 73-77.)

As discussed above, one of ordinary skill in the art would understand a “reciprocal instruction” is monadic. This is consistent with the patent’s only description of the claimed reciprocal instruction. As shown below, the “reciprocal instruction” is described in the ‘140 patent as receiving and operating on a single operand. (RX-005C (Diefendorff) at Q&A 102; JX-006 at 8:16-35, 12:19-26.)

OPCODE	INPUT (scalar or vector)	OUTPUT (replicated scalar or vector)
NOB		
ARL	s	
MOV	v	v
MUL	v,v	v
ADD	v,v	v
MAD	v,v,v	v
RCP	s	s,s,s of v or v or v

(JX-006 at 8:16-35 (emphasis added).)

PUBLIC VERSION

Reciprocal (RCP)

Format:

RCP D[.xyzw],[-]S0[xyzw]

Description:

The present instruction inverts a source scalar into a destination. The source may have one subscript. Output may be exactly 1.0 if the input is exactly 1.0.

(JX-006 at 12:19-26 (emphasis added) (“The present instruction inverts a source scalar”).) Thus, the patent’s description of a reciprocal instruction is inconsistent with NVIDIA’s attempt to construe the term to read on a division instruction, which was known in the art as requiring two operands. (RX-005C (Diefendorff) at Q&A 103; RX-219 at 152; RX-2774 at 8-89, 7-97.)

NVIDIA’s reliance on the IEEE dictionary’s definition of “instruction” is misplaced. In addition to not defining the limitation-at-issue, “reciprocal instruction,” it confirms that an instruction consists of both “an operation and its operand (if any).” As it is undisputed that division and reciprocal instructions operate on a different number of operands, the IEEE definition does not support NVIDIA’s attempt to expand the meaning of “reciprocal instruction” to include a division instruction.

NVIDIA argues “a ‘reciprocal instruction’ is a programming language statement specifying the calculation of a reciprocal.” (CIB at 46.) To the extent NVIDIA is arguing that an instruction is a “programming language statement” NVIDIA’s argument is flawed as such a construction is not supported by the ‘140 patent and is inconsistent with my construction of the limitation “instruction set” as the “complete set of instructions ... *provided* by a given programming language.” (Order No. 20 at 32 (emphasis added).)

PUBLIC VERSION

Accordingly, I find for at least the reasons above, one of ordinary skill in the art at the time of the invention would construe the limitation “reciprocal instruction” to mean “a particular instruction that inverts a single operand.”²

C. Infringement

NVIDIA argues the Accused Products incorporating Mali-T7xx, Mali-T6xx, Adreno 4xx, Adreno 3xx, PowerVR SGX 544, or PowerVR SGX 540 GPUs infringe Claim 14 of the '140 Patent. (CIB at 55.)

The Accused Products include GPUs compliant with the OpenGL ES API. (CX-006C (Aliaga) at Q&A 361-362.) The OpenGL ES API uses a shading language known as OpenGL ES Shading Language (“ESSL”). (*Id.*) The OpenGL ES Shading Language provides a standard API that can be used to define programmable shaders. (*See* CX-343.) Like most programming languages, ESSL provides a set of operators and built-in functions in which programmers can write source code. (RX-3497C (Diefendorff) at Q&A 63.)

NVIDIA’s infringement allegations focus on whether the use of ESSL to program the Accused Products infringes claim 14. (CIB at 55 (“At trial NVIDIA focused its infringement allegations on the use of ESSL to program the Accused Products.”).) In this regard, NVIDIA’s expert, Dr. Aliaga, testified in detail that the use of ESSL in the Accused Products infringes claim 14. (CX-006C (Aliaga) at Q&A 360-366, 369-372.)

² NVIDIA argues this construction of “reciprocal instruction” is incorrect because “particular” and “single” are “extraneous and ambiguous words.” (CIB at 46.) But “single,” which modifies “operand,” is clear on its face and means only “one” operand. With respect to “particular,” NVIDIA tries to twist the construction to argue that “particular instruction” requires the instruction have a “particular format.” (*Id.*) Nothing, however, in the construction I have adopted requires a particular format. Rather, the phrase merely means there must be “one” instruction. (*See* RX-3487C (Diefendorff) at Q&A 45.) This clarifies that a claimed instruction is not satisfied by a programmer cobbling together two or more different instructions.

PUBLIC VERSION

Respondents and the Staff argue that the Accused Products do not infringe claim 14 because ESSL does not include an instruction set that includes each of the nine predetermined instructions required by claim 14. (RIB at 18; SIB at 19.) Specifically, Respondents and the Staff argue ESSL does not include a “reciprocal instruction.” (*Id.*)

For the reasons discussed in detail below, I find NVIDIA’s arguments are not persuasive. Accordingly, based on the testimony of Respondents’ expert, Dr. Diefendorff, and my discussion of the record evidence, *infra*, I find NVIDIA has failed to prove by a preponderance of the evidence that the Accused Products infringe claim 14 of the ‘140 patent.

Claim 14 requires “a predetermined instruction set” that includes the following nine predetermined instructions: (1) a reciprocal instruction; (2) a reciprocal square root instruction; (3) a three component dot product instruction; (4) a four component dot product instruction; (5) a distance instruction; (6) a minimum instruction; (7) a maximum instruction; (8) an exponential instruction; and (9) a logarithm instruction. (JX-006, claim 14).

NVIDIA repeatedly asks me to decide whether “[t]he accused ESSL reciprocal instruction, $1/op1$ ” infringes. (CIB at 47 n.11; *see also id.* at 38 (“whether the ESSL statement ‘ $1/x$ ’ is a reciprocal instruction”), 59 (“ESSL Instruction Syntax” of “ $1/op1$ ”), 60 (“The ESSL statement, $1/x$, is a ‘reciprocal instruction’” ... “ESSL instruction ‘ $1.0/b$ ’ calculates a reciprocal”).) However, the evidence is clear that no such “instruction” or “statement” exists in ESSL. (RX-3497C (Diefendorff) at Q&A 71-72; CX-343.0046 (Section “5.1 Operators” lists “/,” not “1/”); Diefendorff, Tr. 1002:20-1004:15 (explaining that Section 5.1 provides ESSL’s instruction set).) Dr. Aliaga argues that “1/” is the coding symbol for reciprocal in ESSL, *see* CX-006C (Aliaga) at Q&A 366, but the evidence shows what Dr. Aliaga contends is a reciprocal instruction is actually describing how a programmer might use ESSL’s predefined division

PUBLIC VERSION

instruction by providing it a “1” as one of the two operands. (RX-3497C (Diefendorff) at Q&A 71.) As can be seen below from an excerpt of the ESSL specification, there is no “1/” instruction in ESSL’s predetermined instruction set.

Precedence	Operator Class	Operators	Associativity
1 (highest)	parenthetical grouping	()	NA
2	array subscript	[]	Left to Right
	function call and constructor structure field selector, swizzler post fix increment and decrement	() . ++ -	
3	prefix increment and decrement	++ -	Right to Left
	unary (tilde is reserved)	+ - ~ !	
4	multiplicative (remainder reserved)	* / %	Left to Right
5	additive	+ -	Left to Right
6	bit-wise shift (reserved)	<< >>	Left to Right
7	relational	< > <= >=	Left to Right
8	equality	= !=	Left to Right
9	bit-wise and (reserved)	&	Left to Right
10	bit-wise exclusive or (reserved)	^	Left to Right
11	bit-wise inclusive or (reserved)		Left to Right
12	logical and	&&	Left to Right
13	logical exclusive or	^^	Left to Right
14	logical inclusive or		Left to Right
15	selection	? :	Right to Left
16	assignment arithmetic assignments (remainder, shift, and bit-wise are reserved)	=	Right to Left
		+= -=	
		*= /=	
		%= <<= >>=	
17 (lowest)	sequence	,	Left to Right

(CX-343 at 46.) Rather the ESSL specification only provides a “/” instruction, which is a division instruction. (*Id.* at 46, 54 (“The arithmetic binary operators add (+), subtract (-), multiply (*), and divide (/) ...”).) Notably, Dr. Aliaga admitted that what he accuses of infringement is ESSL’s division instruction, a “slash,” “/.” (Tr. 289:5-291:4.) Dr. Aliaga further admitted that this “slash” is “not on its own” a reciprocal instruction. (*Id.* at 289:25-290:2.)

As previously discussed, *supra*, a reciprocal instruction is monadic (*i.e.*, it receives one operand, or input) and a division instruction is dyadic (*i.e.*, it receives two operands or inputs). The evidence shows a programmer must provide two operands, a numerator and a denominator,

PUBLIC VERSION

to the division instruction (“/”) provided by ESSL. (Tr. at 975:20-976:2; CX-343 at 54 (“The arithmetic binary operators add (+), subtract (-), multiply (*), and divide (/) operate on integer and floating-point typed expressions (including vectors and matrices). The two operands must be the same type, or can be a scalar float and the other a float vector or matrix, or one can be a scalar integer and the other an integer vector. Additionally, for multiply (*), one can be a vector and the other a matrix with the same dimensional size of the vector.”).) That the programmer must program a “1” as the numerator as part of a two-operand instruction necessarily means ESSL does not provide a reciprocal instruction. (RX-3497C (Diefendorff) at Q&A 65, 71-72; Diefendorff, Tr. 972:3-9.)

As the plain language of claim 14 and my construction of “instruction set” make clear, the claimed reciprocal instruction must be “predetermined” and “provided by” the accused programming language, here ESSL. For at least the reasons discussed above, I find ESSL does not have a predetermined instruction set with a “reciprocal instruction.” Accordingly, I find NVIDIA has failed to prove by a preponderance of the evidence that the Accused Products infringe claim 14 of the ‘140 patent.

D. Domestic Industry – Technical Prong

NVIDIA contends that its GPUs with Kepler, Fermi, and Maxwell architectures and products incorporating those GPUs practice claim 14 of the '140 patent. (CX-006C (Aliaga) at Q&A 493; CIB at 63-65.) In particular, NVIDIA relies upon ESSL to show it meets the technical prong. (CIB at 63-65.) However, as discussed above with regard to infringement, ESSL fails to include the “reciprocal instruction,” and thus does not meet the limitations of claim 14. Thus, to the extent that NVIDIA relies upon ESSL, I find NVIDIA fails to show that its domestic industry products satisfy the technical prong of the domestic industry requirement.

PUBLIC VERSION

NVIDIA alternatively relies upon OpenGL extensions NV_gpu_program4 and NV_gpu_program5, which it asserts include each of the claimed predetermined instructions. (*Id.* at 64-65.) In this regard, however, NVIDIA fails to cite to any documentary evidence regarding the design and implementation of these extensions in its domestic industry products. (RX-3497C (Diefendorff) at Q&A 176-77.) Instead, NVIDIA relies on the testimony of Dr. Aliaga, who in turn relies entirely upon the uncorroborated testimony of a named inventor. (*See* CIB at 64-65; CX-006C (Aliaga) at Q&A 498-502; CX-005C (Moreton) at Q&A 58.) Thus, I find this evidence is not entitled to much weight. Therefore, I find NVIDIA has failed to show by a preponderance of the evidence that OpenGL extensions NV_gpu_program4 and NV_gpu_program5 disclose each of the claimed predetermined instructions.

Accordingly, for at least the reasons discussed above, I find NVIDIA has failed to prove by a preponderance of the evidence that its domestic industry products practice claim 14 of the '140 patent.

E. Invalidity

Respondents argue Claim 14 is invalid under 35 U.S.C. § 102(f) for improper inventorship and under 35 U.S.C. §§ 102 and 103 as anticipated and obvious. (RIB at 26.)

1. Anticipation – 35 U.S.C. 102

Respondents argue that under NVIDIA's infringement theory, including NVIDIA's proposed construction of the limitation "reciprocal instruction," Renderman on the Horizon860 anticipates claim 14 of the '140 patent. (RIB at 30-31; RRB at 27.) To that end, Respondents' expert, Mr. Diefendorff, testified in detail that Renderman on the Horizon860 meets all the limitations of claim 14.

NVIDIA argues there is no anticipation. (CIB at 60-65.) In particular, NVIDIA argues Respondents argument suffers from a failure of proof because Respondents did not show how the

PUBLIC VERSION

Renderman on Hoizon860 system actually worked. NVIDIA also argues the i860 processor is not a “hardware graphics accelerator.” NVIDIA further argues that the Renderman on Horizon860 does not perform “operations on graphics data” as required by claim 14 of the ‘140 patent.

I have not adopted NVIDIA’s construction of “reciprocal instruction.” (*See supra*, at IV.B.3.b.) Thus, the predicate to Respondents’ anticipation argument is not met. The evidence shows that under my construction, Renderman does not include the claimed “reciprocal instruction.” (RX-005C (Diefendorff) at Q&A 193, 195, 198; RX-2527 at 115, 121.) While a user could write programs using division and square root functions to calculate “1” divided by a number or “1” divided by the square root of a number in the same manner that a user can use division to calculate “1” divided by a number in OpenGL, the RenderMan shading language did not include a “reciprocal instruction” or a “reciprocal square root instruction” as the claim requires. (*Id.*) Thus, under my construction, Renderman on the Hoizon860 does not include all the limitations of claim 14. Accordingly, I find Respondents have failed to prove by clear and convincing evidence that Renderman on the Horizon860 anticipates claim 14 of the ‘140 patent.

NVIDIA also argues that Respondents did not provide any evidence to show how this system actually worked, arguing that as a result Respondents anticipation argument suffers from a lack of proof. (CIB at 66.) I disagree. Respondents actually cite three references that describe the Horizon860. (*See* RIB at 31-32 (citing RX-245; RX-250; RX-2527).) Respondents also cite references describing RenderMan and the i860. (RIB at 31-32 (citing RX-005C (Diefendorff) at Q&A 160-162,174; RX-245 at 1 (“MS-DOS version”); RX-250 at v; RX-255 at 1-6, 8-26, 8-30; RX-265 at 85, 87; RX-2587 at 2 (use of 33-Mhz i860 to accelerate 3D and other graphics applications)).) The evidence shows the Horizon860 was a “board with an i860 chip” that

PUBLIC VERSION

“included ... Pixar’s Photorealistic Renderman,” which “implement[ed]” the “Renderman Interface Specification.” (CIB at 66-67; *see also* RX-005C (Diefendorff) at Q&A 173; Tr. at 239:13-240:10.) NVIDIA argues it is unclear what parts of the Renderman Interface Specification Pixar’s Photorealistic Renderman implemented. But the evidence shows the Graphics Gems III book states that the Horizon860 supports the RenderMan Shading Language (the part of RenderMan relied on by Mr. Diefendorff (RX-005C (Diefendorff) at Q&A 185)), and used that language to generate the book’s cover image. (RX-250 at v; RX-005C (Diefendorff) at Q&A 169-73.)

NVIDIA further argues Renderman on the Horizon860 does not invalidate claim 14 because the i860 processor is not a hardware graphics accelerator. (CIB at 67.) NVIDIA argues the i860 is a general purpose CPU used for general purpose tasks such as running the operating system. (*Id.*)

NVIDIA argues the i860, by itself, is not a hardware graphics accelerator, but whether the i860 is a hardware graphics accelerator is immaterial because Respondents do not make such an argument. Rather Respondents argue, and I so find, the Horizon860 graphics board with the i860 processor clearly was a hardware graphics processor. (RX-005C (Diefendorff) at Q&A 171-81 (Horizon860 was an HGA under both proposed constructions; i860 included specialized graphics hardware); RX-250 at v (“The final rendering of this image was done on a 486 PC/DOS machine with Truevision’s RenderPak™ and Horizon860™ card containing 32 Mbytes of Ram.”).) In fact, NVIDIA senior distinguished engineer, NVIDIA expert witness, and inventor on the ‘140 patent, Dr. Moreton, admitted that “[i]n the context of when it was in the market,” the Horizon860 was a hardware graphics accelerator. (Tr. 149:18-21.) The whole colloquy with Respondents’ counsel is reproduced below.

PUBLIC VERSION

11 Q And I'd like you to be careful. So let's take
12 it step by step, then. The Horizon860 was an accelerator;
13 correct?

14 A Yes.

15 Q Okay. And the Horizon860 was a graphics board;
16 correct?

17 A Yes.

18 Q So it's fair to say that Horizon860 was a
19 hardware graphics accelerator; fair?

20 A In the context of when it was in the market,
21 yes.

22 Q In fact, Horizon860 was used for 3D graphics;
23 correct?

24 A Yes, it was used for 3D graphics.

25 Q Used an Intel i860 processor?

1 A I believe so.

2 Q The i860 processor was a single chip?

3 A Yes.

4 Q The i860 had SIMD graphics extensions?

5 A It did.

6 Q In other words, it had hardware support for
7 graphics functions; right?

8 A Yes.

9 Q The i860 had a separate graphics unit; right?

10 A The SIMD functions were in a separate unit, yes.

11 Q i860 had a separate graphics unit; right?

12 A Yes.

PUBLIC VERSION

13 Q i860 had hardware support for pixel shading?

14 A Very rudimentary, yes.

15 Q i860 had dedicated hardware for floating point
16 calculations?

17 A Yes.

18 Q And floating point calculations are something
19 that are useful for graphics; right?

20 A Yes.

21 Q And, in fact, the i860 was marketed for use in
22 graphics workstations; correct?

23 A Yes, in the context of the time, yes.

24 Q And the i860 was considered to be a graphics
25 processor; right?

1 A At that time, yes.

(*Id.* at 149:11-151:1.) Dr. Moreton's admission is also supported by the prior art's description of the Horizon860 as a 3D graphics "accelerator." (RX-245 at 1; *see also* RX-2587 at 2 ("Horizon860 board" used an "i860 to accelerate 3D and other graphics applications.")). Further, contrary to NVIDIA's argument, the evidence shows a separate CPU, the Intel 486, not the i860 on the Horizon860 board, ran the operating system. (RX-005C (Diefendorff) at Q&A 168-69; RDX-1552.) Thus, for at least the reasons above, I find the Renderman on the Horizon860 is a hardware graphics processor as that term has been construed herein. Accordingly, for the reasons above, I find NVIDIA's argument is not persuasive.

NVIDIA also argues Renderman on Horizon860 does not anticipate because Renderman does not disclose programmable processing of vertex data. (CIB at 68.) Thus, NVIDIA argues Renderman does not perform "operations on graphics data." (*Id.*) NVIDIA's argument is based

PUBLIC VERSION

on its proposed construction of the term “graphics data” as “data used for vertex processing.” However, as discussed *supra*, I have found “graphics data” to not be so limited and have construed the term to mean “data related to graphics.” (*See supra*, at IV.B.3.a.) Thus, under my construction of the term “graphics data” the data need not be limited to data used for vertex processing. In this regard the evidence shows the Renderman Shading Language could be used for many different “geometric transformations” like “special camera projections such as fish eye or IMAX, nonlinear deformations such as bends and twists, or surface displacement functions such as ripples or nubs.” (RX-005C (Diefendorff) at Q&A 187.) Moreover, the evidence shows RenderMan’s “transformation shaders transform a point in space to another point in space” and RenderMan expressly states that a “point” is a “vertex.” (RX-2527 at 59, 113; RX-005C (Diefendorff) at Q&A 187; RDX-1563.) Thus, contrary to NVIDIA’s argument, the evidence shows that the Renderman on Horizon860 does in fact process vertex data. Accordingly, for the reasons above, I find NVIDIA’s argument not persuasive and find the Renderman on Horizon860 performs “operations on graphics data” as claimed in claim 14 of the ‘140 patent.

2. Obviousness – 35 U.S.C. 103

a. Renderman on Horizon860

Respondents argue that under their construction of the limitation “reciprocal instruction” Renderman on Horizon860 combined with the state of the art renders obvious claim 14 of the ‘140 patent. (RIB at 36.) To that end, Respondents expert, Mr. Diefendorff, testified in detail that the Renderman on Horizon860 discloses each of the limitations of claim 14 of the ‘140 patent except the claimed reciprocal instruction and reciprocal square root instruction. (RX-005C (Diefendorff) at Q&A 164, 166-193, 201-211.) With regard to those instructions, Mr. Diefendorff testified that it would have been obvious to one of ordinary skill in the art to add a reciprocal instruction and a reciprocal square root instruction to the instruction set of the

PUBLIC VERSION

RenderMan Shading Language. (*Id.* at Q&A 198-200.) For example, with regard to the reciprocal instruction, Mr. Diefendorff testified:

In my opinion, it would be obvious to add a reciprocal instruction to the RenderMan Shading Language based on the general state of the art, RenderMan itself, and the i860.

A person of ordinary skill had a strong reason to provide a reciprocal instruction in a graphics-tailored instruction set. Before 1999, it was well-known that a reciprocal instruction was useful for graphics processing. In fact, reciprocals are commonly used in calculations in graphics, such as perspective division. As I explained earlier, various factors make reciprocal instructions fairly easy to implement in hardware, while division instructions are more time-consuming and require more silicon. Providing a reciprocal instruction in the RenderMan Shading Language would give programmers the opportunity to directly target the i860's native reciprocal instruction, letting them avoid a more costly division. In fact, in the i860, division is implemented using Newton-Raphson refinement of a reciprocal seed, followed by a multiply. Targeting the i860 reciprocal instruction when appropriate would enable the programmer to avoid considerable time-consuming work. Thus, providing a reciprocal instruction would provide predictable benefits and results.

Adding a reciprocal instruction to RenderMan would be a trivial exercise for any person of ordinary skill. It could be done in only a few lines of code by writing a new function, call it "rcp," and adding it to the library that defines the RenderMan Shading Language instruction set.

(RX-005C (Diefendorff) at Q&A 197; *see also id.* at Q&A 199-200 (Mr. Diefendorff testifying similarly with regard to the claimed reciprocal square root instruction.) Mr. Diefendorff's opinion that it would have been obvious to add reciprocal and reciprocal square root instructions is uncontested. In fact, with respect to Renderman on Horizon860, NVIDIA does not address Respondents' obviousness argument at all and thus under my Ground Rules has waived any such argument.³ And while the Staff asserts that RenderMan on the Horizon860 does not render obvious claim 14, the Staff also fails to address the argument at all. (*See* SIB at 23.)

³ NVIDIA does not address secondary considerations of nonobviousness with respect to Respondents' Renderman on Horizon860 obviousness argument. NVIDIA does, however, raise

PUBLIC VERSION

As discussed, *supra*, the evidence shows that Renderman on Horizon860 satisfies all the limitations of claim 14 except the limitations requiring a reciprocal instruction and a reciprocal square root instruction. NVIDIA's arguments in regards to Respondents' anticipation argument have been found not persuasive and NVIDIA does not address Respondents' obviousness argument at all. Thus, in light of Mr. Diefendorff's testimony, including his unrebutted obviousness testimony, and the reasons I have set forth, *supra*, I find by clear and convincing evidence that claim 14 is obvious in light of Renderman on the Horizon860 in view of the state of the art at the time.

b. "C" Language on TMS34082

Respondents argue that claim 14 is obvious in view of a C programming language, as extended by known graphics libraries, running on the prior art TMS34082 graphics processor ("the TMS chip"). (RIB at 37.) In that regard, Respondents expert, Dr. Diefendorff, testified in detail that claim 14 was obvious. (RX-005C (Diefendorff) at Q&A 248-67.) NVIDIA argues claim 14 is not obvious in view of the C programming language in combination with the TMS chip. (CIB at 68-69; SIB at 23.) In particular, NVIDIA and the Staff argue that Respondents have failed to show that the TMS chip is a "hardware graphics accelerator" as required by claim 14. (*Id.*) NVIDIA also argues that there is extensive evidence of secondary considerations of non-obviousness that cut against a finding of obviousness in this case. (CIB at 69.)

secondary considerations of nonobviousness with regard to Respondents' argument that the "C" programming language on the TMS34082 chip renders obvious claim 14 of the '140 patent. (*See infra*, at IV.E.2.b.(2).) Although NVIDIA has waived the issue, for completeness I note that NVIDIA's secondary considerations of non-obviousness cannot overcome the strong evidence that claim 14 was obvious in light of Renderman on Horizon860 in view of the state of the art at the time, because the products on which NVIDIA relies to show commercial success and widespread praise do not embody the invention of claim 14. (*See supra*, at IV.D.) Thus, NVIDIA's alleged evidence of secondary considerations is irrelevant.

PUBLIC VERSION

(1) Is the TMS34082 a “hardware graphics accelerator”

The Parties’ Positions

NVIDIA argues that like the i860, the TMS34082 is not a hardware graphics accelerator. (CIB at 68.) NVIDIA argues the TMS34082 is a 1991 co-processor with a general-purpose floating point unit (*e.g.*, additional math capability) and some pre-compiled graphics instructions. (*Id.*) NVIDIA argues it did not implement a graphics pipeline with specialized hardware. (*Id.*) NVIDIA argues that it merely provided a floating point unit that could be programmed to implement a graphics pipeline in the same way as a CPU. (*Id.*)

NVIDIA argues that like a CPU (and unlike a hardware graphics accelerator), the TMS34082 is generally programmable to perform any arbitrary function. (*Id.*) NVIDIA argues that to implement OpenGL or another graphics standard, a TMS34082 user must write a software implementation of the pipeline as they would with a CPU. (*Id.*) NVIDIA argues the ’140 patent is not directed at making a generally programmable piece of hardware programmable, but rather at adding a new feature to specialized hardware implementations of graphics pipelines. (*Id.*) NVIDIA argues that this distinction is reinforced by the fact that the TMS34082 is programmed with the C computer language. (*Id.*)

Respondents’ argue the TMS chip was a “Graphics Floating-Point Processor” that performed programmable operations on “graphics data,” including vertex data. (RIB at 38.) Respondents argue that it could be used in an “Accelerator Board” that interfaced with a CPU through the “Texas Instruments Graphics Architecture” (*Id.*) Respondents argue that the TMS chip satisfies even NVIDIA’s construction of an HGA as a “graphics processing unit.” (*Id.* at 39.) Respondents assert that Dr. Moreton admitted that the TMS chip “increased 3D graphics performance” and was “undoubtedly” considered “by some to be a graphics processor.” (*Id.*) According to Respondents, the TMS chip “enhances the throughput of the 3-D graphics pipeline

PUBLIC VERSION

and is tailored for transformation, clipping, and rendering operations.” (*Id.*) Respondents argue that it contained hardware specialized for graphics, including a floating-point unit, sequencer, and dedicated graphics instructions, such as color clipping, visibility testing, and vector math instructions. (*Id.*) Respondents contend NVIDIA’s witnesses incorrectly call the TMS chip’s native instructions “software,” but that the instructions were actually hardware “Built-In Silicon.” (*Id.* at 40.)

Respondents argue that like the Horizon860, NVIDIA’s primary criticism of the TMS chip is that an application writer had to write programs to invoke its graphics instruction set. (*Id.*) However, Respondents note that claim 14 requires a programmable instruction set. (*Id.*) Respondents argue that it is illogical and simply wrong for NVIDIA to argue that the TMS chip is not a hardware graphics accelerator precisely because it satisfies claim 14’s programmability requirement. (*Id.*) Respondents argue that the evidence shows the TMS chip’s hardware graphics instruction set could be used for each stage of the pipeline. (*Id.*)

The Staff argues that while the TMS chip could be used to build a hardware graphics accelerator, it is not by itself a hardware graphics accelerator.⁴ (SIB at 23.)

Discussion

As construed herein, a “hardware graphics accelerator” is “hardware for processing graphics.” Consistent with this construction and as discussed, *supra*, a CPU microprocessor can be employed to function in a system as a hardware graphics accelerator if it is assigned to offloading graphics work from the main CPU, or through the addition of new instructions and

⁴ But for this one conclusory statement by the Staff, the Staff does not independently address this issue. Rather, the Staff adopts NVIDIA’s argument.

PUBLIC VERSION

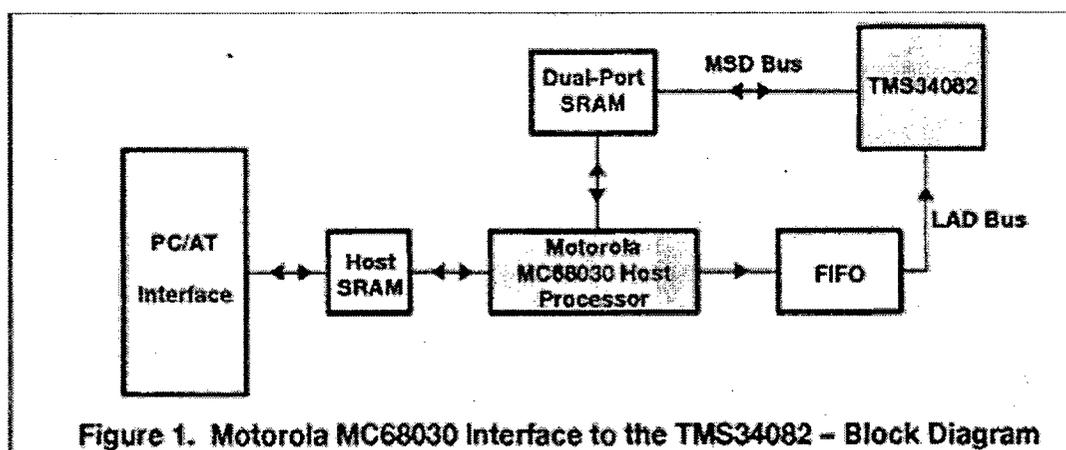
hardware execution units for the purpose of processing graphics data effectively. (*Supra*, at IV.B.3.a.; RX-3497C (Diefendorff) at Q&A 28.)

Here, the evidence shows the TMS chip was sold by Texas Instruments in early 1990 and was designed to be used as an independent processor or as a co-processor to another chip. (RX-005C (Diefendorff) at Q&A 229.) The evidence shows it was primarily designed as a co-processor to TI's TMS34020 Graphics System Processor for accelerating floating-point graphics calculations. (*Id.* at Q&A 223, 227, 229, 232.) The evidence shows the TMS chip was a "Graphics Floating-Point Processor" that contained hardware specially designed to accelerate the processing of graphics data. (*Id.* at Q&A 229, 232; RX-2774 at 1-1.) NVIDIA's expert, Dr. Moreton admitted at the hearing that the field considered the TMS chip to be a "graphics processor" that could be "programmed to perform graphics processing" and "increased 3D graphics performance." (Tr. at 151:7-19.) The TMS chip contained a floating-point core with a sequencer that could execute microcoded graphics-oriented instructions. (RX-005C (Diefendorff) at Q&A 233.) The evidence shows the microcode was etched into on-chip hardware read-only memory and used to implement the graphics pipeline. (*Id.* at Q&A 236, 238 ("These instructions are implemented internally in hardware using microcode and a microsequencer."), 239; RX-2774 at B-70 ("Internal microcode to the TMS34082 is not restricted to the same 32-bit instruction formats so certain internal programs may execute faster than the same operations written with external code can achieve."); RX-233 at 7, 9, 10.) The evidence also shows the TMS34082 was used to accelerate the 3D graphics pipeline and was "tailored for transformation, clipping, and rendering operations." (*Id.* at Q&A 233; RX-233 at 7.) The evidence shows the TMS34082 could also perform, among other things, backface testing, 3-

PUBLIC VERSION

D compares, interpolation, and reflection. ((RX-005C (Diefendorff) at Q&A 236; RX-233 at 8-9.)

As discussed above, the TMS chip could be used as a graphics co-processor. In such a configuration, the evidence shows the host processor would run the operating system and would offload graphics work to the TMS34082 to perform graphics. (RX-005C (Diefendorff) at Q&A 234.) For example, the TMS34082 Designer's Handbook states that "the TMS34082 floating-point processor can be coupled to a Motorola MC68030 microprocessor." (RX-2774 at D-1.) In this configuration, as shown in the figure below, the Motorola MC68030 is the host processor and the TMS chip is the graphics co-processor. (RX-005C (Diefendorff) at Q&A 234; RX-2774 at xii (referring to this configuration as "TMS34082 Accelerator Board").)



(RX-2774 at D-4.)

As described above, the evidence shows the TMS34082 had specialized hardware designed specifically to accelerate graphics processing, particularly floating-point calculations. The evidence shows the TMS34082 contained a floating-point core with a sequencer that could execute microcoded graphics-oriented instructions. The evidence shows these microcoded instructions were imbedded in hardware (*i.e.*, ROM) on the chip. The evidence also disclosed how these hardware instructions implemented the graphics pipeline. Moreover, contrary to the

PUBLIC VERSION

Staff's assertion, the evidence shows an actual configuration with a Motorola MC68030 processor and TMS chip where the Motorola processor serves as the host processor and the TMS chip as a graphics co-processor. In such a configuration there can be no doubt the TMS chip is "hardware for processing graphics."

Accordingly, for at least the reasons above, I find the TMS chip a "hardware graphics accelerator" as claimed in '140 patent.

(2) Secondary considerations of nonobviousness

The Parties' Positions

NVIDIA argues the non-obviousness of the '140 patent is supported by the commercial success and widespread praise of NVIDIA products embodying the invention. (CIB at 69.) NVIDIA asserts the '140 patent exposed the GPU's vertex processing functionality and allowed users to write shading programs to more closely control the shape, appearance, and motions of the objects. (*Id.*) NVIDIA argues that when its GeForce 3 GPU was released in 2001, it was widely recognized for these features. (*Id.*) NVIDIA asserts that it satisfied a long felt need for additional flexibility and programmability in hardware graphics accelerators. (*Id.*) NVIDIA argues that its later releases, including the Tegra 4, continued to enable these "graphics capabilities in mobile devices, and far surpasses the competition, as seen in many of the top mobile GPU benchmarks." (*Id.*) NVIDIA asserts that the inventive programmable architecture has been adopted by other GPU manufacturers – including the manufacturers of the accused GPUs – as well as the APIs used to program them. (*Id.*)

Respondents argue that near-simultaneous invention of claim 14 confirms its obviousness. (RIB at 43.) Respondents argue that in addition to other prior art, GigaPixel, a competitor to NVIDIA, independently invented a hardware graphics accelerator with a programmable instruction set including eight of the nine claimed instructions by March 2000.

PUBLIC VERSION

(*Id.*) Respondents argue that NVIDIA has no response to this evidence. Respondents contend NVIDIA's secondary factors all lack a nexus to claim 14. (*Id.*) Respondents argue that the alleged industry praise noted by NVIDIA merely touts the well-known benefits of programmability, not claim 14's instruction set. (*Id.*) Likewise, Respondents argue NVIDIA does not show claim 14 satisfied a long-felt need or that others have adopted the invention of claim 14. (*Id.*) In fact, Respondents argue NVIDIA has failed to show the alleged success of its products has any connection to the specific instructions of claim 14. (*Id.*)

The Staff does not address NVIDIA's secondary considerations.

Discussion

NVIDIA argues the non-obviousness of the '140 patent is supported by the commercial success and widespread praise of NVIDIA products embodying the invention. However, as discussed *supra*, the products on which NVIDIA relies do not embody the invention of claim 14 as they do not disclose the claimed instructions from the claimed predetermined instruction set. (*See supra*, at IV.D.) Thus, I find NVIDIA has failed to show a nexus between its assertions of commercial success and widespread praise and the claim at issue. *See e.g., Ormco Corp. v. Align Tech., Inc.*, 463 F.3d 1299, 1311–12 (Fed. Cir. 2006) (“Evidence of commercial success, or other secondary considerations, is only significant if there is a nexus between the claimed invention and the commercial success.”) Accordingly, I find NVIDIA's argument with respect to the secondary considerations of nonobviousness not persuasive.

(3) Conclusion

NVIDIA's sole defense is that the TMS chip was not a “hardware graphics accelerator,” but as discussed in detail above the evidence clear shows that it was. NVIDIA does not dispute Mr. Diefendorff's convincing testimony that the combination of the “C” programming language and TMS chip discloses every other limitation of claim 14 of the '140 patent, including every

PUBLIC VERSION

claimed instruction under its construction of reciprocal or the construction that I have adopted herein. Nor does NVIDIA contest it would have been obvious to one of ordinary skill to run the C graphics library on the TMS chip. And while NVIDIA does present alleged evidence of secondary considerations of nonobviousness, as I discussed, supra, there is no nexus between the alleged secondary considerations and claim 14. Thus, the evidence is of no significance. Accordingly, for at least the reasons above, I find the evidence clearly and convincingly establishes that claim 14 of the '140 patent was obvious in light of the combination of the "C" programming language on the TMS34082.

3. Inventorship – 35 U.S.C. 102(f)

The Parties' Positions

NVIDIA argues the evidence shows that Respondents' "do not come close to meeting their burden to prove co-inventorship with clear and convincing evidence." (CIB at 70.)

NVIDIA asserts that Mr. Boyd does not claim to be an inventor and that Respondents cancelled his deposition and submitted no testimony from him about the invention of the '140 patent or anything else. (*Id.*) Instead, NVIDIA argues Respondents premise their non-joinder defense solely on a few lines of text taken out-of-context from Mr. Boyd's February 12, 1999 e-mail and an acknowledgement section in an after-the-fact conference paper describing the commercial embodiment of the '140 Patent invention. (*Id.*)

NVIDIA asserts that it submitted extensive testimony from two actual inventors, detailing how they conceived and developed the '140 Patent without any input from Mr. Boyd. (*Id.*) NVIDIA contends that both inventors specifically disputed that Mr. Boyd was a co-inventor. (*Id.*) NVIDIA asserts that Respondents declined to cross-examine Mr. Lindstrom, one of the testifying inventors, and asked the second testifying inventor only a handful of questions about Mr. Boyd, which did not establish non-joinder. (*Id.*) NVIDIA asserts that it submitted extensive

PUBLIC VERSION

testimony explaining why Mr. Boyd's e-mail was a feature request from Microsoft that provided no definite, concrete guidance on how to achieve a programmable hardware graphics accelerator.

(Id.)

Respondents argue the '140 patent, including claim 14, is invalid under § 102(f) for failing to name Mr. Charles ("Chas") Boyd of Microsoft as a co-inventor for his significant contributions to at least claim 14. (RIB at 26.) Respondents assert that the basic facts of Mr. Boyd's contributions are undisputed. *(Id.)* Respondents assert that two months before NVIDIA's earliest alleged conception, Mr. Lindholm and Dr. Moreton received Mr. Boyd's email disclosing nearly all of claim 14. *(Id.)* Respondents argue that the evidence shows Mr. Papakipos explained to Mr. Lindholm and others that Chas Boyd of Microsoft had the idea of a system for "programmable vertex operations, and per-pixel operations." *(Id. at 27.)*

Respondents argue that Mr. Boyd described a hardware graphics accelerator that would execute these operations in the form of "parallel mimd units." *(Id.)* Respondents contend that at trial, Dr. Moreton admitted that a MIMD unit could execute the instructions disclosed by Mr. Boyd and could be used in a hardware graphics accelerator. *(Id.)* Respondents argue that this evidences conception and contribution of claim 14's "hardware graphics accelerator for receiving graphics data, and performing programmable operations on the graphics data in order to generate output." *(Id.)*

Respondents argue that Mr. Boyd also described to NVIDIA a predefined instruction set, programmable by a user that includes six of the instructions claimed in claim 14. *(Id. at 28.)* Specifically, Respondents argue Mr. Boyd disclosed "vertex shaders," and an API for using those shaders, that provided (1) a three component dot product instruction and (2) four component dot product instruction. *(Id.)* Respondents argue that Mr. Boyd also described

PUBLIC VERSION

instructions for performing operations on texture and pixel data, including (3) a reciprocal instruction, (4) a reciprocal square root instruction, (5) an exponential instruction, and (6) a logarithm instruction. (*Id.*) Respondents assert that Dr. Moreton admitted that these elements constitute a “significant portion” of claim 14. (*Id.*)

Respondents argue that NVIDIA’s responses to this evidence of Mr. Boyd’s conception and contribution to the ’140 patent are irrelevant to the issue of joint inventorship. (*Id.*)

Respondent assert that NVIDIA argues that in later correspondence Microsoft used parts of NVIDIA’s design documents for its DirectX 8 specification, and that Mr. Boyd’s email describes additional features not claimed in the ’140 patent. (*Id.*) Respondents argue this evidence does nothing to undermine the express evidence that Mr. Boyd’s February 1999 email disclosed a significant portion of claim 14 to NVIDIA. (*Id.* at 28-29.) Respondents assert that it is tellingly that NVIDIA does not argue that it conceived of the invention before Mr. Boyd’s email. (*Id.* at 29.)

Respondents argue that contrary to NVIDIA’s argument, Mr. Boyd’s email was not just a mere “request” for further development, but rather discloses extensive details, including a programming model for using the instructions within the DirectX application program interface. (*Id.*) Respondents assert that Mr. Boyd’s email provides more detail on his programming model and hardware graphics accelerator than NVIDIA’s own conception documents and the ’140 patent itself. (*Id.*) Respondents also argue Dr. Moreton conceded at trial that Mr. Boyd wanted to “collectively develop a programmable vertex capability” and that his email included “musings about things that it might include and behaviors it might have.” (*Id.*) Respondents contend that given those admissions and the clear disclosure of the Boyd email, more than clear and

PUBLIC VERSION

convincing evidence exists showing that Mr. Boyd contributed to the conception of at least claim 14 of the '140 patent. (*Id.*)

Respondents contend Mr. Boyd's significant contribution to the claimed invention of the '140 patent requires he be named a co-inventor. (*Id.*) Respondents argue NVIDIA's failure to name him as such invalidates the entire patent. (*Id.*)

The Staff believes that the evidence shows that claim 14 of the '140 patent is invalid under 35 U.S.C. § 102(f) for failure to name Mr. Boyd of Microsoft as an inventor, and therefore unenforceable. (SIB at 23.) The Staff argues that the documentary evidence of record clearly shows that Mr. Boyd contributed to the conception of the claimed invention. (*Id.*) The Staff asserts that months before NVIDIA's alleged conception of claim 14 of the '140 patent, Mr. Boyd sent NVIDIA his idea for a hardware graphics accelerator providing programmable vertex and pixel operations. (*Id.* at 23-24.) The Staff asserts that Mr. Boyd's ideas were received by Erik Lindholm, a named inventor of the '140 patent, no later than March 2, 1999. (*Id.* at 24.) The Staff argues that Mr. Boyd proposed a hardware graphics accelerator for DirectX 8 that performed programmable vertex operations. (*Id.*) The Staff argues that in his email, Mr. Boyd provided substantial details related to the instruction sets that would be supported, including identification of six of the nine required instructions ultimately claimed by the '140 patent. (*Id.*) Thus, the Staff argues that Mr. Boyd made substantial contribution to the conception of claim 14 of the '140 patent.

The Staff argues that NVIDIA's argument that Mr. Boyd's ideas simply reflected well-known principles is unconvincing. (*Id.*) The Staff argues that Mr. Boyd clearly did far more than suggest a pie-in-the-sky idea that he did not know how to implement. (*Id.*) The Staff argues he provided concrete ideas that he envisioned for Microsoft DirectX 8, including an

PUBLIC VERSION

outline of instructions programmable features that would be implemented. (*Id.*) The Staff contends that in response, NVIDIA offers only the uncorroborated testimony of NVIDIA employees, and conclusory opinions of its experts. (*Id.* at 25.) The Staff argues such testimony is entitled to little weight. (*Id.*)

The Staff believes that single most reasonable inference to be drawn from the evidence adduced at the hearing is that Mr. Boyd made substantial contributions to the conception of claim 14 of the '140 patent, and is thus an unnamed inventor. (*Id.*) Accordingly, the Staff argues claim 14 is invalid under 35 U.S.C. § 102(f). (*Id.*) The Staff argues that since the Commission has no power to correct inventorship, the Commission must determine that the '140 patent is currently unenforceable. (*Id.*) Moreover, the Staff argues that as an unnamed inventor, Mr. Boyd and/or Microsoft owns an undivided common interest in the '140 patent, and thus NVIDIA lacks all substantial rights in the patent. (*Id.*) The Staff argues Mr. Boyd and/or Microsoft are therefore necessary parties to this investigation. (*Id.*)

Discussion

Inventorship of a patent is presumed to be correct. *Hess v. Advanced Cardiovascular Sys., Inc.*, 106 F.3d 976, 980 (Fed. Cir. 1997). Overcoming this presumption to show that a patent is defective under 35 U.S.C. §102(f) for non-joinder requires proof by clear and convincing evidence. *Eli Lilly & Co. v. Aradigm Corp.*, 376 F.3d 1352, 1358 (Fed. Cir. 2004); *Certain Microprocessors, Components Thereof, And Products Containing Same*, Inv. No. 337-TA-781, I.D., 2012 WL 6883205, at **159-60 (Dec. 14, 2012). A review of all of the evidence must yield an “abiding conviction that the truth of a factual contention is ‘highly probable.’” *Certain Elec. Digital Media Devices And Components Thereof*, Inv. No. 337-TA-796, I.D., 2012 ITC LEXIS 2638, at *33 (Oct. 24, 2012).

PUBLIC VERSION

Respondents allege that the '140 Patent is invalid under 35 U.S.C. §102(f) “for failing to name Mr. Charles (‘Chas’) Boyd of Microsoft as a co-inventor for his significant contributions to claim 14.” According to Respondents, before the inventors listed on the '140 Patent conceived of claim 14, Mr. Boyd sent an e-mail (RX-2823C) to NVIDIA “disclosing nearly all of claim 14,” including the concept of “programmable vertex and pixel operations performed on [a hardware graphics accelerator]” using “a predefined instruction set containing six of claim 14’s nine instructions.” Respondents argue this disclosure made Mr. Boyd “at least [an unnamed] co-inventor of the '140 Patent.” The Staff agrees.

I do not find Respondents have established, by clear and convincing evidence, that Mr. Boyd should have been named as a co-inventor for Claim 14 of the '140 patent. The only arguably probative non-opinion evidence I find Respondents have presented is a 1999 e-mail in which Mr. Boyd provided a list of features he would like NVIDIA to develop before the conception date. (RX-2823C at 702468; Tr. (Moreton) at 159:12-21.) The problem with the email, is that in isolation, it provides no information as to whether Mr. Boyd originated any of the concepts in Claim 14. Instead, one very possible meaning I can take from the entire email is that Mr. Boyd was making suggestions on how NVIDIA might develop the technologies (a programmable shader) useful to Microsoft, his employer. (Tr. (Moreton) at 169:12-14.)

I do not find the fact that the named inventors thanked Mr. Boyd, among others in a 2001 conference paper to be probative. Instead, the meaning of thanking Mr. Boyd is just too ambiguous. As NVIDIA cogently argues, there were others that were thanked, yet Respondents accuse none of these as being omitted inventors. (CIB at 76.)

I find the unchallenged and unrebutted testimony of Mr. Lindstrom and Dr. Moreton on the question of Mr. Boyd’s role and the relevance of what Mr. Boyd suggested in his e-mail to be

PUBLIC VERSION

convincing. Both of these named inventors testified Mr. Boyd had no role in devising or developing any of the innovative technology of the '140 patent and that the desires Mr. Boyd expressed for programmable vertex processing were known in the industry. (CX-2543C (Lindholm) at Q&A 3-10; CX-2133C (Moreton) at Q&A 39-40.) Mr. Lindholm even explained that Mr. Boyd's suggestions were not particularly helpful, sentiments echoed in a different manner by Dr. Moreton. (CX-2543C (Lindholm) at Q&A 9-10; CX-2133C (Moreton) at Q&A 39-40.) In short, I find their testimony that Mr. Boyd contributed very little of substance, if anything at all, to be credible.

I further note that one of the most convincing witnesses presented by either party, Mr. Charles Diefendorff, (perhaps to partly explain why Claim 14 was obvious), opined that the instructions contained in Claim 14 were known for graphics processing and thus Mr. Boyd could not have invented them. (RX-005C at Q&A 321; *see also* CIB at 75.) This testimony is rather consistent with that of Mr. Lindholm and Dr. Moreton. Moreover, the only meaning I can take from Mr. Diefendorff's testimony (on this particular point) is that Mr. Boyd could not have invented the features Respondents claim he did. Instead, the hard part or the key to the invention, as argued by NVIDIA, is making the instructions programmable by permitting the programmer to have access to the granular operations performed by the GPU during vertex processing and Mr. Boyd had no role in that according to Dr. Moreton. (*See* CX-003C (Lindholm) at Q&A 95-96; CX-2133C at Q&A 16 (Moreton).) Hence, Mr. Diefendorff's testimony is inimical to Respondent's inventorship arguments, regardless about what he opined about what Mr. Boyd's email allegedly disclosed.

Mr. Boyd has not testified. Since his failure to testify can be reasonably argued to be adverse to both parties, I find his failure to testify to be probative of nothing.

PUBLIC VERSION

Instead of agreeing with Respondents' claim that Mr. Boyd was a co-inventor, I find the evidence indicates it is more likely than not that Mr. Boyd made suggestions, on behalf of Microsoft, as a possible business partner to NVIDIA, on features he (Microsoft) would desire. (See CX-2133 (Moreton) at Q&A 39.) This is not an inventorship. See *Garrett Corp v U.S.*, 422 F.2d 874, 881 (Ct. Cl. 1970); *Shatterproof Glass Corp. v. Libbey-Owens Ford Co.*, 758 F.2d 613, 624 (Fed. Cir. 1985); and *Morgan v. Hirsch*, 728 F.2d 1449, 1452 (Fed. Cir. 1984).

V. U.S. PATENT NO. 6,690,372

U.S. Patent No. 6,690,372 ("the '372 patent") is titled "System, Method and Article of Manufacture for Shadow Mapping." (JX-002). The '372 patent issued on February 10, 2004 on an application filed on December 5, 2000. The '372 patent claims priority to U.S. Patent No. 6,532,013 ("the '013 patent"), which was filed on May 31, 2000. The '372 patent incorporates by reference the entire specification of the '013 patent. (JX-002 at 1:9-13.) The '372 patent lists Walter E. Donovan and Liang Peng as inventors. (*Id.*) There are 25 claims. (*Id.*) NVIDIA alleges the Accused Products infringe claims 23 and 24. NVIDIA relies upon claim 23 for purposes of the proving the technical prong of the domestic industry requirement.

The '372 patent generally relates to techniques for calculating the effects of shadows in computer graphics pipelines. (JX-002 at 1:24-25.) The patent states:

A major objective in graphics rendering is to produce images that are so realistic that the observer believes the image is real... One important way to make images more realistic is to determine how objects in a scene cast shadows and then represent these shadows in the rendered image. Shadows enhance the realism of an image because they give a two-dimensional image a three-dimensional feel.

(JX-002 at 1:28-38.)

Claim 24 depends from claim 23, which depends from claim 22, which depends from independent claim 21. Claims 21, 22, 23, and 24 read as follows:

PUBLIC VERSION

Claim 21. A system for performing shading calculations in a graphics pipeline, comprising:

- (a) logic for performing a first shading calculation in order to generate output;
 - (b) logic for saving the output; and
 - (c) logic for performing a second shading calculation using the output in order to generate further outputs;
- wherein the first and second shading calculations together include a plurality of decoupled variables.

Claim 22. The system as recited in claim 21, wherein the system includes a shading module for performing the first shading calculation in order to generate the output.

Claim 23. The system as recited in claim 22, wherein the system includes a texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output.

Claim 24. The system as recited in claim 23, wherein the system includes a feedback loop coupled between an input and an output of the shading module for performing the second shading calculation using the texture information from the texture look-up module in order to generate further output.

(372 patent at 14:29-52, 16:30-54.)

A. Level of Ordinary Skill in the Art

A person of ordinary skill in the art of the '372 patent would have “at least a four-year degree in Electrical Engineering, Computer Engineering, Computer Science, or equivalent, as well as at least two years of experience in graphics processing including developing, designing or programming software or hardware for graphics processing units, hardware graphics accelerators or other graphics processing systems.” (Order No. 20: Construing Terms of the Asserted Patents (“*Markman* Order”) at 62 (April 2, 2015).)

B. Claim Construction

1. Order No. 20: Construing Terms of the Patent - “shading calculation”

In Order No. 20, I construed the phrase “shading calculation” to mean “a computation of a value concerning the appearance of a surface.” (*Markman* Order at 72.)

PUBLIC VERSION

2. Disputed Term - “texture look-up module coupled to the shading module” (Claim 23)

The Parties’ Positions

NVIDIA argues that the heart of the parties’ dispute is whether claim 23 requires the “shader-to-texture” configuration described in the ’372 Patent. (RRB at 9.) NVIDIA contends that Respondents ignore the context of the claim, including the bulk of the claim language and the specification, in order to argue that the plain and ordinary meaning of “coupled” – in isolation – does not require any particular direction. (*Id.*) NVIDIA argues that in contrast, it considers the entire claim element in context to recognize that a person of ordinary skill would understand claim 23 to require the novel “shader-to-texture” configuration described in the ’372 Patent. (*Id.* at 10.) NVIDIA argues that with the ’372 Patent, NVIDIA reorganized the pipeline to a “shader-to-texture” configuration that added coupling where the shader module generates texture coordinates that are sent to a texture fetch module. (*Id.*) NVIDIA asserts that this is in the claim language and in the specification. (*Id.*)

NVIDIA argues that Respondents’ construction of “coupled” in the abstract, divorced from the rest of the claim and the specification, represents a fundamental misunderstanding of claim construction law. (*Id.*) NVIDIA argues that contrary to Respondents’ approach, proper claim construction requires taking the specification and the surrounding claim language into account. (*Id.* at 11.)

NVIDIA contends that contrary to Respondents’ arguments, NVIDIA is not importing limitations from the preferred embodiment. (*Id.*) NVIDIA argues that it is simply reading the term “coupled” in the context of the surrounding claim language, as required by the Federal Circuit. (*Id.*) NVIDIA argues that the invention requires coupling whereby the shader can send texture coordinates to a texture look-up module (*i.e.*, the claimed shader-to-texture

PUBLIC VERSION

configuration). (*Id.*) NVIDIA asserts this is expressly stated in the claims. (*Id.*) NVIDIA argues that while in the abstract the term “coupled” may not specify a direction, the language of claims 23 clearly does. (*Id.*) Specifically, NVIDIA argues the claim states that the shading module generates “the output,” and the coupling enables the texture look-up module to “us[e] texture coordinates *associated with the output*” to “retrieve texture information.” (*Id.*) In other words, NVIDIA argues, the claim discloses that the texture module is coupled to the shading module for a purpose. (*Id.* at 11-12.) NVIDIA argues that purpose is to “retriev[e] texture information using texture coordinates associated with the output.” (*Id.* at 12.) NVIDIA contends that this requires the “shader-to-texture” configuration. (*Id.*) According to NVIDIA, the texture look-up module cannot use texture coordinates associated with the output if it does not receive the texture coordinates, or at least the output they are associated with, from the shading module. (*Id.*)

NVIDIA asserts that because Respondents interpret “coupled to” in a vacuum, their construction is inconsistent with the rest of the claim. (*Id.* at 13.) NVIDIA argues that under Respondents’ proposed construction, the texture look-up module cannot receive “texture coordinates associated with the output,” and therefore cannot use such coordinates. (*Id.*) NVIDIA asserts that Respondents, in an attempt to account for this inconsistency, argue that *any* texture coordinates used by the texture look-up module are “associated with the output,” even if they do not come from the shading module. (*Id.*) But NVIDIA argues Respondents’ own expert conceded on cross-examination that the “output” referred to in claim 23 is the output of the shading module. (*Id.*) NVIDIA argues that Respondents’ construction is nonsensical as it would mean that the claimed “texture coordinates” can be “associated” with a shading calculation that has not yet occurred and an output does not yet exist. (*Id.*) Moreover, NVIDIA asserts the

PUBLIC VERSION

specification is clear that the texture operations are “a function of” the shading calculation, meaning they depend on the output. (*Id.* at 14.) NVIDIA then maintains the preferred and only embodiments of the hardware graphics pipeline of claims 23 and 24 are disclosed in Figure 4 and its accompanying description. (*Id.*) According to NVIDIA, these embodiments show the novel “shader-to-texture” configuration in which the shading module sends texture coordinate information to the texture look-up module. (*Id.*) NVIDIA argues that all of the embodiments disclosed that the shading module and the texture module are coupled in both directions. (*Id.*) NVIDIA argues that Respondents’ proposed construction is completely divorced from the specification. (*Id.* at 17.) NVIDIA argues that the ‘372 patent clearly disclaims the single direction (*i.e.*, linear) configuration proposed by Respondents. (*Id.*) NVIDIA asserts that in light of this disclaimer, one of ordinary skill in the art would understand claim 23 to require the novel shader-to-texture configuration. (*Id.*)

Respondents argue that “coupled” should be construed according to its plain meaning to allow a connection in one or more directions. (CIB at 46.) Respondents assert that to avoid invalidity, NVIDIA seeks to limit the claims to what is shown in one exemplary figure in the patent. (*Id.*) Specifically, Respondents assert that NVIDIA argues “coupled,” in the context of the texture look-up module only, requires coupling “in two directions,” *i.e.*, bidirectional coupling, and that the shading module “send texture coordinates to the texture module and receive filtered texture colors back—all during a single pass through the graphics pipeline.” (*Id.*) Respondents argue NVIDIA’s construction should be rejected because it is at odds with the plain language of the claims and an improper attempt to limit the claims to a preferred embodiment. (*Id.* at 46-47.)

PUBLIC VERSION

Respondents argue that NVIDIA's construction of "coupled" in the context of the texture look-up module is inconsistent with the ordinary meaning of the term. (*Id.* at 47.) Respondents also argue that it is inconsistent with the term's use in other places in the claims. (*Id.*)

Respondents assert that Dr. Aliaga admitted at the hearing that the term "coupled" should be given its ordinary meaning. (*Id.*) Respondents argue the plain meaning of the term allows a connection in one or more directions. Respondents argue that their expert Dr. Crawfis explained that "coupled" "would be either going one way or two ways." (*Id.*) Respondents also contend that Dr. Aliaga conceded that the plain meaning of "coupled" "in the context of the '372 Patent" is that "two elements are coupled when on[e] element can send a signal to the other." (*Id.*) Respondents also argue that the patent itself describes one-way connections as "coupled." (*Id.*)

Respondents argue there is nothing in the intrinsic evidence that would indicate that the term "coupled" should not be given its plain meaning. (*Id.*) Respondents assert that NVIDIA argues that because the texture coordinates must be "associated with" the output, the shader module must send its output to the texture look-up module, and that somehow requires bidirectional coupling. (*Id.*) But, Respondents argue, there is no such requirement in the claim; it only requires that the shading calculation output be "associated with" texture coordinates. (*Id.*)

Respondents also contend NVIDIA improperly relies on Figure 4 of the '372 patent to try to limit the claims and, in doing so, commits the "cardinal sin" of reading a preferred embodiment into the claims. (*Id.* at 48.) Respondents argue that the Figure 4 embodiment is the only place in the '372 patent that describes coupling in two directions between the shader module and texture module, where the shader may send in a single pass texture coordinates to the texture module. (*Id.*) But, Respondents argue, that figure is expressly described as "one embodiment" of the invention. (*Id.*) Respondents argue the patent consistently emphasizes the

PUBLIC VERSION

exemplary nature of Figure 4. (*Id.*) Respondents argue that the patent states that texture coordinate information “may be sent to the texture look-up module 408,” but does not state that the shading module and texture look-up module must be coupled in two directions and occur in a single pass. (*Id.* at 48-49.) Respondents also assert the specification provides embodiments where the output is recirculated through the graphics pipeline. (*Id.* at 49.) In these embodiments, Respondents argue the texture look-up module receives texture coordinates via other components, not directly from the shading module. (*Id.*)

Respondents further argue that NVIDIA’s disclaimer argument is without merit. (*Id.*) Respondents assert that NVIDIA argues that in an incorporated patent, U.S. Patent 6,532,013, applicants disclaimed a “linear pipeline” that uses multipass rendering, and that this limits the construction of “coupled” in the ’372 patent claims to coupling in “two directions.” (*Id.*) Respondents argue there is no legal basis to read a disclaimer into claim 23 of the ’372 patent. (*Id.*) According to Respondents, the ’013 patent is a different patent, with different claims, and with one exception, different named inventors. (*Id.*) Further, Respondents assert NVIDIA’s disclaimer argument relies on a prior art figure in the ’013 patent that was deleted from the ’372 patent. (*Id.*) Respondents argue that there is no discussion of the alleged non-linear pipelines, or single-pass rendering, or the exclusion of multipass rendering in either the ’372 patent or its prosecution history. (*Id.*) On the contrary, Respondents argue that Figure 4 of the ’372 patent shows one-way coupling and the patent discloses an embodiment involving multipass rendering. (*Id.*) Respondents thus argue there has been no clear disavowal of ’372 patent claim scope. (*Id.*)

Respondents argue in sum that there is no basis for NVIDIA’s new argument that allowing “the shader module to make a request to the texture module, and then deliver information back” “was the critical invention that made claim 23 be allowed.” (*Id.*)

PUBLIC VERSION

Respondents argue that during the Markman proceedings, NVIDIA identified a completely different structure—the feedback loop of dependent claim 24—as an allegedly “important” part of the invention. (*Id.*) Respondents assert that NVIDIA now abandons this position because Respondents’ prior art discloses several feedback loops. (*Id.*) Respondents argue that NVIDIA’s new argument fares no better, and its unduly narrow “coupled” construction should be rejected. (*Id.*)

Discussion

In its pre-hearing brief, NVIDIA asserted that the construction of the phrase “texture look-up module coupled to the shading module ...” was in dispute. (*See* CPHB at 81-84.) NVIDIA interpreted this phrase to require the claimed coupling to be in “both directions.” (*Id.* at 46-47; *see also* CX-2127C (Aliaga) at Q&A 123.) In its initial post-hearing brief NVIDIA appears to recast this construction as a so-called “shader-to-texture configuration.” The term “shader-to-texture” is a term made up by NVIDIA and used for the first time in its post-hearing brief. The term was never considered by its expert and cannot be found anywhere in the ‘371 patent. In its post-hearing reply brief, NVIDIA provides some clarity to its initial brief, stating:

Respondents contend that NVIDIA seeks to import “bi-directional” into the term “coupled,” but that mischaracterizes the issue. The issue is whether the language of claim 23, when read as whole in the context of the patent, requires “shader-to-texture” coupling. It does.

(CRB at 33.) Thus, it appears the dispute remains over the proper construction of the term “coupled.” To that end the parties’ positions are relatively straight forward. NVIDIA’s position appears to be that the connection (*i.e.*, coupling) between the texture look-up module and the shading module must be such that the texture look-up module can receive texture coordinate information from the shading module. Thus, according to NVIDIA’s position the coupling must allow data to flow from the shading module to the texture look-up module. Respondents’ and

PUBLIC VERSION

the Staff's position appears to be that the term "coupled" should get its plain and ordinary meaning and that the plain and ordinary meaning is satisfied as long as there is a connection between the texture look-up module and the shading module. Thus, according to Respondents' position, the claim language would be satisfied regardless of whether the connection permitted data to flow from the texture look-up module to the shading module or from the shading module to the texture look-up module.

NVIDIA's construction of "coupled" in the context of the texture look-up module is inconsistent with the ordinary meaning of the term, including its use in other places in the claims. The plain meaning of the term allows a connection in one or more directions. Respondents' expert Dr. Crawfis explained that "coupled" "would be either going one way or two ways." (Tr. at 921:25-922:4; RX-3495C (Crawfis) at Q&A 36.) NVIDIA's expert, Dr. Aliaga, also conceded that "[g]enerally, two elements are coupled when one [sic] element can send a signal to the other." (CX-2127C (Aliaga) Q&A 66; Tr. at 251:14-16 ("coupled' just means one or two-way communication.")) Moreover, Dr. Aliaga admitted that "[t]he term itself does not necessarily imply a direction." (CX-2127C (Aliaga) at Q&A 68.) Even the '372 patent itself describes one-way connections as "coupled." For example, both claims 10 and 25 require a "combiner module coupled to the output of the shading module" and Figure 4 clearly shows the combiner module (410) connected to the shading module (406) with a one-way connection. (See JX-002, claims 10, 25, Fig. 4.)

NVIDIA argues that the claim language itself requires "the texture look-up module receives and uses texture coordinates from the shading module." (CIB at 10.) But as discussed above, there is no dispute that the plain and ordinary meaning of the term "coupled" does not require any "directionality" in the connection between the shading module and the texture look-

PUBLIC VERSION

up module. Moreover, the fact that texture coordinates are “associated with” the output does not add directionality to the connection or require the texture look-up module to “receive and use” the output. (RX-002C (Crawfis) at Q&A 192.) As Dr. Crawfis testified, texture coordinates can be “associated with” the output of a shading module if they are operating on the same pixel. (*Id.* at Q&A 157; Tr. 916:17-20.) The ‘372 patent’s specification also similarly states that “associated with” simply means data related to a pixel(s). (See JX-002 at 6:57-67 (describing colors “associated with” a quad, or group of pixels, in multiple passes).)

Nothing in the intrinsic evidence warrants a departure from the ordinary meaning of “coupled.” NVIDIA argues that because the texture coordinates must be “associated with” the output, the shader module must send its output to the texture look-up module, and that somehow requires directional coupling from the shader module to the texture look-up module. (See Tr. at 252:19-253:16 (interpreting “associated with” to mean “[i]t’s in the output, it’s with the output, it’s a function of the output”). But claim 23 imposes no such requirement. The claim only requires that the shading calculation output be “associated with” texture coordinates. (See Tr. at 916:17-20 (explaining texture coordinates used by the texture look-up module can be “associated with” the output of the shading module if the two modules are “both working on the same fragment”); RX-002C (Crawfis) at Q&A 157.

NVIDIA’s construction, in effect, tries to replace the phrase “associated with,” with the phrase “receives and uses.” But the evidence shows that the patentees knew how to limit a claim to “using the output”—as they did in claim 21 where the second calculation “us[es] the output” of the first calculation. (See JX-002, claim 21.) They did not, however, impose the same limitation on the texture look-up module in claim 23. Moreover, even if I were inclined to accept NVIDIA’s “receive and use” theory, which I am not, it still does not justify limiting the

PUBLIC VERSION

claims as NVIDIA proposes. Contrary to NVIDIA's construction, the '372 patent discloses that in multi-pass rendering, the texture look-up module can "receive and use" texture coordinates from a shader module in a second pass through the pipeline. (RX-3495C (Crawfis) at Q&A 34; RDX-398.)

NVIDIA relies extensively on the embodiment illustrated in Figure 4 of the '372 patent to support its limiting claim construction. (*See e.g.*, CX-2127C (Aliaga) at Q&A 123 (relying on Figure 4 to "require[]" coupling in "both directions"), 36, 55-56.) However, as a general rule the particular examples or embodiments discussed in the specification are not to be read into the claims as limitations. Here, Figure 4 is expressly described as "one embodiment" of the invention and there is no clear indication that the applicant intended to limit the invention of claim 23 to this embodiment. (*See* JX-002 at 5:65-67 ("FIG. 4 illustrates a hardware implementation for programmable shading in accordance with one embodiment of the present invention."), 6:48-49 (Figure 4 is "illustrative" implementation); *see also Epos Techs. Ltd. v. Pegasus Techs Ltd.*, 766 F.3d 1338, 1341 (Fed. Cir. 2014) ("[I]t is improper to read limitations from a preferred embodiment described in the specification...absent a clear indication" of an intention to do so.)) Notably, the specification discloses other embodiments where the output is recirculated through the graphics pipeline (so-called multipass) such that the texture look-up module receives texture coordinates via other components, not directly from the shading module. (Tr. at 254:25-255:4 (admitting "[t]here are embodiments of multipass where texture coordinates come from other places," including files on the computer); *see also* RX-3495C (Crawfis) at Q&A 34; RDX-398.)

NVIDIA argues that in a prior art figure in the '013 patent, which is incorporated by reference into the '372 patent, the applicants disclaimed "linear" pipelines, or what NVIDIA

PUBLIC VERSION

calls “texture-to-shader” topologies. (See CIB at 7, 28.) NVIDIA argues this disclaimer limits the construction of the term “coupled” in the ’372 patent claims. (See e.g., CX-2127C (Aliaga) at Q&A 140.) A disclaimer of claim scope must be “clear and unmistakable.” *Omega Eng’g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1325-26 (Fed. Cir. 2003). However, as described above, Figure 4 of the ’372 patent shows one-way coupling (in addition to two-way) and the patent discloses an embodiment involving multipass rendering. Moreover, the ’013 patent itself does not support a disclaimer of “linear” or “texture-to-shader” configurations. The ’013 patent merely states that prior art pipelines, including both “linear” and “non-linear” types, enabled one texture fetch and texture calculation per rendering pass and were “static in nature.” (CX-1952C at 2:49-57.) The ’013 patent proposes to solve that problem by either using a “feedback loop” or adding more shading and texture look-up modules to “constitute at least four logical modules.” (*Id.* at 3:28-42.) Contrary to NVIDIA’s argument, neither proposed solution requires a “shader-to-texture” configuration that “use[s] the output of a shading calculation to retrieve texture information.” Thus, for at least the reasons above, I find there has been no clear disavowal of ’372 patent claim scope.

In short, the actual language of claim 23 requires only that the texture look-up module be “coupled” to the shading module and use texture coordinates “associated with” the shading module output—for example, related to the same pixel or fragment. The evidence shows that the plain and customary meaning of the term “coupled” does not require a specific directional connection between the texture look-up module and the shading module. Rather, under the plain meaning of the term, the connection between the texture look-up module and the shading module can be in any direction or both directions. Respondents have not shown that the applicant acted as his/her own lexicographer to give a special meaning to the term “coupled” and I have found

PUBLIC VERSION

Respondents' disclaimer argument not persuasive. Thus, I find no basis to depart from the term's plain and ordinary meaning, which I note is consistent with the language of the claims and the specification, to require the limiting construction proposed by NVIDIA.

C. Infringement

NVIDIA argues that the Accused Products with Adreno, PowerVR, and Mali GPUs infringe claim 23 of the '372 patent. (CIB at 14.) NVIDIA argues the Accused Products with Mali GPUs also infringe claim 24.

1. Claim 23

NVIDIA argues that the Accused Products infringe claim 23.⁵ To this end, NVIDIA's expert, Dr. Aliaga, testified in detail how each limitation of claim 23, claim 22, and claim 21 is met by the accused Adreno, Mali, and PowerVR GPUs. (CX-006C (Aliaga) at Q&A 94-211, 241-257, 280-293, 311-326). The Staff argues that with respect to those products analyzed by Dr. Aliaga NVIDIA has proven infringement of claim 23. (SRB at 19-20.)

Respondents argue that: (1) certain products do not infringe because NVIDIA failed to present infringement evidence with respect to those products; (2) NVIDIA has failed to show that the Accused Products have a "shading module" as claimed; and (3) the Accused Products do not provide a "texture look-up module" coupled to a "shading module". (RIB at 54-63.)

⁵ NVIDIA cites to its pre-hearing brief in support of its argument that the Accused Products meet all the elements of claim 21. (See CIB at 14 (citing NVIDIA PrHB at 85.) Such an incorporation by reference is a violation of my ground rules and therefore I am striking the citation from NVIDIA's opening brief. (See Ground Rule 15.1.1 ("The Post-hearing briefs shall not incorporate anything by reference ...").) Respondents argue NVIDIA's conclusory statement that "the accused products meet each element of claim 21" and citation to over 100 Q&As from Dr. Aliaga's witness statement in support is also an improper incorporation by reference. Knowing NVIDIA's allegations of infringement, Respondents did not dispute in their pre-hearing brief, at the hearing, or in their initial post-hearing brief that the Accused Products satisfy the limitations of claim 21. Therefore I do not find in this instance NVIDIA's citation to Aliaga's witness statement improper.

PUBLIC VERSION

a. Sufficiency of NVIDIA's infringement evidence

The Parties' Positions

NVIDIA asserts that the Accused Samsung devices with Adreno, Mali and PowerVR GPUs infringe the '372 patent and that I should recommend a limited exclusion order prohibiting the importation of infringing products. (CIB at 22.) NVIDIA argues that Respondents seek to evade this exclusion order by arguing that only the fifty representative products identified by Dr. Aliaga can be found to infringe, because allegedly "NVIDIA has failed to accuse any software on over 40 Samsung products" listed in RDX-408C. (CIB at 22-23.) NVIDIA argues that critically, Respondents have not identified differences in the design or operation of these supposedly "non-accused" articles that would even possibly suggest a different infringement outcome. (*Id.* at 23.) NVIDIA argues that to the contrary, these devices appear to have the same Adreno, PowerVR, and Mali GPUs. (*Id.*) Thus, NVIDIA contends that upon a finding of violation, all infringing articles should be excluded. (*Id.*)

Respondents assert that there is no dispute that claims 23 and 24 require shader software to satisfy the performing and saving limitations: "performing a first shading calculation," "saving the output," and "performing a second shading calculation." (RIB at 55.) Respondents argue that NVIDIA failed to present evidence that over forty Samsung products practice these limitations. (*Id.*) Respondents argue NVIDIA's failure to prove the limitations for Claims 23 and 24 means, these products—which are specifically identified in RDX-408C and include Mali and Adreno GPUs, have not been proven to infringe. (*Id.*)

The Staff acknowledges Respondents argument, but does not address the argument in its initial brief. (SIB at 28.) In its reply brief the Staff agrees with Respondents that there is a failure of proof regarding some forty plus accused Samsung Products. (SRB at 19.)

PUBLIC VERSION

Discussion

There is no dispute that at least four limitations in claim 21 (on which asserted claims 23 and 24 depend) require specific software. (See JX-002, claim 21 (“logic for performing a first shading calculation...,” “logic for saving the output,” “logic for performing a second shading calculation...,” and “decoupled variables”); CX-006C (Aliaga) Q&A 63, 467 (logic requires software); RX-3495C (Crawfis) at Q&A 60.) NVIDIA does not argue otherwise in its briefs.

NVIDIA’s infringement expert for the ‘372 patent, Dr. Aliaga, failed to analyze software in over forty of the accused Samsung products. Thus, NVIDIA failed to present evidence regarding over forty accused Samsung products (identified in exhibit RDX-408C), and thus those products have not been shown to infringe at least the following limitations of claim 21: “logic for performing a first shading calculation...,” “logic for saving the output,” “logic for performing a second shading calculation...,” and “decoupled variables.” (RX-3495C (Crawfis) at Q&A 61-65; RDX-408C.)

NVIDIA now argues that the products that were analyzed by Dr. Aliaga are “representative products,” and thus any remedy in this investigation should still nevertheless cover the forty plus accused products that Dr. Aliaga did not analyze. (See CIB at 22-23). NVIDIA’s logic is defective. It is NVIDIA that has not presented any evidence in support to show that the analyzed products are in fact representative. Indeed, NVIDIA does not cite to any evidence that would support this belated contention. (*Id.*) Moreover, Dr. Aliaga admitted that “the software installed on the device varies somewhat from device to device.” (CX-006C (Aliaga) at Q&A 96.)

NVIDIA also incorrectly tries to shift the burden to Respondents and suggest that it is Respondents who need to show that the “non-accused articles” are different. (CIB at 22-23.)

PUBLIC VERSION

But “complainant cannot rely on an assumption that all of the accused products include similar structure to shift the burden to the accused infringer.” *Certain Integrated Circuits, Chipsets, & Prods. Containing Same Including Televisions, Media Players, & Cameras*, Inv. No 337-TA-709, Final Init. and Rec. Det., 2011 WL 1836230, at *24 (U.S.I.T.C. Apr. 4, 2011). NVIDIA has the burden of proof on infringement. Because NVIDIA did present evidence regarding the forty plus accused products identified in RDX-408C, NVIDIA has not met its burden with respect to those products.

Accordingly, I find NVIDIA has failed to prove those products identified in RDX-408C infringe claim 23 of the '372 patent.

b. Do the Accused Products have a “shading module” and a “texture look-up module” coupled to a “shading module”?

Claim 22 adds a shading module limitation. Claim 23, which depends from claim 22, adds a limitation requiring “a texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output.”

(1) Adreno GPUs

[

] (JX-113C at QNVIDIAITC932_00030854 (annotated); *see also* CX-006C (Aliaga) at Q&A 283-285, 287-289.)

PUBLIC VERSION

As Dr. Aliaga testified, the Adreno [] receives and uses texture coordinates [] to retrieve texture information []. (CX-006C (Aliaga) at Q&A 283-285, 287-289.) At the hearing, on cross-examination, Dr. Crawfis admitted that he was “comfortable” calling []. (Tr. (Crawfis) at 907:6-8.) Dr. Crawfis further admitted that [] (RX-3495C (Crawfis) at Q&A 160.) Thus, I find the evidence shows the Adreno GPU includes the claimed “shading module” and “texture look-up module,” as well as the “texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output.”

PUBLIC VERSION

Respondents argue the []
], *see* RIB at 61, but NVIDIA no longer alleges the []
]. (*See* CIB.) Thus, I find this argument irrelevant. In any event, as discussed above, the
evidence shows the [] (See e.g., CX-006C
(Aliaga) at Q&A 283-285; Tr. (Crawfis) at 907:6-8; JX-116C at QNVIDIAITC932_00033249
[]

Respondents also argue “NVIDIA has not identified the required ‘logic for saving the
output’” and thus, [] (RIB at 62.) Respondents,
however, never made this argument in their pre-hearing brief, *see* RPHB at 96-97 and 101-102,
and thus I deem it waived. (Order No. 2 at Ground Rule 11.2 (Contentions not set forth in the
pre-hearing brief “shall be deemed abandoned or withdrawn.”).) Moreover, claim 22 merely
requires “logic for saving the output” of “a first shading calculation.” (JX-002, Claim 22.) As
Dr. Aliaga testified, the Adreno GPUs []

[] (CX-006C (Aliaga) at Q&A 281 ([]
] Respondents argue the []
] *see* RIB at 62, but there is no requirement in the
claim that the logic be outside the shading module, it only has to save the output of shading
calculations, which is generated inside the shading module. (JX-002, claim 22.) As Dr. Aliaga
testified, the output of the shading calculation is saved [] before being sent to
other modules. (CX-006C (Aliaga) at Q&A 281.)

PUBLIC VERSION

(2) PowerVR GPU

As shown below, Imagination's technical documents prove that PowerVR GPUs include a USSE shading module (shown in the figure below in orange) and a Texture Co-Processor [] texture look-up module (shown in the figure below in lavender). (CX-047C at IMGPLC-932INV0001781, IMGPLC-932INV0001782 (annotated); CX-006C (Aliaga) at Q&A 311-323).)

In fact, at the hearing, on cross-examination, Respondents' expert, Dr. Crawfis, agreed the "USSE is a shading module, yes." (Tr. (Crawfis) at 913:13-15.) The evidence also shows the PowerVR GPU includes the claimed "texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output." (CX-047C at IMGPLC-932INV0001782 (annotated); CX-006C (Aliaga) at Q&A 311-323).) As seen in the picture above, the arrow (highlighted in blue) shows "texture lookup" information is output by

PUBLIC VERSION

the shading module to be received and used by the texture look-up module to retrieve texture information [

]. (CX-006C (Aliaga) at Q&A 319-323.)

Accordingly, I find the evidence shows the PowerVR GPUs include the claimed “shading module” and “texture look-up module,” as well as the “texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output.”

Respondents argue NVIDIA’s infringement contentions must fail because the arithmetic logic units (“ALUs”) alone within the USSEs are not shading modules, *see* RIB at 60, but NVIDIA no longer alleges the ALUs alone within the USSEs are also shading modules. Thus, I find this argument irrelevant.

Respondents also argue that the PowerVR GPUs do not infringe if I accept “Nvidia’s new construction of ‘shading module,’ which requires hardware for generating colors without receiving them from other elements.” (RIB at 60.) NVIDIA does not make this claim construction argument in its post-hearing brief. In fact, NVIDIA argues the asserted claims do not require “interpolating colors per fragment.” (CRB at 28.) Thus, I find Respondents’ argument irrelevant.

Respondents further argue that NVIDIA has failed to show that the texture look-up module is coupled to the shading module. (RIB at 61.) In particular, Respondents argue the accused ALUs are not coupled to the Texture Co-Processor either directly or indirectly and that to reach the [

] (*Id.*) As previously discussed, NVIDIA is no longer asserting that the ALUs alone are the shading module. Rather NVIDIA argues the USSE is the

PUBLIC VERSION

shader module. In this regard, the evidence is clear that the Texturing Co-processor is coupled to the USSE. (CX-047C at IMGPLC-932INV0001781, IMGPLC-932INV0001782 (annotated); CX-006C (Aliaga) at Q&A 319-323.)

(3) Mali GPU

As seen in the figure below, the Mali GPUs include an Arithmetic Pipeline (shown in the figure below in orange) and a Texture Pipeline (shown in the figure below in lavender).

(JX-043C at 13 (annotated).) For the Accused Devices with Mali GPUs, the evidence shows that the “shading module” limitation is met by the “Arithmetic Pipeline” and the “texture look-up module” is met by the “Texturing Pipeline.” (CX-006C (Aliaga) at Q&A 243-253.)

Specifically, the evidence shows that the arithmetic operations (*i.e.*, shading calculations) for shader programs written in OpenGL are executed, by design, in the Arithmetic Pipeline. (*Id.* at

PUBLIC VERSION

Q&A 244, 245.) The evidence also shows that texture look-ups using texture coordinates are performed by the Texturing Pipeline. (*Id.* at Q&A 250, 252.)

As clearly illustrated in the picture above, the evidence shows the Arithmetic Pipeline (shading module) and Texturing Pipeline (texture look-up module) are [

]. (*Id.* at Q&A 250, 251; JX-0043C at ARM_ITC_932_004969; CDX-006 at 77.) Referring again back to the figure, the evidence shows that [

]. (CX-006C (Aliaga) at Q&A 250-252.) Accordingly, I find the evidence shows the Mali GPUs include the claimed “shading module” and “texture look-up module,” as well as the “texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output.” (JX-043C at 13; CX-006C (Aliaga) at Q&A 243-253.)

Respondents make several related arguments based on the premise that the Arithmetic Pipeline cannot “perform shading calculations.” (RIB at 55-58.) Respondents argue that the Arithmetic Pipeline does not perform a shading calculation because it just performs arithmetic and is only one subcomponent that must work with dozens of other subcomponents in the Tripipe to perform the claimed shading calculation. (RIB at 56.) However, respondents argument relies on a construction of “shading calculation” that is inconsistent with the construction I have adopted for that term. Under my construction, a “shading calculation” is “a computation of a value concerning the appearance of a surface.” Thus, a shading calculation is a “computation” (*i.e.*, an arithmetic operation), and there is no dispute the Arithmetic Pipeline performs such operations. (Tr. (Larri) at 661:6-663:2 [

PUBLIC VERSION

] (emphasis added); RX-012C (Larri) at Q&A 119 [

] (emphasis added).)

Respondents also argue that a “shading module” cannot be a “subcomponent” that performs only a portion of the shading calculations. (RIB at 56-58.) According to Respondents, the Arithmetic Module must perform all of the operations that are invoked when executing a shading program (*e.g.*, fetching instructions, loading source data, *etc.*) in order to be a shading module. (*Id.*) As Dr. Aliaga explained, this is simply incorrect. The evidence shows that when executing a line of code from an OpenGL shader program, the GPU will perform a number of different operations, including, among other things, shading calculations. (CX-006C (Aliaga) at Q&A 122-223 (identifying various OpenGL code segments that will cause the accused devices to perform shading calculations and texture look-ups, among other things).) The shading module performs the shading calculations, *i.e.* the arithmetic operations, and other components may perform a number other operations (*e.g.*, creating threads, scheduling threads, texture look-ups, clipping, rasterization, *etc.*). The fact that other operations are invoked as part of the process does not mean that the shading module cannot perform shading calculations or that it is only performing a portion of the shading calculations.

c. Conclusion

As set forth in detail above, I find the Accused Products with Adreno, PowerVR, or Mali GPUs include the claimed “shading module” and “texture look-up module,” as well as the “texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output.” Accordingly, I find for at least the reasons above that NVIDIA has shown by a preponderance of the evidence that the Accused Products

PUBLIC VERSION

with Adreno, PowerVR or Mali GPUs, except those listed in RDX-408C, infringe claim 23 of the '372 patent.

2. Claim 24

NVIDIA contends that the Accused Products with Mali GPUs also infringe claim 24 of the '372 patent. (CIB at 20-22.) To that end, Dr. Aliaga opines that each limitation of claim 24 is met by those products. (CX-006C (Aliaga) at Q&A 256.) Respondents and the Staff argue that the Accused Products do not include a feedback loop as required by claim 24. (RIB at 59; SIB at 31-32.)

The Parties' Positions

NVIDIA argues that the Arithmetic Pipeline shading module is part of the "Tri-pipe" and as shown in the Mali Manual, there is a feedback loop between the output of the shading module and the input to the shading module. (CIB at 21.) NVIDIA argues that the [

] (*Id.*)

Respondents assert that NVIDIA contends the claimed feedback loop is a "register file" (*i.e.*, RAM) []. (RIB at 59.) But, Respondents argue, RAM is not a feedback loop. (*Id.*) Respondents argue that a feedback loop has a fixed structure and operation: it is a loop that feeds back an output to an input in a specified order. (*Id.*) Respondents argue that by contrast, the Mali GPU's register file is RAM that, just as with any other memory, can store data from multiple sources in a number of arbitrary registers and provide random access to multiple storage locations at any point in time. (*Id.*) Respondents argue that its function is to save and access any stored data—[]—so that programs can be executed in a fully-programmable processor. (*Id.*) Respondents contend one of ordinary

PUBLIC VERSION

skill in the art understands that RAM is fundamentally different than a feedback loop because it is not a “loop” that feeds output back to an input. (*Id.*)

The Staff asserts that NVIDIA argues that the “feedback loop” limitation of claim 24 is met in the Mali GPUs because the alleged shading module can [

]. (SIB at 31-32.) However, the Staff argues, a register file is not a feedback loop. (*Id.* at 32.) The Staff argues that although the Accused Products enable the shader modules to save and load data, they do not contain feedback loops as claim 24 requires. (*Id.*)

Discussion

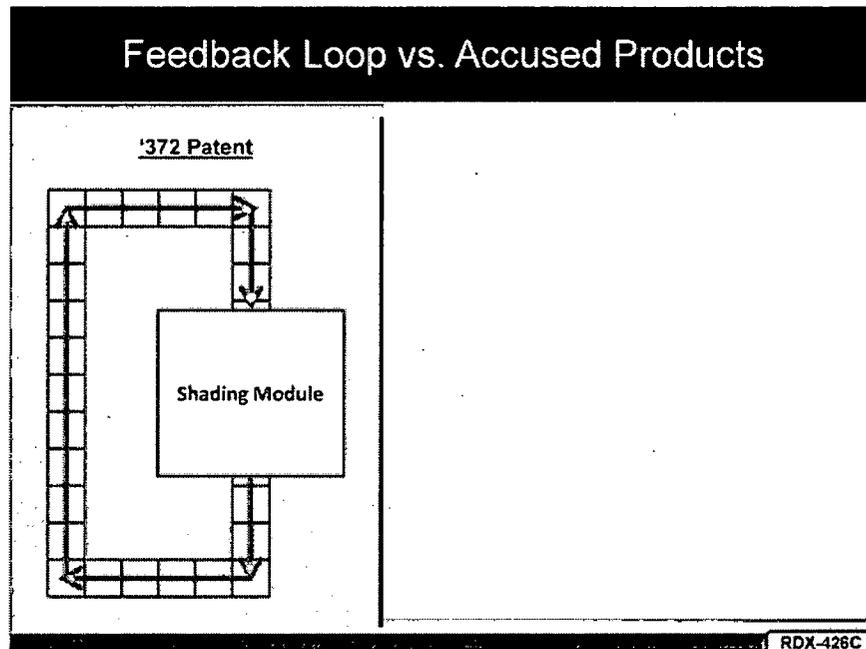
Claim 24 requires a feedback loop coupled between an input and an output of the shading module. Referring to the figure below from the Mali Manual, NVIDIA argues the Mali GPUs include the claimed feedback loop.

(JX-043C at 9 (annotated).) Specifically, NVIDIA states that “there is a feedback loop (colored yellow) between the output of the shading module (the arrow leading down and away from the [] and the input to the shading module (the arrow leading into the top of that box).” (CIB at 21.)

PUBLIC VERSION

The claimed “feedback loop” is a specific hardware structure known to one of ordinary skill in the art as “a structure (loop) that feeds back an output to an input.” (RX-3495C (Crawfis) at Q&A 98.) At the hearing, Dr. Aliaga admitted that “feedback loop” has to be a “loop” and has to “provide[] information from the output to the input.” (Tr. at 258:3-14.) This is consistent with the ’372 patent, which describes a FIFO buffer that moves data through a set path as an exemplary feedback loop. (JX-002 at 6:52-56; RX-3495C (Crawfis) at Q&A 99.) As Dr. Crawfis explained, a “standard hardware implementation” of a FIFO buffer “would be a set of registers daisy chained together that follow this very linear path.” (Tr. at 929:5-14.)

The evidence shows that RAM, on the other hand, randomly stores data in any sequence from any number of sources. (CX-006C (Aliaga) at Q&A 256.) The evidence shows that one of ordinary skill would not consider such a structure a “loop” that feeds an output back to an input. The differences are illustrated in the figure below.



(RDX-426C.) As shown on the right of the slide, in RAM, control logic can retrieve multiple pieces of data simultaneously from many different registers, just like any other memory. (RX-

PUBLIC VERSION

3495C (Crawfis) at Q&A 96-98; *see also* RDX-424C; RDX-425C.) The evidence shows that such use of a register file was known well before 1999 and was known to be a distinctly different structure than the claimed feedback loop. (RX-3495C (Crawfis) at Q&A 106.)

NVIDIA asserts that Respondents' arguments are based on an "incredibly detailed and unsupported" definition of feedback loop. (*See* CIB at 21.) But as discussed above, NVIDIA's own expert conceded the correctness of Respondents' construction. (*See* Tr. 258:3-14 ("feedback loop" has to be a "loop" and "provide[] information from the output to the input").) NVIDIA's construction, on the other hand, which can cover any random access memory, eviscerates the word "loop" from the claimed "feedback loop."

NVIDIA asserts that Respondents' expert, Dr. Crawfis, acknowledged that the feedback loop could include a FIFO (first in-first out) register and that "a standard hardware implementation" would be "a set of registers daisy chained together," arguing that because the Mali architecture also includes a "register file" in its feedback loop it infringes claim 24. (CIB at 21.) But the evidence shows NVIDIA is confusing two different uses of registers. While the evidence shows a register itself is not a feedback loop, multiple registers can be arranged in a way to form a feedback loop. As Dr. Crawfis described, a FIFO is a particular "set of registers daisy-chained together that follow this very linear path"—*i.e.*, a feedback loop. (Tr. 929:10-14.) Contrary to NVIDIA's argument, this is the only register-configuration the '372 patent identifies as a "feedback loop." (JX-2 at 6:52-56; RX-3495C (Crawfis) Q&A 99.) The evidence shows a random access register file, on the other hand, uses registers (as well as more complicated control logic) in a different way to allow access from various structures, at various times, in any order. Accordingly for the reasons above I find NVIDIA's argument not persuasive.

PUBLIC VERSION

Moreover, claim 24 requires not just a feedback loop, but rather a “feedback loop coupled between an input and an output of the shading module *for performing the second shading calculation using the texture information from the texture look-up module in order to generate further output.*” (emphasis added.) Neither NVIDIA, nor its expert, Dr. Aliaga, make any attempt to show how these many other claim 24 limitations are practiced by the Mali GPUs.

Accordingly, I find for at least the reasons above that NVIDIA has failed to show by a preponderance of the evidence that the Accused Products with Mali GPUs infringe claim 24 of the '372 patent.

D. Domestic Industry – Technical Prong

NVIDIA argues its products embodying its Fermi, Kepler, and Maxwell architectures combined with graphics software meet the technical prong of the domestic industry requirement with respect to the '372 patent. (CIB at 23-25.) To that end, NVIDIA’s expert, Dr. Aliaga, testified in detail that NVIDIA’s products when combined with graphics software such as NVIDIA’sameworks software practice at least claims 21-23 of the '372 patent. (CX-006C (Aliaga) at Q&A 463-487.) The Staff agrees. (SIB at 32.)

Respondents argue that NVIDIA has not shown that its products practice claims 10, 23, and 24.⁶ (RIB at 62.) However, Respondents do not contest that NVIDIA’s products embodying Fermi, Kepler, and Maxwell architectures practice at least claims 21 and 22 of the '372 patent. Under Commission precedent, NVIDIA need only show that it practices a single claim of the '372 patent to satisfy the technical prong of the domestic industry requirement. *Certain Ammonium Octamolybdate Isomers*, Inv. No. 337-TA-477, Comm’n Op. at 55 (August 28, 2003)

⁶ NVIDIA does not assert that its products embodying Fermi, Kepler, and Maxwell architectures practice claim 10 or claim 24. (CIB at 23-25.)

PUBLIC VERSION

“In order to satisfy the technical prong of the 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.”)

Accordingly, based on the unrebutted testimony of Dr. Aliaga, NVIDIA has shown by a preponderance of the evidence that the products embodying Fermi, Kepler, and Maxwell architectures practice at least claims 21 and 22 of the ‘372 patent. Therefore, I find NVIDIA has satisfied the technical prong with regard to the ‘372 patent.

E. Invalidity

Respondents argue that claims 23 and 24 are anticipated by: (1) “McCool” (RX-2792), which is an article titled “Texture Shaders” that was published on August 8, 1999; (2) U.S. Patent No. 6,236,413 (“the ‘413 patent”) (RX-179), which was filed on August 14, 1998 and issued on May 22, 2001; and (3) “Ackerman” (RX-180), which is a 1993 publication titled “An Architecture for High Performance Rendering Engine.” Respondents also argue that claims 23 and 24 are obvious in light of the ‘413 patent.

1. Anticipation

a. “McCool”

Respondents argue that claims 23 and 24 are anticipated by an article titled “Texture Shaders” that was published on August 8, 1999 (“McCool”). Because McCool was published before the May 31, 2000 priority date of the ‘372 patent, McCool is prior art under 35 U.S.C. § 102(a).

NVIDIA had claimed in its pre-hearing brief that McCool was not prior art because the ‘372 inventors conceived the subject matter of the claims on August 6, 1999, two days before McCool’s August 8, 1999 publication date, and diligently reduced it to practice. This issue was hotly contested in Respondents’ pre-hearing brief and at the hearing. NVIDIA relegated this

PUBLIC VERSION

argument to a footnote in its post-hearing brief, to wit: “McCool is also not prior art because the ’372 Patent was conceived earlier and diligently reduced to practice.” (CIB at 34 n.10.) NVIDIA then string cites to 25 Q&As from Montrym’s direct witness statement, 23 Q&As from Montrym’s rebuttal witness statement, the entire rebuttal witness statement of Kovanis, 45 Q&As from Aliaga’s rebuttal witness statement, and pages 91-93 of NVIDIA’s pre-hearing brief. (*Id.*) Because this issue was contested, NVIDIA was obliged to develop its argument in order to preserve it. (Ground Rule 15.1.1 “Any factual or legal issues not addressed in the post-hearing briefs shall be deemed waived.”) NVIDIA did not develop this argument in its brief or provide any explanation of the evidence to which it string cites. It is NVIDIA’s responsibility to develop their argument, not my responsibility to ferret it out. Accordingly, I find NVIDIA did not properly preserve its prior-conception argument. Thus, I find the argument waived.⁷

(1) Claim 23

The Parties’ Positions

NVIDIA argues that like Ackermann and the ‘413 Patent, McCool discloses the same, old fashioned “texture-to-shader” arrangement and does not disclose the claimed “shader-to-texture” configuration. (*Id.* at 31.) NVIDIA argues that this is shown by the one-way arrows between the alleged texture look-up module and the alleged shading module. (*Id.*) NVIDIA asserts that Dr. Crawfis conceded “the output of what [he] say[s] is a shading module, is going to the fragment operations circle” and not to the texture look-up module, and that data flows from the texture look-up module but not back. (*Id.*) Thus, NVIDIA argues McCool does not disclose “a texture look-up module coupled to the shading module for retrieving texture information

⁷ Even if it were not waived, for the reasons set forth in the Staff’s initial post-hearing brief, which I adopt herein, I would find that NVIDIA failed to carry its burden to show an earlier date of conception and diligent reduction to practice. (SIB at 33-35.)

PUBLIC VERSION

using texture coordinates associated with the output” of the shading module. (*Id.*) NVIDIA asserts that McCool’s one way arrows demonstrate the alleged texture look-up modules cannot receive or use “texture coordinates associated with the output” of the shading module to “retriev[e] texture information.” (*Id.*) In fact, NVIDIA argues, McCool goes so far as to explain that “[s]hader programs do not *request* texture samples....” (*Id.*)

NVIDIA argues McCool is not in the claimed “shader-to-texture” configuration where the texture look-up module receives texture coordinates from the shading module. (*Id.* at 32.) NVIDIA asserts that Respondents’ “multi-pass” argument fails for the reasons it set forth with respect to Ackermann and the ‘413 Patent – namely, the fact McCool would require multiple passes to even potentially perform such texture operations proves it is not in the claimed “shader-to-texture” configuration of claim 23. (*Id.*)

NVIDIA asserts that Respondents rely on a single sentence from § 6.6 of McCool to argue that one of ordinary skill could modify McCool to add the claimed “shader-to-texture” configuration. (*Id.*) NVIDIA argues that this sentence confirms that McCool would not meet the elements of claim 23 unless it was somehow modified. (*Id.*) NVIDIA argues McCool, however, does not teach how such modifications could be made. (*Id.*) NVIDIA argues that to the contrary, McCool criticizes attempts to modify its teaching (in the manner Respondents propose) by stating that “an additional feedback channel” would “require that a new shader be selected, which inhibits some of the coherence based optimizations mentioned above.” (*Id.*)

Moreover, NVIDIA argues that as a matter of law, McCool does not anticipate because it does not enable one of ordinary skill in the art to practice the claimed invention, *i.e.*, a graphics pipeline with a “shader-to-texture” configuration. (*Id.*) NVIDIA argues that Respondents cannot meet their high burden to show McCool is enabled because Dr. Aliaga opines that McCool is not

PUBLIC VERSION

enabled, Dr. Crawfis does not dispute that McCool is not enabled, and McCool specifically states that it does not enable the allegedly anticipating configuration. (*Id.* at 33.) NVIDIA asserts that Dr. Aliaga clearly testified that McCool would not enable the invention of claim 23. (*Id.*) NVIDIA argues that this evidence is undisputed since Dr. Crawfis did not testify that McCool would enable one of ordinary skill to practice claim 23 or 24. (*Id.*) NVIDIA argues that to the contrary, Dr. Crawfis agreed McCool is “sort of at a high level and doesn’t have a lot of detail.” (*Id.*) NVIDIA argues that the experts’ opinion that McCool is “high level” and incomplete is consistent with the McCool reference, itself, which states that it includes “a simple preliminary ‘plausibility’ design for single-pass shading and some global architectural consequences.” (*Id.* at 33-34.) NVIDIA argues the “simple preliminary” single-pass design does not anticipate and McCool’s musings on “global architectural consequences” were to reject the modifications Respondents’ propose. (*Id.*) NVIDIA asserts that McCool states that using such a feedback channel would eliminate some required functionality and therefore would not work, *i.e.*, is not enabled. (*Id.* at 34.) NVIDIA argues that critically McCool states that the proposed approach would require “that a new shader be selected” instead of reusing the same shader as shown in the ’372 patent, and that this would “inhibit[] some of the coherence-based optimizations” of McCool. (*Id.*) NVIDIA argues that Dr. Aliaga’s testimony explained why a person of ordinary skill in the art would not have found McCool enabling, particularly with respect to the changes Respondents propose. (*Id.*) NVIDIA asserts that Respondents chose not to cross-examine Dr. Aliaga on this testimony and their expert did not offer any contrary opinions. Thus, NVIDIA argues, all of the evidence shows that McCool is not enabling and Respondents cannot meet their heavy burden of proving it is prior art. (*Id.*)

PUBLIC VERSION

Respondents argue that McCool is prior art under § 102(a). (RIB at 69.) Respondents assert that the '372 patent claims nothing new over McCool, arguing that McCool anticipates claim 23 under all parties' constructions. (*Id.*) Respondents argue that McCool discloses flexible and programmable hardware for performing iterative shading calculations and applying multiple textures to achieve realistic graphics. (*Id.*) Respondents argue that using McCool's hardware complex shading effects that "currently take multiple passes can be rendered in one pass." (*Id.* at 69-70.)

Respondents assert that NVIDIA concedes McCool discloses the claimed first and second shading calculations, decoupled variables, and texture look-up module. (*Id.* at 70.) In addition, Respondents argue McCool's texture shader is a shading module that performs iterative shading calculations for multitexturing. (*Id.*) Respondents also argue that the texture shader is coupled to the texture look-up module. (*Id.*)

Respondents contend that NVIDIA's only arguments in response to McCool's anticipatory disclosure are based on its narrow claim constructions, which Respondents argue are improper because they seek to limit the claims to a preferred embodiment of the invention. (*Id.* at 71.) Respondents also argue that NVIDIA's arguments are factually incorrect. (*Id.*)

Respondents assert that NVIDIA does not dispute that McCool's "texture lookup and filter" modules are the claimed "texture look-up modules" and that they retrieve texture information using texture coordinates. (*Id.* at 73.) Respondents argue that the texture lookup modules in McCool are directly coupled to the texture shader as described in a preferred embodiment (Fig. 4) of the '372 patent. (*Id.*) Respondents also argue that the texture coordinates used by the texture look-up modules are also "associated with the output" of the texture shader at least because they are "both working on the same fragment or both working on

PUBLIC VERSION

the same pixel, so they are associated...” (*Id.*) Respondents assert that NVIDIA does not dispute that McCool anticipates under these plain meaning interpretations of the claim. (*Id.*)

Respondents stress that NVIDIA’s only argument is based on its construction of the term “coupled” to require coupling in “both directions” in a single pass through the pipeline—a construction based on an alleged “disclaimer” of other types of coupling. (*Id.* at 73-74.)

Respondents argue that there was no such disclaimer, and that NVIDIA’s construction of “coupled” to require “two directions” is improperly based on the two-directional arrows in Figure 4 of the ’372 patent. (*Id.* at 74.)

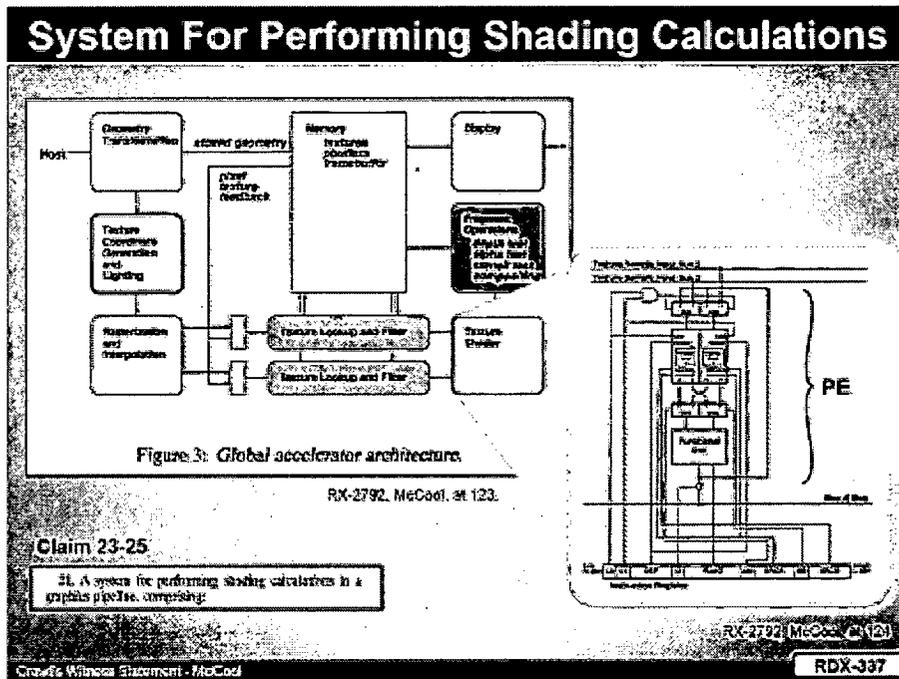
Respondents argue that even if NVIDIA’s construction applies, McCool still anticipates because it “clearly teaches bidirectional coupling” in section 6.6. (*Id.*) Respondents argue that section 6.6 expressly discloses that, in addition to receiving textures from the texture lookup module, the texture shader may also send “texture coordinates” that are “fed back to the texture lookup unit(s)” using an “additional feedback channel.” (*Id.*) Respondents assert that Dr. Aliaga tries to avoid this disclosure by arguing that McCool teaches away from using it. (*Id.* at 75.) Respondents argue that argument is both factually incorrect, and, even if true, not relevant to anticipation. (*Id.*) Respondents also contend there is no support for NVIDIA’s argument that McCool does not enable coupling in two directions. (*Id.*) Respondents argue that McCool expressly discloses how to implement bidirectional coupling with “an additional feedback channel” using a “packet approach” that “requires that a new shader be selected on a per-fragment basis.” (*Id.*)

The Staff argues that the evidence shows that McCool anticipates claim 23 of the ’372 patent. (SIB at 35.) The Staff contends that NVIDIA’s arguments are rooted in the preferred embodiment of the claimed invention, not in the claims. The Staff argues that “[a]s Respondents

show, each of these prior art references discloses a shader module and texture look-up module as claimed.” (*Id.*) The Staff asserts that NVIDIA claims that McCool simply disclose static, linear texturing, as was known in the prior art at the time of the '372 patent. (*Id.*) The Staff argues that the evidence, however, shows otherwise. (*Id.*) The Staff contends Dr. Crawfis demonstrated that McCool discloses flexible, iterative shading techniques, as opposed to the static, linear texturing that was identified in related the related '013 patent as prior art. (*Id.*)

Discussion

McCool discloses flexible and programmable hardware for performing iterative shading calculations and applying multiple textures to achieve realistic graphics. (RX-2792 at Abstract; see also RX-002C (Crawfis) at Q&A 98; CX-2127C (Aliaga) at Q&A 176.) Using McCool’s hardware—including a programmable “texture shader” and texture lookup modules—complex shading effects that “currently take multiple passes can be rendered in one pass.” (RX-2792 at Abstract.) The McCool architecture is shown in Figures 3 and 4, annotated below in RDX-337.



PUBLIC VERSION

(RDX-337.) According to Respondents, the above figure shows the claimed texture look-up module in purple and the shading module in yellow. Respondents' expert, Dr. Crawfis, testified in detail that McCool satisfies each and every limitation of claim 23 under the plain meaning of the terms. (RX-002C (Crawfis) Q&A 105-128; Crawfis Tr. at 919:22 ("McCool clearly anticipates claim 23."))

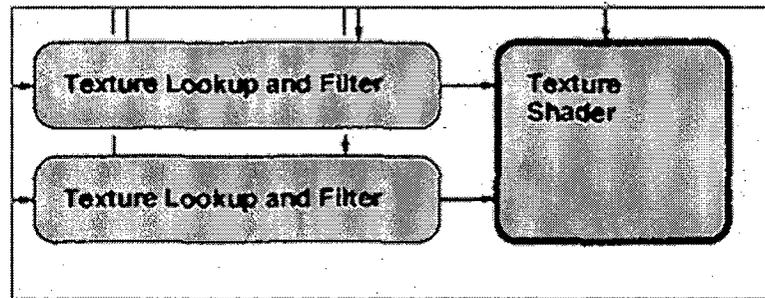
NVIDIA raises two arguments. NVIDIA argues that McCool does not disclose the "shader-to-texture"⁸ configuration required by claim 23 and that McCool is not enabled. As discussed in more detail below, I disagree with both arguments.

(a) Does McCool disclose a "texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output"?

NVIDIA does not dispute that McCool's "texture lookup and filter" modules are the claimed "texture look-up modules" and that they retrieve texture information using texture coordinates. (RX-2792 at Fig. 3, 123, 125; RX-002C (Crawfis) at Q&A 123-25, 127; RDX-347-48.) Moreover, as illustrated below, the evidence shows the texture lookup modules (shaded in purple) in McCool are directly coupled to the texture shader. (RX-002C (Crawfis) at Q&A 123; RX-2792 at 123; Tr. at 922:5-11 (texture look-up and shading module in Figure 3 are "clearly" coupled under plain meaning).)

⁸ As discussed, *supra*, the term "shader-to-texture" is a term coined by NVIDIA and is not found in the '372 patent or its claims. That term has connotations tied to NVIDIA's proposed claim construction and thus to avoid any confusion I will refer to the actual language of the claim.

McCool



(RDX-347.) The evidence also shows the texture coordinates used by the texture look-up modules are “associated with the output” of the texture shader at least because they are “both working on the same fragment or both working on the same pixel, so they are associated...” (Tr. at 916:17-20.) Accordingly, under the construction I have adopted herein for the phrase “texture look-up module coupled to the shading module,” the evidence clearly shows that McCool discloses the claimed “texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output.” NVIDIA does not appear to dispute such.

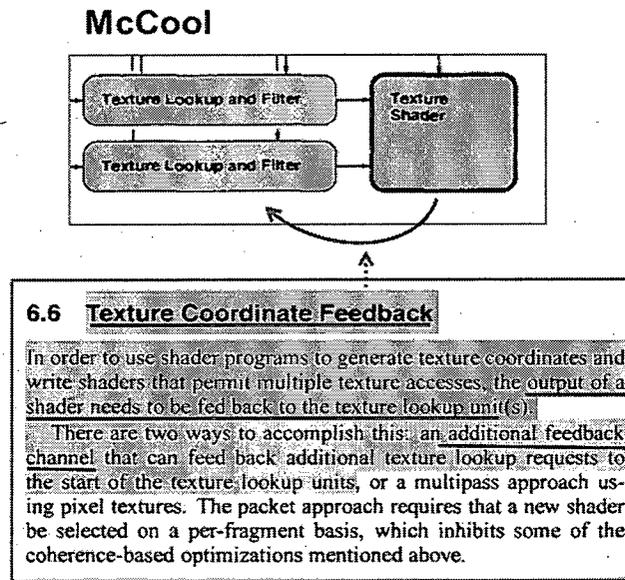
NVIDIA’s only argument is based on its construction of the term “coupled” to require directional coupling from the shader to the texture look-up module. I have not adopted this construction of the term “coupled.” Nevertheless, the evidence shows that even under NVIDIA’s construction McCool still anticipates because it “clearly teaches bidirectional coupling” in section 6.6. (Tr. at 922:12-17.) Section 6.6 of McCool states in relevant part:

6.6 Texture Coordinate Feedback

In order to use shader programs to generate texture coordinates and write shaders that permit multiple texture accesses, the output of a shader needs to be fed back to the texture lookup unit(s).

There are two ways to accomplish this: an additional feedback channel that can feed back additional texture lookup requests to the start of the texture lookup units, or a multipass approach using pixel textures. The packet approach requires that a new shader be selected on a per-fragment basis, which inhibits some of the coherence-based optimizations mentioned above.

This section expressly discloses that, in addition to receiving textures from the texture lookup module, the texture shader may also send “texture coordinates” that are “fed back to the texture lookup unit(s)” using an “additional feedback channel.” (RX-2792 at 123, 125; RX-002C (Crawfis) at Q&A 123, 125; RDX-347, 348.) As shown in red below, the “additional feedback channel” creates bidirectional coupling and thus even satisfies NVIDIA’s construction. (See e.g., CX-2127C (Aliaga) at Q&A 36.)



NVIDIA concedes that McCool “potentially discloses connecting the Texture Shader back to the texture lookup module [] in the phrase ‘an additional feedback channel.’” (CX-

PUBLIC VERSION

2127C (Aliaga) at Q&A 189.) In this configuration, the evidence also shows the texture coordinates used by the texture look-up modules are “associated with the output” of the texture shader. (RX-002C (Crawfis) at Q&A 125 (“Page 125 discloses that the texture shader can be used to generate texture coordinates, which are then sent to the texture lookup module to retrieve texture information for the next calculation. McCool discloses that this can be done using multiple passes through the pipeline or an additional feedback channel.”).) Accordingly, even under NVIDIA’s construction, the evidence clearly shows that McCool discloses the claimed “texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output.”

(b) Is McCool enabled?

NVIDIA’s argument that McCool is not enabled is tied to its proposed construction of “texture look-up module coupled to the shading module.” I have not adopted NVIDIA’s construction and therefore find NVIDIA’s argument not persuasive.

Although I have not adopted NVIDIA’s claim construction, as discussed, *supra*, even under NVIDIA’s construction I have found McCool discloses the limitations of claim 23. That finding is based at least in part on McCool’s disclosure of an additional feedback loop in section 6.6. NVIDIA argues that McCool criticizes and thus teaches away from the additional feedback channel because it requires alleged “modifications” and inhibits some “optimizations.” (CIB at 31-32.) However, even if true, this is not relevant to anticipation. *Rasmusson v. SmithKline Beecham Corp.*, 413 F.3d 1318, 1326 (Fed. Cir. 2005) (teaching away is irrelevant to anticipation). Regardless of disparagement, as discussed above, McCool’s disclosure of the elements of claim 23 anticipates the claim. *Verizon Servs. Corp. v. Cox Fibernet Va., Inc.*, 602 F.3d 1325, 1337 (Fed.Cir. 2010) (anticipation is found if reference teaches each and every limitation and “[i]t is well settled that utility or efficacy need not be demonstrated”). Moreover, I

PUBLIC VERSION

disagree with NVIDIA's argument. The evidence shows that McCool does not criticize the use of the "additional feedback channel" in favor of multi-pass rendering, but rather describes the "additional feedback channel" as one of two ways to accomplish texture coordinate feedback. (RX-2792 at 125 (describing bidirectional coupling with "an additional feedback channel" or a "multipass approach"); *see also* Tr. 922:15-17 ("So McCool clearly teaches the bidirectional coupling in his section 6.6 where he talks about two possible ways of doing dependent texturing."))

NVIDIA describes the "additional feedback channel" as a modification. But contrary to NVIDIA's argument, the evidence shows that the "additional feedback loop" is an expressly-identified "implementation," (*i.e.*, embodiment) of McCool's hardware where the shader requests textures. (RX-2792 at 123 (Section 6 is titled "Implementations").) Thus, NVIDIA's argument relies on an erroneous assumption that the disclosure of multiple examples renders one example less anticipatory, and thus should be rejected. This, of course, is contrary to the law. *Leggett & Platt, Inc. v. VUTEk, Inc.*, 537 F.3d 1349, 1356 (Fed. Cir. 2008); *see also Krippelz v. Ford Motor Co.*, 667 F.3d 1261, 1268 (Fed. Cir. 2012) (reversible error to hold a prior art figure did not anticipate where text of reference disclosed additional, anticipatory, embodiments).

NVIDIA also argues that McCool does not enable the disclosed "shader-to-texture configuration." (CIB at 32.) Enablement under section 102 requires that "suggestions be enabled to one of skill in the art." *Novo Nordisk Pharms., Inc. v. Bio-Tech. Gen. Corp.*, 424 F.3d 1347, 1355 (Fed. Cir. 2005). This is a different standard than enablement under section 112 because it does not require proof that a PHOSITA would be able to "use" the invention. *See id.* ("While section 112 provides that the specification must enable one skilled in the art to 'use' the invention, section 102 makes no such requirement as to an anticipatory disclosure.") (internal

PUBLIC VERSION

quotations and citations omitted). Under this standard, NVIDIA's contention amounts to arguing that a PHOSITA—defined in this case as including a person with a four-year degree in electrical engineering and two years of experience in graphics hardware—would not understand McCool's disclosure of a simple feedback channel connecting a shader module to a texture look-up module to be enabled. Such arguments, however, are belied by the express disclosures of McCool.

As described, *supra*, McCool explicitly discloses a “texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output” via an “additional feedback channel that can feed back additional texture lookup requests” from the shader to the texture look-up module using a “packet approach.” (*See supra*, at V.E.1.a.(1)(a); *see also* RX-2792 at 125.) Figure 3 of McCool also specifically illustrates other bidirectional connections to the Texture Shader, including from the Fragment Operations module. (RX-2792 at 123 (showing arrow in two directions from Texture Shader to Fragment Operations).) These are not “musings” as NVIDIA argues in its brief, but rather, the evidence shows, a specific hardware implementation using a flexible and programmable SIMD processor and known elements such as a feedback channel. (RX-002C (Crawfis) at Q&A 122 (“[A] person of ordinary skill would understand McCool's Texture Shader to be a flexible, programmable device that could perform numerous operations.”), 125; *see also* RX-2792 at 123, 125.)

NVIDIA does not offer any contrary evidence of undue experimentation to rebut these disclosures. NVIDIA only contends, without support, that the “additional feedback channel” in McCool “would require ‘that a new shader be selected’ instead of reusing the same shader as shown in the ‘372 patent.” (CIB at 34.) But the evidence shows McCool discloses the same type of shader described in Figure 4 of the ‘372 patent: a Texture Shader that generates texture coordinates to be “fed back” to the texture look-up module. (RX-2792 at 125.) (Dr. Aliaga

PUBLIC VERSION

essentially concedes as much, testifying that “McCool potentially discloses connecting the Texture Shader back to the texture lookup module” using an “additional feedback channel.” CX-2127C, (Aliaga) at Q&A 189.)

NVIDIA contends that Dr. Crawfis never opined that McCool was enabled. (CIB at 33-34.) However, as described above, Dr. Crawfis repeatedly testified how McCool implements bidirectional coupling in a single pass. (*See e.g.*, RX-002C (Crawfis) at Q&A 125 (Section 6.6 “discloses that the texture shader can be used to generate texture coordinates, which are then sent to the texture lookup module,” and “McCool discloses that this *can be done* using multiple passes through the pipeline or an additional feedback channel.”), Q.128 (“The feedback channel specifically “allows communication between the texture lookup and shading modules in a single-pass through the pipeline.”). Dr. Crawfis also explained that McCool specifically “teaches two ways to do it. One is to add an additional feedback channel that goes directly from the output of the shading module to the input of the texture lookup units.” (Tr. at 922:20-923:10; *see also id.* at 923:24-924:3 (“I think it’s very clear in section 6.6 that McCool now is going to support not only multitexturing but dependent texturing and the bidirectional coupling”).)

I find the above discussed disclosures from McCool more than sufficient to establish enablement. Indeed, NVIDIA cannot reasonably argue McCool is not enabling when the evidence showed that McCool provided more specifics than the ‘372 patent itself. (Tr. 897:11-15 (McCool “has more specifics in it than, say, the ‘372 patent.”).)

(c) Conclusion

For at least the reasons discussed above, I find Respondents have shown by clear and convincing evidence that McCool anticipates claim 23 of the ‘372 patent.

PUBLIC VERSION

(2) Claim 24

Respondents argue claim 24 is anticipated by McCool only under NVIDIA's construction of the term "feedback loop." (*See* RRB at 69.) I have not adopted NVIDIA's construction of the term "feedback loop." (*See supra*, at V.C.2.) Therefore, under a proper construction of the term "feedback loop" I find Respondents have failed to prove by clear and convincing evidence that claim 24 is anticipated by McCool.

b. U.S. Patent No. 6,236,413

Respondents argue that claims 23 and 24 are anticipated by U.S. Patent No. 6,236,413 ("the '413 patent"). The '413 patent was filed by Silicon Graphics Inc. on August 14, 1998 and issued on May 22, 2001. Because the '413 patent was filed prior to the May 31, 2000 priority date of the '372 patent, the '413 patent is prior art under 35 U.S.C. § 102(e).

(1) Claim 23

The Parties' Positions

NVIDIA argues that the '413 patent describes the old "texture-to-shader" configuration and does not teach the claimed "shader-to-texture" configuration. (CIB at 29.) NVIDIA argues that this can be seen in Figure 2 of the '413 patent, which shows the arrow going from the alleged texture look-up module to the supposed shading module. (*Id.*) NVIDIA argues that Cte, the output the alleged shading module, goes to "light environment unit" (206), "per-pixel lighting unit" (204) and recirculation pipe "TECTR" (211). (*Id.*) NVIDIA argues that Cte does not go to either the "Texture address unit" (201) or the "Texture filter unit" (202), which Dr. Crawfis argues are, together, the texture look-up module. (*Id.*) Thus, NVIDIA argues, as Dr. Crawfis conceded, "the output of the shader module, Cte,... doesn't go into the texture lookup module." (*Id.*)

PUBLIC VERSION

NVIDIA contends that Dr. Crawfis's instruction to apply an overly broad claim construction to manufacture an invalidity argument led him to opine that multiple and inconsistent portions of the '413 Patent were the shading and texture look-up modules of claim 23. (*Id.*) For instance, NVIDIA argues he opined that the shading module was either (i) "[P]er-pixel lighting unit" (204) or (ii) "[T]exture environment unit (203) or possibly (iii) "[T]exture filter unit" (202). (*Id.*) NVIDIA argues that Dr. Crawfis similarly opined that the claimed texture look-up module was (i) just "[T]exture address unit" (201) or (ii) that together with "[T]exture Filter unit" (202) or, possibly, (iii) those two in combination with "[T]exture environment unit" (203). (*Id.* at 29-30.) NVIDIA argues that Dr. Crawfis's substantial uncertainty over what constitutes the claimed shading and texture look-up modules casts significant doubt on the credibility of his invalidity opinions. (*Id.* at 30.)

NVIDIA contends that Respondents' repeat their single sentence "multi-pass" argument. (*Id.*) NVIDIA argues that the fact that the '413 patent's pipeline requires multiple passes to even potentially perform such texture operations proves it is not in the claimed "shader-to-texture" configuration. (*Id.*) NVIDIA argues that the alleged "shading module" and "texture look-up module" in the '413 Patent pipeline are coupled in the wrong direction for the wrong purpose. (*Id.*) NVIDIA also argues that Respondents' "multi-pass" argument is contradicted by the '413 patent specification, which teaches that the texture coordinates are generated within what Respondents contend is the claimed texture look-up module (the combination of texture address unit (201) and texture filter unit (202)) and are not "associated with the output" of the shading module. (*Id.*) Thus, NVIDIA argues, Respondents' "multi-pass" argument fails because regardless of the number of times you go through the pipeline, the alleged texture look-up

PUBLIC VERSION

module never receives texture coordinates associated with the output of the shading module as required by claim 23. (*Id.*)

Respondents argue the '413 patent is prior art under §§ 102(a) and (e). (RIB at 76.) Respondents argue the '413 patent anticipates claims 23 and 24 under the plain and ordinary meaning of the claims. (*Id.* at 77.) The '413 patent improves graphic processing with feedback loops (called “recirculation pipes”) that “implement complex graphics operations” and “iterative” shading, including multitexturing. (*Id.*) Respondents assert that NVIDIA does not dispute that the '413 patent discloses the claimed software, including first and second shading calculations and decoupled variables, as well as the claimed hardware texture lookup module. (*Id.*) Respondents argue that the “Texture Environment Unit” or “TEU” is a “shading module” that performs iterative shading to generate a color (Cte). (*Id.* at 78.) Respondents argue the TEU is directly connected to the texture look-up module. (*Id.*) Respondents assert that NVIDIA’s response to these disclosures is to rely on the same narrow claim constructions that it did for McCool. (*Id.*) Respondents argue that NVIDIA’s arguments are legally improper and factually incorrect.

Respondents argue that the Texture Address Unit and Texture Filter Unit in the '413 patent are together the claimed “texture look-up module” that retrieves texture information using texture coordinates. (*Id.* at 80.) Respondents argue the TEU, the claimed “shading module,” is coupled to this texture look-up module as shown in Figure 2 by the arrow labeled Ct directly connecting those two modules. (*Id.*) Respondents further argue that because the TEU output color corresponds to the same fragment as the texture coordinates used by the texture look-up module, the texture coordinates in the '413 patent are “associated with” the output. (*Id.*)

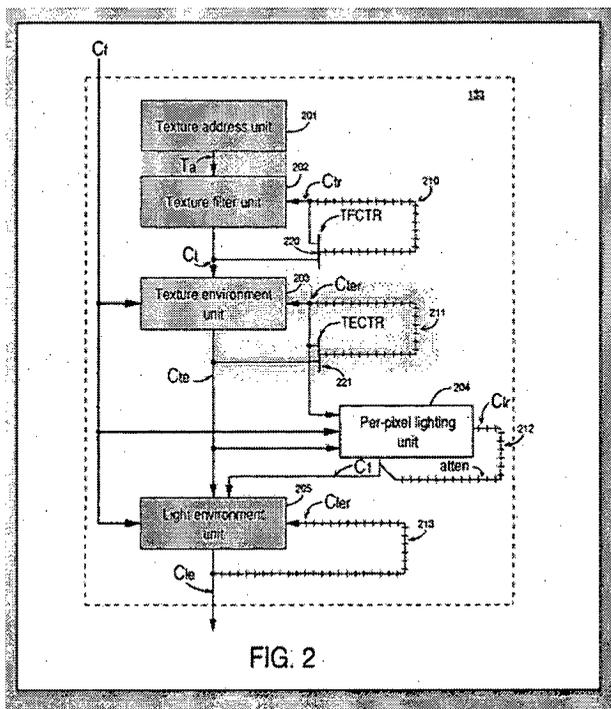
PUBLIC VERSION

Respondents maintain that NVIDIA is factually incorrect about the operation of the '413 patent hardware. (*Id.*) Respondents argue that the TEU is capable of outputting texture coordinates for use in the texture look-up module in multipass rendering. (*Id.*) Respondents claim that this is the type of rendering that Dr. Aliaga admitted was one embodiment of the '372 patent. (*Id.*)

The Staff argues that the evidence shows that the '413 patent anticipates claim 23 of the '372 patent. (SIB at 35.) The Staff contends that NVIDIA's arguments are rooted in the preferred embodiment of the claimed invention, not in the claims. The Staff says that "[a]s Respondents show, each of these prior art references discloses a shader module and texture look-up module as claimed." (*Id.*) The Staff asserts that NVIDIA contends that the '413 patent simply disclose static, linear texturing, as was known in the prior art at the time of the '372 patent. (*Id.*) The Staff argues that the evidence, however, shows otherwise. (*Id.*) The Staff argues that Dr. Crawfis demonstrated that the '413 patent discloses flexible, iterative shading techniques, as opposed to the static, linear texturing that was identified in related the related '013 patent as prior art. (*Id.*)

Discussion

The evidence shows that the '413 patent improves graphic processing with feedback loops (called "recirculation pipes") that "implement complex graphics operations" and "iterative" shading, including multitexturing. (RX-179 at 7:16-25; *see also* RX-002C (Crawfis) at Q&A 171-74; RDX-372.) The '413 patent thus addresses the same problem as the '372 patent: "a hardware graphics accelerator that allows flexible and iterative shading calculations and texture lookups." (*See* CX-006C (Aliaga) at Q&A 52.) Figure 2 of the '413 patent shows the hardware used for iterative shading:



(RDX-374.) According to Respondents, the above figure shows the claimed texture look-up module in purple, shading module in orange, and the feedback loop in green. Respondents' expert, Dr. Crawfis, testified in detail that the '413 patent satisfies each and every limitation of claim 23 under the plain meaning of the terms. (RX-002C (Crawfis) Q&A 105-128; Crawfis Tr. at 919:22 ("McCool clearly anticipates claim 23.").)

NVIDIA raises only one argument. NVIDIA argues that the '413 patent does not disclose the "shader-to-texture" configuration required by claim 23. I disagree.

The evidence shows the Texture Address Unit and Texture Filter Unit in the '413 patent (shown in purple above) are together the claimed "texture look-up module." (RX-179 at 8:15-27, Fig. 2; RX-002C (Crawfis) at Q&A 190-193; RDX-382-84.) The evidence shows that these units retrieve texture information using texture coordinates. (*Id.*) The evidence also shows that the TEU, the claimed "shading module," is coupled to this texture look-up module as shown in Figure 2 above by the arrow labeled Ct. (RX-179 at Fig. 2; RDX-382; RX-002C (Crawfis) at

PUBLIC VERSION

Q&A 190.) As is clear from the Figure, the claimed shading module and texture look-up module are directly connected together. (*Id.*) Further, because the TEU output color C_{te} corresponds to the same fragment as the texture coordinates used by the texture look-up module, the evidence shows the texture coordinates in the '413 patent are "associated with" the output. (Tr. at 916:4-20; RX-002C (Crawfis) at Q&A 190-191.) NVIDIA argues that the texture address unit in the '413 patent, not the shading module, is the only unit that can generate a texture coordinate. (*See* CIB at 30.) But Dr. Aliaga concedes that the '413 patent could also use the output of a shading calculation to retrieve a texture. (*See* RX-002C (Crawfis) at Q&A 193; RDX-384.) In fact, the evidence shows any module capable of outputting a color, such as the '413 patent's shading module, is necessarily capable of outputting a texture coordinate. (*See e.g.*, RX-2792 at 122 (color is also a texture coordinate as it can "be used either as a high-precision RGBA [red, green, blue, alpha] value, 4D texture coordinate, or a single depth value").) Thus, I find the evidence clearly shows that the '413 patent discloses a "texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output" under the construction that I have adopted herein for the phrase "texture look-up module coupled to the shading module." Accordingly, for at least the reasons above, I find Respondents have shown by clear and convincing evidence that the '413 patent anticipates claim 23 of the '372 patent.

(2) Claim 24

The Parties' Positions

NVIDIA do not contest in their opening brief that the '413 patent discloses the additional limitations of claim 24.

Respondents contend that NVIDIA does not dispute that the '413 patent discloses a feedback loop as required by claim 24. (RIB at 80.) Respondents argue that the recirculation

PUBLIC VERSION

pipe in the '413 patent is a feedback loop under the plain meaning of the term. Respondents argue that the recirculation pipe is a FIFO, which is the same structure the '372 patent identifies as a feedback loop. (*Id.*)

The Staff argues that the '413 patent discloses the use of recirculation pipes “to implement complex graphics operations by using a ‘multi-pass’ operation.” (SIB at 36.) The Staff argues that the multi-pass operation involves the recirculation of data through one or more of recirculation pipes 210-213 in order to perform iterative processing on the data, thereby implementing more complex operations than possible in a ‘single’ pass.” (*Id.*) Thus, the Staff argues claim 24 is anticipated. (*Id.*)

Discussion

As shown in Figure 2, *supra*, the '413 patent discloses a recirculation pipe 211 (shown in green) that is coupled between an output of the Texture Environment Unit 203 (*i.e.*, shader module) and an input of the Texture Environment Unit. (RX-002C (Crawfis) at Q&A 196-97; RDX-385-86; RX-179 at 3:5-16, 7:16-25, claims 3 and 23, Fig. 2; Tr. at 928:16-21.) The evidence shows the recirculation pipe is a FIFO, which is the same structure the '372 patent identifies as a feedback loop. (RX-002C (Crawfis) at Q&A 182; Tr. at 928:22-929:1, 929:10-14; RX-179 at 6:65-67.) Thus, I find the '413 patent clearly discloses the claimed feedback loop under the plain meaning of the term.

As stated above, NVIDIA does not contest that the '413 patent discloses the additional limitations of claim 24. Accordingly, I find for at least the reasons above that Respondents have shown by clear and convincing evidence that the '413 patent anticipates claim 24 of the '372 patent.

PUBLIC VERSION

c. “Ackerman”

Respondents argue that claims 23 and 24 are anticipated by an article written by Ackerman et al. titled “An Architecture for High Performance Rendering Engine, Rendering, Visualization and Rasterization Hardware” that was published in 1993 (“Ackerman”). (RX-180.) Ackerman incorporated by reference an article written by Cook titled “Shade Trees” (“Cook”) that was published in 1984. Because Ackerman and Cook were published more than one year before the May 31, 2000 priority date of the ‘372 patent, Ackerman and Cook are prior art under 35 U.S.C. § 102(b).

(1) Claim 23

The Parties’ Positions

NVIDIA argues that Ackermann does not anticipate claims 23 or 24. (CIB at 26.) NVIDIA contends that while Respondents argue Ackermann anticipates claims 23 and 24, their own expert disagreed and testified “I would not say Ackermann anticipates.” (*Id.*) NVIDIA argues that Dr. Crawfis’s admission at the hearing is even more significant as it is irreconcilably opposite his witness statement, when he testified “[i]t is my opinion that Ackerman anticipates claims... 23-2[4] of the ‘372 Patent.” (*Id.*)

NVIDIA argues that even under Dr. Crawfis’ “overly broad” construction, Ackermann does not disclose the claimed “shader-to-texture” configuration. (*Id.* at 27.) NVIDIA argues that Dr. Crawfis’s demonstrative and testimony makes clear, the output of Ackermann’s alleged shading module “doesn’t go to the texture lookup module.” (*Id.*) Thus, NVIDIA argues, Ackermann does not have “a ‘texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output’” of the shading module. (*Id.*)

PUBLIC VERSION

NVIDIA asserts that Respondents do not contend that Ackermann has the claimed “shader-to-texture” configuration. (*Id.*) NVIDIA argues that instead, Respondents speculate – in a single sentence – that “Ackermann could retrieve textures using coordinates associated with the output of a shading calculation through multi-pass rendering.” (*Id.*) Respondents contend this argument fails for three reasons. (*Id.*)

First, NVIDIA argues that Ackermann makes absolutely no reference to “multi-pass rendering.” (*Id.* at 28.) Second, NVIDIA argues that even if Ackermann did teach multi-pass (which, it does not), it still would not anticipate, because this modification of Ackermann would be no different than the prior art described and disclaimed in the ’372 Patent. (*Id.*) For instance, NVIDIA argues the ’013 patent, incorporated by reference into the ’372 Patent, criticized such prior art pipelines because they “only enable one texture fetch and texture calculation per rendering pass.” (*Id.*) NVIDIA argues that even in a hypothetical multi-pass topology, Ackermann does not disclose the claimed “shader-to-texture” configuration because the “texture look-up module” is not “coupled to the shading module for retrieving texture information using texture coordinates associated with the output” of the shading module. (*Id.*) According to NVIDIA, it is coupled in the opposite direction for a different purpose. (*Id.*) Third, NVIDIA argues Ackermann shows that the output of the alleged shading module (labeled “RGB α ”) is not “texture coordinates” and is not received or used by the alleged texture look-up module “for retrieving texture information” as required by claim 23. (*Id.*)

Respondents argue that Ackermann (and the Cook reference incorporated therein) is prior art under section 102(b) and that under NVIDIA’s infringement theory that an ALU can be a “shading module,” Ackermann anticipates claims 23 and 24 of the ’372 patent. (RIB at 82.) Respondents argue regarding the claimed hardware that NVIDIA concedes that Ackermann

PUBLIC VERSION

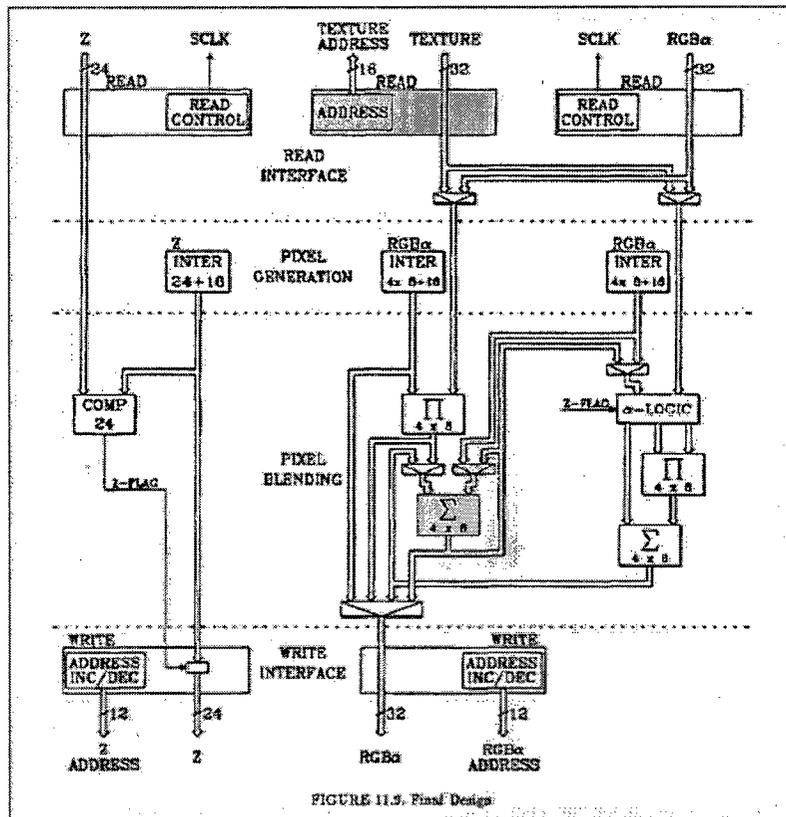
discloses a feedback loop and texture look-up module. (*Id.*) Respondents argue that Ackermann also discloses the same type of subcomponent—an adder—that NVIDIA identifies as a shading module for purposes of infringement. (*Id.*) Respondents contend that in the context of the prior art NVIDIA argues that the adder is not a shading module because it “only adds numbers.” (*Id.* at 82-83.) But, Respondents maintain, there is no relevant difference between an adder and an ALU. (*Id.* at 83.) Respondents reason that if NVIDIA’s proposal that such a subcomponent can be a shading module is adopted for infringement, Ackermann also anticipates the claims. (*Id.*)

Respondents contend that NVIDIA tries to distinguish Ackermann by arguing that the texture look-up module is not coupled to the adder because it is not coupled in two directions. (*Id.*) Respondents argue that such argument is based on NVIDIA’s unduly narrow claim constructions. (*Id.*) Respondents argue that like the ’413 patent, the Ackermann architecture can retrieve textures using texture coordinates in the output of a shading calculation through multipass rendering. (*Id.*)

The Staff argues that the evidence shows that Ackermann anticipates claim 23 of the ’372 patent. (SIB at 35.) The Staff contends that NVIDIA’s arguments are rooted in the preferred embodiment of the claimed invention, not in the claims. (*Id.*) The Staff argues that “[a]s Respondents show, each of these prior art references discloses a shader module and texture look-up module as claimed.” (*Id.*) The Staff asserts that NVIDIA contends that Ackermann simply disclose static, linear texturing, as was known in the prior art at the time of the ’372 patent. (*Id.*) The Staff argues that the evidence, however, shows otherwise. (*Id.*) The Staff argues that Dr. Crawfis demonstrated that Ackermann discloses flexible, iterative shading techniques, as opposed to the static, linear texturing that was identified in related the related ’013 patent as prior art. (*Id.*)

Discussion

The evidence shows that the '413 patent improves graphic processing with feedback loops (called "recirculation pipes") that "implement complex graphics operations" and "iterative" shading, including multitexturing. (RX-179 at 7:16-25; see also RX-002C (Crawfis) at Q&A 171-74; RDX-372.) The '413 patent thus addresses the same problem as the '372 patent: "a hardware graphics accelerator that allows flexible and iterative shading calculations and texture lookups." (See CX-006C (Aliaga) at Q&A 52.) Figure 11-9 of Ackermann shows the hardware used for iterative shading:



(RDX-367.) According to Respondents, the above figure shows the claimed texture look-up module in purple, shading module under NVIDIA's infringement theory in orange, and the feedback loop in green. (*Id.*) Respondents' expert, Dr. Crawfis, testified in detail that under

PUBLIC VERSION

NVIDIA's infringement theory, Ackermann satisfies each and every limitation of claim 23.

(RX-002C (Crawfis) at Q&A 134-166.)

NVIDIA raises only one argument. NVIDIA argues that Ackermann does not disclose the "shader-to-texture" configuration required by claim 23. I disagree.

The evidence shows the adder element (shown in orange above) is the claimed shader module. (RX-002 (Crawfis) at Q&A 152) The evidence shows the adder adds two numbers, and that in some cases one of the two numbers will be the output of the previous calculation. (*Id.*) The evidence shows the adder is used to perform iterative calculations, which include modifying a pixel's appearance based on a texture and other inputs. (*Id.* at Q&A 154.) The evidence shows that the output of the adder is the color of a pixel. (*Id.* at Q&A 152.)

The evidence shows that Ackerman also discloses the claimed texture look-up module (shown in purple above) and that this texture look-up module is coupled to the shading module under the plain meaning of "coupled."⁹ (*Id.* at Q&A 155, 156, 161.) Further, the evidence shows the texture look-up module receives texture information using texture coordinates associated with the output of a shading calculation. (*Id.* at 155, 157.) Specifically, Section 11.4.3.3 of Ackermann discloses texture-mapping where the vertex data are interpreted as addresses that point into a texture space, which are texture coordinates. (*Id.* at Q&A 157.) Because there is a correspondence between the fragment that is being operated on and the interpolated texture coordinates, the evidence shows the texture coordinates are associated with the output (*i.e.*, fragment or pixel) of the first shading calculation. (*Id.*; *see also id.* at Q&A 159-61 (the Ackermann architecture can retrieve textures using texture coordinates in the output of a

⁹ Properly construed in accordance with its plain and ordinary meaning, the term "coupled" include both directly coupled and indirectly coupled. Nothing in the '372 patent indicates otherwise.

PUBLIC VERSION

shading calculation through multipass rendering.); RX-180 at 165.) Thus, I find the evidence clearly shows that Ackermann discloses a “texture look-up module coupled to the shading module for retrieving texture information using texture coordinates associated with the output” under the construction that I have adopted herein for the phrase “texture look-up module coupled to the shading module.”

NVIDIA argues that Ackermann is “no different than the prior art described and disclaimed in the ‘372 Patent.” (*See* CIB at 28.) However, unlike that art, the evidence shows Ackermann discloses a programmable rendering engine, which includes a “pixel-processing block” to perform iterative shading, that is designed to execute “configurable and multiple stage shading algorithms using iterative and/or recursive techniques.” (RX-002C (Crawfis) at Q&A 139-140.)

Accordingly, for at least the reasons above, I find Respondents have shown by clear and convincing evidence that Ackermann anticipates claim 23 of the ‘372 patent.

(2) Claim 24

The evidence shows Ackermann discloses the claimed feedback loop. (RX-002C (Crawfis) at Q&A 162-163; RX-180.) This feedback loop is illustrated above in Figure 11-9. (*See* RDX-367.) The feedback loop is highlighted in green and shows the loop coupled between an input and an output of a shading module. (RX-002C (Crawfis) at Q&A 162-163; RX-180; RDX-367.) Section 11.4.5.2 of Ackermann describes how the results of a first calculation are stored in a first-in first-out, or “FIFO,” register, and later fed back for use in the second internal stage of the pixel modification stage. (RX-002C (Crawfis) at Q&A 162-163; RX-180) Notably, a FIFO is the same structure the ‘372 patent identifies as a feedback loop. Specifically, Section 11.4.5.2 states at page 172 that “We want to implement such a feedback loop for each of the two internal stages of the pixel modification stage,” thereby allowing iterative pixel accumulation

PUBLIC VERSION

and blending. (*Id.*) Moreover, the evidence shows that Ackermann discloses how with this feedback loop, the second shading calculation uses the texture information received from the texture lookup module to generate further output. (*Id.*) Thus, I find Ackermann clearly discloses the claimed feedback loop under the plain meaning of the term.

As stated above, NVIDIA does not contest that Ackermann discloses the additional limitations of claim 24. Accordingly, I find for at least the reasons above that Respondents have shown by clear and convincing evidence that Ackermann anticipates claim 24 of the '372 patent.

2. Obviousness

Respondents argue that “to the extent Nvidia’s construction of ‘coupled’ is adopted, which McCool expressly discloses, the '413 patent also renders the claims obvious.” (RIB at 83.) Under the claim constructions I have adopted herein, I have found that claims 23 and 24 are anticipated by the '413 patent. Accordingly, Respondents’ obviousness argument is moot.

VI. U.S. PATENT NO. 7,038,685

U.S. Patent No. 7,038,685 (“the '685 patent”) is titled “Programmable Graphics Processor for Multithreaded Execution of Programs.” (JX-005, '685 patent.) The '685 patent issued on May 2, 2006 on an application filed on June 30, 2003, and lists John Erik Lindholm as the inventor. (*Id.*) There are 45 claims. (*Id.*) NVIDIA alleges infringement of claims 1 and 15 of the '685 patent. Those claims read as follows:

Claim 1. A graphics processor for simultaneous multithreaded execution of program instructions associated with threads to process at least two sample types including dynamic load balancing of sample types among the threads comprising:

at least one multithreaded processing unit that includes a thread control unit, the thread control unit having:

a thread storage resource (TSR) configured to store thread state data for each of the threads to process the at least two sample types, wherein the at least two sample types are assigned to the threads for multithreaded

PUBLIC VERSION

execution based on an allocation priority among the at least two sample types, and

a programmable computation unit (PCU) for processing the sample types, the thread control unit assigning one of the at least two sample types to the PCU based on the stored thread state data for dynamically balancing the number of samples assigned to the threads on a cycle to cycle basis.

Claim 15. A graphics processor including a multithreaded processing unit adapted for dynamically controlling the number of threads allocated to processing each sample of a plurality of sample types comprising:

a thread control unit configured to store pointers to program instructions associated with threads for simultaneously processing samples of different sample types,

at least one programmable computation unit (PCU) configured to process the samples under control of the program instructions, each of the threads being allocated the sample types based on a priority assigned to the sample type.

(JX-005 at 17:17-35, 18:26-37.)

A. Level of Ordinary Skill in the Art

The Parties' Positions

The parties dispute the level of ordinary skill in the art with respect to the '685 patent.

(See CIB at 86; RIB at 87). Specifically, NVIDIA argues that a person of ordinary skill in the art of the '685 patent would have “at least a four-year degree in Electrical Engineering, Computer Engineering, Computer Science, or equivalent, as well as at least two years of experience in graphics processing including developing, designing or programming software or hardware for graphics processing units, hardware graphics accelerators or other graphics processing systems.”

(CIB at 86.) NVIDIA also argues that Respondents expert, Dr. Fussell, testified that NVIDIA’s proposed level of ordinary skill in the art is “a perfectly fine position” and that “someone reading the patent” would understand that it only describes “a hardware device with specialized graphics hardware.” (*Id.*)

PUBLIC VERSION

Respondents note that the parties' experts dispute the level of ordinary skill, but argue that no expert suggests that applying one level of ordinary skill or another affects any dispute. (RIB at 87.) Respondents' expert, Dr. Fussell, opines that a person of ordinary skill in the art of the '685 patent would have "at least a four year degree in computer science, electrical engineering, or computer engineering and five or more years in the field of computer hardware architecture research and/or development." (RX-009C (Fussell) at Q&A 30).

The Staff disagrees with NVIDIA's proposed person of ordinary skill in the art and supports the definition proposed by Respondents. (SIB at 39.) The Staff argues that NVIDIA's proposal would include someone with a computer science degree and two years of game programming, but without any substantial coursework or experience in computer hardware of any sort, much less computer graphics hardware. (*Id.*) The Staff states that it has a hard time understanding how someone without any substantial knowledge and experience in computer architecture could be considered one of ordinary skill in the art of a patent regarding the design of dynamic load balancing in a multithreaded processing unit. (*Id.*) The Staff argues that the '685 patent clearly requires expertise beyond computer graphics programming. (*Id.*) The Staff argues that in light of the technical field of the '685 patent and Dr. Fussell's testimony, a person of ordinary skill in the art would have "at least a four year degree in computer science, electrical engineering, or computer engineering and five or more years in the field of computer hardware architecture research and/or development."

Discussion

I agree with the Staff that in light of the technical field of the '685 patent, a person of ordinary skill in the art would have to have expertise beyond computer graphics programming. However, I disagree with Respondents and Staff that the patents are so complex as to require five

PUBLIC VERSION

years of experience. Thus, I find one of ordinary skill in the art of the '685 patent would have “at least a four year degree in computer science, electrical engineering, or computer engineering and two or more years in the field of computer hardware architecture research and/or development.”

B. Claim Construction

1. Agreed construction - “dynamically controlling...”

The preamble of claim 15 includes the phrase “dynamically controlling the number of threads allocated to processing each sample of a plurality of sample types.” This phrase in the preamble is limiting as it was added during prosecution to overcome certain prior art. (JX-11 at 113.) No party argues otherwise. The parties agree the limitation “dynamically controlling the number of threads allocated to processing each sample of a plurality of sample types” means “dynamically adjusting the number of threads allocated to process each type of sample, among two or more sample types.” (CIB at 88, RIB at 88.)

2. Disputed construction - “graphics processor”

The term “graphics processor” is found in the preambles of the asserted claims of the '685 patent. The parties dispute whether the term “graphics processor” is limiting. The parties all state that if the term is found to be limiting that it should be given its plain and ordinary meaning, however, the parties disagree on what the plain and ordinary meaning should be. (*Id.*). The parties’ proposed constructions are as follows:

Term	Proposed Constructions		
	Complainant	Respondents	Staff
“graphics processor”	Pre-Hearing Brief: GPU (graphics processing unit)	Not limiting as used in the preamble.	hardware capable of processing graphics data
	Post-Hearing Brief: “A processor that includes at least some specialized graphics hardware”	However, if construed the term means “hardware capable of processing graphics data”	(not limiting as used in preamble)

The Parties' Postions

NVIDIA contends that its experts Dr. Doggett and Dr. Dally, as well as Respondents' expert Dr. Fussell, are all in agreement that the specification and claims are limited to processors that have at least some specialized hardware for graphics processing. (CIB at 86.) Thus, NVIDIA argues, its construction, which gives effect to this understanding, should be adopted. (*Id.* at 86-87.)

NVIDIA asserts that prior to the hearing Respondents' contended that a graphics processor did not require specialized graphics hardware. (*Id.* at 87.) Specifically, NVIDIA argues that Respondents suggested that a graphics processor could be a general purpose processor that is merely "capable" of processing graphics data, *e.g.*, a CPU that can "sometimes" be a graphics processor and "sometimes" be just a general purpose processor. (*Id.*) NVIDIA argues that at the hearing, however, Respondents' expert, Dr. Fussell, testified that a "graphics processor" – as that term is used in the '685 patent – must include at least some "specialized graphics hardware," which a general purpose processor does not have. (*Id.*) NVIDIA argues that Dr. Fussell further supported NVIDIA's construction when he conceded that the '685 patent is directed only to a single device with hardware dedicated to graphics processing. (*Id.*)

NVIDIA argues that the specific field of the invention relates to "multithreaded processing, and more particularly to processing graphics data in a programmable graphics processor." (*Id.*) NVIDIA contends that in describing the need for the invention, the patent states "it would be desirable to provide improved approaches to processing different types of graphics data to better utilize one or more processing units within a graphics processor." (*Id.*) NVIDIA further asserts that the summary of the invention, without reference to an embodiment, describes the invention as "A graphics processor for multithreaded execution of program

PUBLIC VERSION

instructions associate with threads to process at least two sample types....” (*Id.*) NVIDIA argues that the patent could not be clearer that it is directed to graphics processors, and not general purpose processors. (*Id.* (citing *General Electric Co. v. Nintendo of America, Inc.*, 179 F.3d 1350 (Fed. Cir. 1999).)

Respondents contend that the term “graphics processor” appears only in the preamble of the claims 1 and 15 of the ’685 Patent. (RIB at 88.) According to Respondents, “Generally, the preamble does not limit the claims.” (*Id.*) Respondents argue that NVIDIA failed to explain in its pre-hearing brief or any witness statement why this phrase would be limiting at all, but regardless, the term is not limiting. (*Id.*) Respondents argue the body of claims 1 and 15 set forth a complete invention, and “graphics processor” does not provide antecedent basis for any term in the claim body. (*Id.*) Thus, accordingly to Respondents, this term is not limiting. (*Id.*)

Respondents argue that if construed, Respondents’ construction of “hardware capable of processing graphics data” represents the plain and ordinary meaning of the term. (*Id.*) Respondents contend that by contrast, NVIDIA’s proposed construction of “GPU (graphics processing unit)” is incorrect and injects uncertainty. (*Id.*) Respondents assert that NVIDIA claims it “coined the term ‘GPU’” well before filing of the ’685 patent in conjunction with the launch of products it claims embody the now-dropped ’488 Patent and “everyone in the computer industry accepted and acknowledged the term GPU as describing this new computer hardware.” (*Id.*) Respondents argue that the term GPU was “accepted and acknowledged” before filing of the ’685 Patent, and NVIDIA chose not to use it in the specification or the claims demonstrates the impropriety of NVIDIA’s construction. (*Id.* at 88-89.) Furthermore, Respondents argue “GPU” injects uncertainty into the claims. Respondents contend that NVIDIA’s pre-hearing brief lists at least four characteristics of a “GPU” without explaining

PUBLIC VERSION

whether those characteristics are intended as limitations. (*Id.*) Respondents argue that given this uncertainty, construing a graphics processor as “GPU” would be inappropriate and would only complicate the meaning of the simple term “graphics processor.” (*Id.*)

The Staff asserts that as a threshold matter, the parties dispute whether the term “graphics processor” in the preambles of the asserted claims of the '685 patent is limiting. (SIB at 40.)

The Staff contends that in their pre-hearing brief, NVIDIA failed to state with particularity any basis as to why the term “graphics processor” in the preambles of the asserted claims should be limiting. (*Id.*) The Staff argues that the bodies of the claims recite complete structure, and do not refer to the graphics processor in the preamble. (*Id.*) Furthermore, the Staff argues nothing in the bodies of the asserted claims breathe life and meaning to their preambles. (*Id.*) Thus, according to the Staff, the preambles are not limiting. (*Id.*)

The Staff argues that should the term “graphics processor” nonetheless be construed, Dr. Fussell opines that the plain and ordinary meaning of the term “graphics processor” is simply “hardware for processing graphics” and that this construction is consistent with the claim language and with the specification. (*Id.* at 41.)

Discussion

As a threshold matter, I must determine whether the term “graphics processor” in the preambles of the asserted claims of the '685 patent is limiting. NVIDIA argues that the term “graphics processor” is limiting because (1) it “is the subject of the invention of the patent and is necessary to give meaning to the claims”; and (2) it “recites structure emphasized as important by the specification.” (*See* CIB at 76.) I note that NVIDIA failed to state in its pre-hearing brief with particularity any basis as to why the term “graphics processor” in the preambles of the asserted claims should be limiting.

PUBLIC VERSION

“Generally, the preamble does not limit the claims.” *Allen Eng’g Corp. v. Bartell Indus., Inc.*, 299 F.3d 1336, 1346 (Fed. Cir. 2002). However, if a preamble “recites essential structure or steps, or if it is ‘necessary to give life, meaning, and vitality’ to the claim,” then the preamble can limit the scope of a claim. *Catalina Mktg. Int’l, Inc. v. Coolsavings.com, Inc.*, 289 F.3d 801, 808 (Fed. Cir. 2002) (citation omitted). Likewise, “when reciting additional structure or steps underscored as important by the specification, the preamble may operate as a claim limitation.” *Catalina Mktg.*, 289 F.3d at 808. “Conversely, a preamble is not limiting ‘where a patentee defines a structurally complete invention in the claim body and uses the preamble only to state a purpose or intended use for the invention.’” *Id.* (quoting *Rowe v. Dror*, 112 F.3d 473, 478 (Fed. Cir. 1997)). “[W]hether to treat a preamble as a claim limitation is determined on the facts of each case in light of the claim as a whole and the invention described in the patent.” *Bicon, Inc. v. Straumann Co.*, 441 F.3d 945, 952 (Fed. Cir. 2006) (quoting *Storage Tech. Corp. v. Cisco Sys., Inc.*, 329 F.3d 823, 831 (Fed. Cir. 2003)).

Here, I find the bodies of asserted claims 1 and 15 define structurally complete inventions directed to the multithreaded processing of different types of graphics data. Contrary to NVIDIA’s argument, I do not read the specification as limiting the claims to processing graphics data in a graphics processor. Nor do I read the specification as *underscoring the importance* of processing graphics in a graphics processor. To the contrary, the specification states that “[o]ne or more aspects of the invention generally relate to multithreaded processing, and more particularly to processing graphics data in a programmable graphics processor.” (‘685 patent at 1:7-9.) Thus, the applicant makes clear in the specification that the invention is not limited to processing graphics data in a programmable graphics processor. Rather, the invention is drawn broader to multithreaded processing. (See ‘685 Patent at 1:36-38 (“A method and apparatus for

PUBLIC VERSION

processing and allocating threads for multithreaded execution of graphics programs is described.”.) Accordingly, I find the term “graphics processor” in the preamble is not limiting.

Nevertheless, assuming *arguendo* the term “graphics processor” is found limiting, I would construe the term to mean “hardware capable of processing graphics data.” NVIDIA’s newly proposed construction of graphics processor as “a processor that includes at least some specialized graphics hardware” is without any support in the intrinsic record. NVIDIA relies solely on testimony elicited from Dr. Fussell at the hearing to support its construction. The testimony that NVIDIA cites does not discuss the claims but, instead, concerns the disclosures in “the patent”—*i.e.*, the preferred embodiments contained in the ‘685 patent specification. (Tr. at 797:25-798:24.) In this context, Dr. Fussell understandably agreed that the ‘685 patent specification disclosed “a hardware device with specialized graphics hardware,” not that the term “graphics processor” is limited to a preferred embodiment from the specification. (*Id.* at 795:15-796:7.) Dr. Fussell explained that NVIDIA’s distinction ignores the reality that “some generations of [graphics] hardware, for example, used general purpose processors that were specifically programmed to do functions in the graphics pipeline that was -- that are now done in GPU” as “part of a dedicated high-performance graphics system.” (*Id.* at 813:16-815:5.)

NVIDIA relies on *Gen. Elec. Co. v. Nintendo Co.* in support of its argument that the term “graphics processor” is limiting. *Gen. Elec. Co. v. Nintendo Co.*, 179 F.3d 1350 (Fed. Cir. 1999). In that case, the relevant claim language found only in the preamble recited: “by mapping bits from a display location in a memory associated with a computer onto the raster.” *Nintendo*, 179 F.3d at 1361-62. In that context, the court concluded that the preamble was limiting because “the inventors were working on the particular problem of displaying binary data on a raster scan display device.” *Id.* By contrast, the stated “problem” addressed by the ‘685 patent was having

PUBLIC VERSION

separate pixel and vertex processors. (JX-005 at 1:15-28; *see also* CIB at 13 (“The ’685 Patent claims a sophisticated unified shader architecture based on multithreaded processing units that can process different types of samples simultaneously, such as pixel and vertex data. This includes assigning samples to threads based on an allocation priority among sample types. This allocation priority is part of thread creation and can be fixed, dynamic or programmable. The invention allows the mix of vertices and pixels being processed by the multithreaded processing units to change during operation, *i.e.* dynamic load balancing in the execution of the threads.”).) Because the solution of a single unit performing pixel and vertex processing is reflected in the elements recited in the body of claims 1 and 15, I find *Gen. Elec. Co. v. Nintendo Co.* does not support NVIDIA’s argument.

C. Infringement

NVIDIA alleges that the Accused Products containing the Mali T6xx series, Mali T7xx series, Adreno 3xx series, and Adreno 4xx series GPUs directly infringe claims 1 and 15 of the ’685 patent. (CIB at 85.)

1. Claim 1

a. Mali GPUs

NVIDIA’s experts, Dr. Doggett and Dr. de la Inglesia, testified in detail that the Accused Products with Mali GPUs (Mali T-6xx and T-7xx) infringe claim 1 of the ’685 patent. (CX-008C (Doggett) at Q. 86-96, 98, 445, 466, 498, 519; CX-007C (de la Iglesia) at Q. 77-83, 90, 93, 101-107, 114, 117.) Respondents and the Staff argue that the Accused Products do not infringe, because they do not satisfy the following limitations of claim 1:

- (1) “wherein the at least two samples are assigned to the threads for multithreaded execution based on an allocation priority among the at least two sample types” (“limitation 1[d]”); and

PUBLIC VERSION

- (2) “the thread control unit assigning one of the at least two sample types to the PCU based on the stored thread state data for dynamically balancing the number of samples assigned to the threads on a cycle to cycle basis” (“limitation 1[e]”).

(RIB at 89-90; SIB at 42.)

To better understand the parties’ infringement and non-infringement arguments it is important to first explain the basic operation of the relevant aspects of the ARM Mali GPU.

The Mali GPUs (Mali T-6xx and T-7xx) are all based on the Midgard architecture. (RX-12C (Larri) at Q&A 18.) The figure below shows how information flows from an application to the Mali GPU.

(See JX-042C at 21.) The process begins when an application running on the device needs to perform graphics processing. (RX-12C (Larri) at Q&A 22.) To render an image, the application will issue a series of OpenGL ES calls that are received by the Mali GPU device driver. (*Id.*) As shown in the figure, both the application and the driver run on the application processor. (*Id.*) After receiving an OpenGL ES call, the driver generates graphics “jobs.” (*Id.* at Q&A 27.) In the Mali GPUs, the device driver creates five different job types: fragment, vertex, geometry,

PUBLIC VERSION

compute and tiler.¹⁰ (JX-042C at 25.) The driver writes a data structure called a “job descriptor” for each job into memory that is shared between the driver (running on the CPU) and the GPU hardware. (RX-012C (Larri) at Q&A 27, 34.) The driver next writes to memory-mapped control registers on the GPU that instruct the Job Manager in the GPU to retrieve and process the jobs. (*Id.* at Q&A 28.) The Job Manager reads a job from memory and breaks it down into smaller “tasks.” (*Id.*) The Job Manager then distributes these tasks to the “shader cores.” (*Id.* at Q&A 29.)

As shown in the figure below, the Job Manager within the GPU contains three job slots numbers 0, 1, and 2.

¹⁰ The ARM documentation uses the word “fragments” instead of “pixels.” These terms are used interchangeably and both are considered “sample types.” (See the ‘685 patent at 3:62-65.)

PUBLIC VERSION

(JX-042C at 24.) Each job slot can read a job descriptor from memory and process the job descriptor. (RX-012C (Larri) at Q&A 34.) The driver software determines [] by following a set of rules encoded into the driver software based on []. (Id. at Q&A 34, 35.) The driver software then writes to the GPU control registers [] and to direct the job slot []. (Id. at Q&A 34.)

The “cores” (a.k.a. processors) are what actually perform the required computing tasks. (Id.) As shown in the block diagram below, each shader core contains a “fragment thread creator” and a “generic thread creator.”

(JX-042C at 28.) Fragment thread creators are designed to receive fragment tasks from the Job Manager and to create fragment threads. (RX-012C (Larri) at Q&A 30.) Generic thread creators are designed to receive other tasks from the Job Manager, such as vertex tasks and to create the

PUBLIC VERSION

appropriate threads for those tasks. (*Id.*) The thread creators create the threads (*i.e.*, a sequence of program instructions along with various state information) necessary to process the tasks received from the Job Manager. (*Id.* at Q&A 31.) Each shader core also contains a “Tri-Pipe” execution unit that is responsible for executing instructions. (*Id.* at Q&A 30.) Once a thread has been created by either the fragment thread creator or generic thread creator, it will be sent to the Tri-Pipe for execution. (*Id.* at Q&A 32.) After execution the shader core writes the resulting image data to the frame buffer, which is usually used to provide the data to be displayed on the devices screen. (*Id.* at Q&A 33.)

(1) Do the Accused Products meet limitation 1[d]?

NVIDIA’s argument can be best summarized as follows:

The ’685 Patent teaches that any mechanism can be used to achieve the claimed “thread allocation priority,” which “may be fixed, programmable, or dynamic.” In Mali GPUs, thread allocation priority based upon sample type occurs by creating a greater throughput for vertices than for pixels. This is done by having two “job slots” that can send vertex tasks to the thread creators and only one “job slot” that can send fragment tasks to the thread creators. An extra vertex job slot results in vertex tasks getting priority over fragments because tasks are distributed from the three job slots to the shader cores on a [], so vertex tasks can be chosen for processing twice as often. As discussed above, the Job Manager reads a job from memory and breaks it down into smaller “tasks.”

(CIB at 102 (internal citation omitted).) In support of its argument, NVIDIA relies on the testimony of its experts and certain ARM technical documentation. (*See id.* at 102-104) For example, NVIDIA cites to Exhibit JX-042C, which includes the following passage:

(JX-042C at 26.) [

PUBLIC VERSION

]

(Tr. at 629:19-630:5; RX-12C (Larri) at Q&A 64; *see also* Tr. at 629:3-11; 636:9-637:7, 639:14-25; 681:10-23.) Although the Mali GPU has the inherent flexibility to theoretically use [

] for vertex jobs, the evidence decidedly shows that the Accused Products with Mali GPUs are not capable of using [] for vertex jobs because the Mali GPU device driver prevents such use. (RX-012C (Larri) at Q&A 40-60; RX-001C (Conte) at Q&A 174-83.)

Specifically, the evidence shows ARM provides Respondent Samsung with the Verilog RTL code for the Mali GPU and software device driver. (RX-012C (Larri) at Q&A 14.) The Verilog RTL code is used to synthesize the Mali GPU. (*Id.*) The device driver, which is shared across GPU models, “configures and controls” the Mali GPU. (*Id.* at Q&A 14, 20; Tr. at 667:17-668:18, 673:7-10; JX-042C at 21 (Figure 2-2).) The driver software determines which [

] by following a set of rules encoded into the driver software based on the type of []. (*Id.* at Q&A 34, 35.) Jobs originate from either OpenGL or OpenCL calls from the application processor. The types of jobs being sent to the job slots depend on whether OpenGL or OpenCL jobs are being processed. (*Id.* at Q&A 36.)

Applications that use OpenGL can produce various graphics related jobs such as vertex and

PUBLIC VERSION

fragment jobs and applications that use OpenCL can produce general-purpose compute jobs. (*Id.* at Q&A 37.)

Mr. Larri testified in detail about the driver code. (*Id.* at Q&A 43-58.) For example,

[

] (RX-012C (Larri) at Q&A 49-53.) [

] (*Id.* at

Q&A 55-57.) [

] (*Id.* at Q&A 58-59.)

I found Mr. Larri's detailed walk through of the driver code to be very credible. I note NVIDIA made no attempt at the hearing to question Mr. Larri on his detailed analysis of this source code and likewise NVIDIA's initial post-hearing brief does not address the substance of Mr. Larri's analysis. Mr. Larri's testimony shows that regardless of whether the job being sent originated from an OpenGL or OpenCL application call, in no instance is [] used to process vertex jobs. (*Id.* at Q&A 40, 42, 59, 60.) The evidence shows only [] is used to process vertex jobs. (*Id.* at Q&A 42.)

Because only one job slot is used to process vertices and one job slot is used to process fragments (*i.e.*, pixels) there can be no priority of vertex jobs over pixels as NVIDIA contends. In fact, because the processing of vertices and fragments would be on a one-to-one basis there can be no priority of any kind. As admitted by NVIDIA's expert, Dr. Doggett, this fact defeats NVIDIA's infringement assertions:

PUBLIC VERSION

Q. Okay. And if the Court - I want to be very clear about this. If the Court were to determine the factual record does not establish that that second job slot is used for vertex, then there would be a one-to-one ratio and no priority for vertex over fragment; right?

A. If the second job slot was somehow disabled, then yeah, there would be a one-to-one ratio and no priority.

(Tr. at 393:22-394:4.) Accordingly, for at the reasons above, I find NVIDIA has failed to show that the Accused Products with Mali GPUs meet limitation 1[d].

NVIDIA spends much of its Post-Hearing Brief discussing an alleged “hardware configuration” of Mali GPUs, *see* CIB at 98-100, 102-104 (making assertions regarding “hardware configuration,” how “the hardware is configured,” and what is “used in the RTL”), but there is no such thing as an allegedly infringing “hardware configuration” of the Mali GPU itself. As explained above, it is the driver software that is responsible for assigning jobs to job slots, not the hardware. The evidence shows that the accused Samsung mobile devices containing Mali GPUs are configured and controlled by ARM driver software. (Tr. at 667:17-668:18, 673:7-10; RX-012C (Larri) at Q&A 34; RX-01C (Conte) at Q&A 175.) Because the driver does not allow use of [] for vertex data, the evidence shows Samsung devices with Mali GPUs are neither configured for nor capable of operating in the accused manner. (RX-012C (Larri) at Q&A 40-60; RX-01C (Conte) at Q&A 174-83.) Mr. Larri consistently explained that the statements relied upon by Nvidia describe only the theoretical flexibility of the GPU hardware by itself, and do not address how ARM driver software configures that hardware to actually operate within mobile devices. (*See e.g.*, Tr. at 629:3-11; 636:9-637:7, 639:14-25; 681:10-23.) Thus, NVIDIA’s references to ARM documents that do not describe how Samsung products are configured by the ARM driver software provide no basis to understand actual operation of these accused products.

PUBLIC VERSION

NVIDIA argues for the first time in its post-hearing brief that the Mali GPU hardware allegedly infringes even though the driver prevents use of [] for vertex processing. (See CIB at 106-109.) This argument was not made in NVIDIA's pre-hearing brief and is therefore waived. (See Order No. 2, Ground Rule 11.2.) But in any case, I find NVIDIA's argument not persuasive. The Asserted Claims are not pure apparatus claims. They contain functional language that must be shown to be satisfied in order to secure a finding of infringement. Moreover, the language of claim 1 explicitly requires that the "at least two sample types *are assigned* to the threads for multithreaded execution based on an allocation priority." (emphasis added.) This is active language that cannot be met by a passive chip. As discussed above, the evidence is clear that the Mali GPU by itself cannot assign anything; the driver configures and controls the GPU.¹¹

NVIDIA cites a number of cases in support of their new theory, but I find none of them support NVIDIA's argument. One such case is *Silicon Graphics, Inc. v. ATI Techs., Inc.*, 607 F.3d 784, 794 (Fed. Cir. 2010). While it does not support NVIDIA's argument, I find its contrast to the facts at issue here instructive. In *Silicon Graphics v. ATI*, the Federal Circuit found that the accused chip was "designed in such a way as to enable a user of that product to utilize the [allegedly infringing] function without having to modify the product." *Silicon Graphics, Inc. v. ATI Techs., Inc.*, 607 F.3d 784, 794 (Fed. Cir. 2010) (internal citations omitted). The absence of operating system software was found immaterial because "[n]othing in the record suggests that the Microsoft Windows operating system provides anything other than a way to activate the

¹¹ Although NVIDIA appears to be running away from the driver code, it seems to acknowledge its relevance. (See CIB at 55 ("[a]ll of the GPUs also come with 'drivers.'"), 57 ("This information is provided to the GPUs through a 'driver' supplied with the GPUs by the GPU vendor."), 99 (acknowledging that the "Job Manager receives 'jobs' from the device driver").)

PUBLIC VERSION

accused product.” *Id.* at 795. By contrast, this case presents a drastically different scenario, where the driver is necessary to configure and control the Accused Products with Mali GPUs. Without the driver software the Mali GPU is for all intents and purposes no more than just a proverbial paperweight.

(2) Do the Accused Products meet limitation 1[e]?

As discussed *supra*, there is no priority between vertex operations and fragment operations in the assignment of operations to threads in the Accused Products with Mali GPUs. The evidence shows the shader cores [] when passing threads from thread creators to the Tri-pipe. (RX-0012C (Larri) at Q&A 99.) The evidence shows the shader cores do not []. (*Id.* at Q&A 100; RX-01C (Conte) at Q&A 206; Tr. at 406:5-407:1.) The evidence shows the same [] is followed to select new threads to enter the Tri-pipe []. (RX-0012C (Larri) at 100.) Thus, the evidence shows Mali GPUs do not dynamically balance or control the number of samples types assigned to threads as required by limitation 1[e].

(3) Conclusion

I found above that the Accused Products with Mali GPUs do not satisfy either limitation 1[d] or limitation 1[e] of claim 1 of the ‘685 patent. Accordingly, I find for at least the reasons above that NVIDIA has failed to prove by a preponderance of the evidence that the Accused Products with Mali GPUs infringe claim 1.

b. Adreno GPUs

NVIDIA’s experts, Dr. Doggett and Dr. de la Inglesia, testified in detail that the Accused Products with Adreno (A3X and A4X) GPUs infringe claim 1 of the ‘685 patent. (CX-008C

PUBLIC VERSION

(Doggett) at Q&A 245-254; CX-007C (de la Iglesia) at Q&A 13-25, 32-38, 45-55, 62-68.)

Respondents argue that the Accused Products do not infringe, because they do not satisfy the following limitations of claim 1:

- (1) “wherein the at least two samples are assigned to the threads for multithreaded execution based on an allocation priority among the at least two sample types” (“limitation 1[d]”); and
- (2) “the thread control unit assigning one of the at least two sample types to the PCU based on the stored thread state data for dynamically balancing the number of samples assigned to the threads on a cycle to cycle basis” (“limitation 1[e]”).

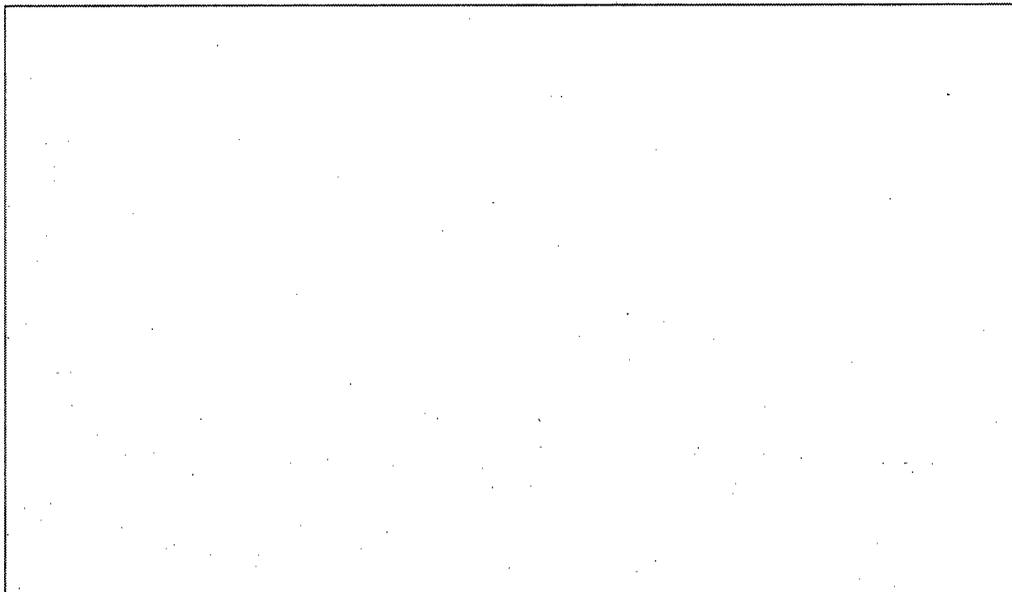
(RIB at 102-08.) The Staff only partially agrees with Respondents, arguing the Accused Products do not infringe claim 1 because they do not satisfy limitation 1[e]. (SIB at 42.)

(1) Do the Accused Products meet limitation 1[d]?

As confirmed at the hearing by Respondents’ expert Dr. Fussell and Qualcomm engineer, Mr. Du, the Accused Products with Adreno GPUs include at least two sample types that are assigned to threads based on an allocation priority among the at least two sample types.

[

]



PUBLIC VERSION

(JX-122C at 110 (annotated); *see also* RX-006C (Du) at Q&A 76-78 and 127; Tr. (Du) at 563:6-9, 560:8-10, 550:22-551:23, 544:16-545:12, 560:8-10.) As clearly set forth in the figure above,

[]

(JX-122C at 110.) At the hearing Mr. Du, a Qualcomm engineer and author of the []

RTL code and relevant technical documentation, testified:

[

]

(Tr. (Du) at 563:6-9.) [

] (*Id.*; *see also* 550:22-551:33; 542:5-18; 542:12-15 ([

]).

Respondents' expert Dr. Fussell [

]. (Tr. at 804:3-14, 805:23-806:7, 807:22-809:3, 806:17-807:2.) For example,

Dr. Fussell testified:

[

]

(Tr. (Fussell) at 805:23-806:7.) Dr. Fussell also testified:

[

]

(Tr. at 804:3-14.) [

] Accordingly, I find NVIDIA has shown by a preponderance of the evidence that the Accused Products with Adreno GPUs meet limitation 1[d].

In their post-hearing brief, Respondents mint a new term – “system condition rules” – to rebrand what their documents and witnesses admit are [] This term accompanies a new non-infringement theory that the assignment of samples to threads is not

[

] (See RIB at 105-06.) This argument never appeared in Respondents’ pre-hearing brief and has therefore been waived. (See Order No. 2 (Ground Rule 11.2).) In any event the argument is without merit.

(2) Do the Accused Products meet limitation 1[e]?

The only evidence or explanation NVIDIA provides in its initial post-hearing brief that the Accused Products with Adreno GPUs meet limitation 1[e] is:

PUBLIC VERSION

In addition, the Adreno GPUs can [

]. *Id.*; JX-122C.110; Hrg. Tr. (Du) at 561:16-562:13, RX-6C (Du) at Q. 126.

(CIB at 91.)¹²

Dr. Doggett provides nothing but a series of citations to documents without explanation and a conclusory statement to support his assertion that A3X and A4X “dynamically balanc[e] the number of samples assigned to the threads on a cycle to cycle basis.” (*See* CX-08C (Doggett) at Q&A 58-60; see also Q&A at 247-48, 296-97.) Dr. Doggett’s analysis is so lacking it is not even clear what, if any, functionality Dr. Doggett is alleging meets this element. NVIDIA’s citations to Dr. de la Iglesia fair even worse as they do not discuss dynamically balancing the number of samples at all. (*See* CX-07C (de la Iglesia) at Q&A 13-15, 45-47.) And without any explanation it is entirely unclear how the citations to Dr. Du show this limitation is met. Further, the evidence shows [

] (RX-3498C (Fussell) at Q&A 58-60.) .) Thus, the evidence shows Adreno GPUs do not dynamically balance or control the number of samples types assigned to threads as required by limitation 1[e].

(3) Conclusion

I found above that the Accused Products with Adreno GPUs do not satisfy limitation 1[e] of claim 1 of the ‘685 patent. Accordingly, I find for at least the reasons above that NVIDIA has failed to prove by a preponderance of the evidence that the Accused Products with Adreno GPUs infringe claim 1.

¹² The “*Id.*” in the citation at the end of the quote references “CX-8C (Doggett) at Q. 247-48 and 296-97; CX-7C (de la Iglesia) at Q. 13-15 and 45-47.” (CIB at 91.)

PUBLIC VERSION

2. Claim 15

a. Mali GPUs

NVIDIA argues that the disputed claim elements of 1[d] and the latter clause of 1[e] are met by Mali GPUs for the same reasons as the corresponding claim elements 15[b] and 15[preamble], respectively. (CIB at 112.) Respondents do not raise any new arguments as to claim 15 other than those already discussed in connection with claim 1. Therefore, for the same reasons discussed in detail, *supra*, with regard to claim 1, I find NVIDIA has failed to show that the Accused Products with Mali GPUs satisfy all the limitations of claim 15. Accordingly, I find NVIDIA has failed to show by a preponderance of the evidence that the Accused Products with Mali GPUs infringe claim 15 of the '685 patent.

b. Adreno GPUs

Claim 15 requires that “each of the threads is allocated sample types based on a priority of the sample type,” meaning that “each of those assignments for each of the threads” must be based on a priority assigned to the sample type. (Tr. (Fussell) at 819:21-820:7.) NVIDIA fails to address the word “each” in its analysis of claim 15. In fact, NVIDIA does not even identify “each” as part of the disputed phrase. (CIB at 92.)

Regardless, the evidence shows that this claim element is not met. Mr. Du explained how thread assignment/allocation works in Adreno A3X and A4X: [

] (Tr. (Du) at 542:10-

15.) Mr. Du explained [

] (*Id.* at 562:6-13; *see*

also id. at 562:17-21 ([

]).

Dr. Fussell explained [

PUBLIC VERSION

] (Tr. (Fussell) at 806:3-7; 804:3-12; *see also* RX-3498C (Fussell) at Q&A 63.) [

] it cannot be said, and certainly NVIDIA has failed to show, that each of the assignments for each of the threads is based on a priority assigned to the sample type as required by claim 15. Thus, for at least the reasons above, I find NVIDIA has failed to show that the Accused Products with Adreno GPUs meet all the limitations of claim 15. (*See* RX-3498C (Fussell) at Q&A 62-63.) Accordingly, I find NVIDIA has failed to show by a preponderance of the evidence that the Accused Products with Adreno GPUs infringe claim 15.

D. Domestic Industry - Technical Prong

NVIDIA contends that its domestic industry products practice claim 15 of the '685 patent. (CIB at 129.) To that end, NVIDIA's expert, Dr. Doggett, testified in detail that NVIDIA's products embodying the Kepler, Maxwell, and Fermi architectures practice the '685 patent. (CX-008C (Doggett) at Q&A 633-685). Neither Respondents nor the Staff contests this point. Thus, I find the evidence shows that NVIDIA's domestic industry products embodying the Kepler, Maxwell, and Fermi architectures practice claim 15 of the '685 patent. Accordingly, I find that NVIDIA has satisfied the technical prong of the Domestic Industry requirement with respect to the '685 patent.

E. Invalidity

1. Anticipation (35 U.S.C. § 102) – The ATI Unified Shader

The Parties' Positions

Respondents contend that under Section 102(g)(2) the ATI Unified Shader anticipates claims 1 and 15 of the '685 patent. (RIB at 109-123). Specifically, Respondents contend that the

PUBLIC VERSION

evidence shows that Andrew Gruber, then a principal architect at ATI Technologies, and a small group of engineers at ATI conceived of the ATI Unified Shader in November 2000—roughly two and a half years before the filing of the '685 patent and twenty months before NVIDIA's alleged conception of the '685 patent. (*Id.* at 110.) Respondents argues Mr. Gruber and others at ATI were diligent in reducing the invention to practice (1) through patent applications that matured into U.S. Patent Nos. 7,239,322 and 6,897,871 (“the ATI Patents”) and (2) through ATI's Xenos chip, which launched as part of the Microsoft Xbox 360 in November 2005. (*Id.*) Respondents also contend that that in addition to being conceived before the '685 patent and diligently reduced to practice, the ATI Unified Shader meets each limitation of claims 1 and 15. Thus, Respondents argue the ATI Unified Shader anticipates claims 1 and 15 of the '685 patent.

NVIDIA argues that the ATI Unified Shader is not prior art. NVIDIA argues that ATI never reduced the invention to practice, much less was “diligent” in reducing the ATI Unified Shader to practice as required under the law. Moreover, NVIDIA argues that Respondents failed to prove by clear and convincing evidence that each of the claim elements was conceived as of the November 2000 date of conception alleged by Respondents. Specifically, NVIDIA argues that Respondents failed to prove conception with regard to the allocation priority elements of the Asserted Claims, limitations 1[d] and 15[b]. (CIB at 113-122).

The Staff argues Respondents have not shown conception of the asserted claims prior to June 27, 2003 with respect to the dynamic load-balancing limitation of those claims. Thus, the Staff argues, the ATI Unified Shader does not qualify as prior art to the '685 patent. The Staff also argues that Respondents have failed to show that the ATI Unified Shader anticipates the asserted claims. In particular, the Staff argues that the ATI Unified Shader does not disclose claim 1[d], claim 1[e], or claim 15[b] of the '685 patent.

PUBLIC VERSION

Discussion

A patent claim is invalid if, before the patentee's invention, "the invention was made in this country by another inventor who had not abandoned, suppressed, or concealed it." 35 U.S.C. § 102(g)(2). A party may establish prior invention under Section 102(g)(2) by proving that another inventor (1) was the first to conceive of the invention and (2) then exercised reasonable diligence in reducing that invention to practice. *E.g., Tyco Healthcare Grp. LP v. Ethicon Endo-Surgery, Inc.*, 774 F.3d 968, 975 (Fed. Cir. 2014).

a. Conception

"The test for conception is whether the inventor had an idea that was definite and permanent enough that one skilled in the art could understand the invention." *Tyco Healthcare Grp. LP*, 774 F.3d at 974-75 (internal quotation marks omitted). Although an inventor must have "a specific, settled idea, the inventor need not know that his invention will work for conception to be complete. He need only show that he had the idea." *Id.*

The '685 patent has an effective filing date of June 30, 2003. (JX-005.) In its opening brief, NVIDIA argues that the '685 patent was conceived on July 3, 2002 and constructively reduced to practice on June 30, 2003. (*See* CIB at 113.) However, the only support NVIDIA provides for the alleged July 3, 2002, date of conception is a citation to its pre-hearing brief. NVIDIA's citation to its pre-hearing brief violates my Ground Rules and therefore I am striking the citation. (*See* Ground Rule 15.1.1 ("Any factual or legal issues not addressed in the post-hearing briefs shall be deemed waived. The post-hearing briefs shall not incorporate anything by reference, but may include pinpoint citations to legal authority or the evidentiary record.").) Thus, I find NVIDIA provides no support for its alleged July 3, 2002, date of conception. Accordingly, to qualify as prior art against the '685 patent under 102(g)(2), the relevant features of the ATI Unified Shader that meet the claim limitations of the Asserted Claims of the '685

PUBLIC VERSION

patent must have been conceived prior to June 30, 2003 filing date of the '685 patent, and subsequently diligently reduced to practice.

Here, Respondents argue only that the ATI Unified Shader was conceived in November 2000. (See RIB at 108-109; RX-9C (Fussell) at 92.) The only exhibit from November 2000 relied on by Respondents is Exhibit RX-376C, titled "R400 Architecture Proposal" (version 0.1). Respondents rely on a number of materials dated between 2001 and 2003 to allegedly prove conception. In their reply brief Respondents attempt to argue that even though these materials are dated after November 2000, they still prove a conception date earlier than that of the '685 patent. Respondents did not make this argument in its pre-hearing brief or initial post-hearing brief. To the contrary, Respondents and their expert consistently state that the ATI Unified Shader was conceived in November 2000. Thus I find Respondents have waived any argument that asserts the ATI Unified Shader had a date of conception other than November 2000.

Respondents improperly incorporate in their post-hearing reply brief portions of their pre-hearing brief. (RRB at 91 ("For these reasons, as well as those stated in Respondents' Pre-Hearing Brief at pages 65 through 69, Respondents have demonstrated prior conception.") Incorporating by reference in post-hearing briefing is forbidden as it circumvents the page limits I have set and unfairly prejudices the other parties. (See Ground Rule 15.1.1 ("Any factual or legal issues not addressed in the post-hearing briefs shall be deemed waived. The post-hearing briefs shall not incorporate anything by reference, but may include pinpoint citations to legal authority or the evidentiary record.").) Accordingly, I am striking the citation to Respondents' pre-hearing brief.

As discussed in more detail below, I find Respondents' have failed to show the claim limitations of the Asserted Claims of the '685 patent were conceived in November 2000 and thus

PUBLIC VERSION

failed to prove the ATI Unified Shader constitutes prior art to the '685 patent. For example, in arguing conception of the "at least one multiprocessing unit" limitation of claim 1, the only documentary evidence cited by Respondents is Exhibit RX-426C, titled, "*R400 Sequencer Specification SQ*" (Version 2.02). (See RX-009C (Fussell) at Q&A 123.) However, this document was not created until May 13, 2002, and Respondents make no effort to tie this document back to November 2000. (RX-426C.) Thus, Respondents have failed to show this element was conceived in November 2000. Likewise, in arguing conception of a thread control unit that stores state information for each thread as required by claims 1 and 15, Respondents' expert, Dr. Fussell, again only referred to the May 13, 2002, R400 Sequencer Specification SQ. (See RX-009C (Fussell) at Q&A 124.) While Respondents cite to page 9 of RX-376C to allegedly show conception of this element, Respondents provide absolutely no discussion or argument explaining how the text on page 9 discloses conception of a thread control unit that stores state information for each thread. (See RIB at 112.) Thus, Respondents have failed to show these elements were conceived in November 2000.

With regard to the allocation priority elements of the Asserted Claims, limitations 1[d] and 15[b], Respondents' expert, Dr. Fussell, testified that there were three ways the ATI Unified Shader met these limitations. First, with reference to the R400 Architecture Proposal dated November 2000 (Exhibit RX-376C), Dr. Fussell testified that [

] (RX-009C at Q&A 125; RX-376C at 10.) Even taking Dr. Fussell's testimony as true that [] such testimony is insufficient to show conception of limitations 1[d] and 15[b]. These limitations explicitly require more than just a

PUBLIC VERSION

priority of one sample type over another. Limitation 1[d] requires that the “at least two sample types are assigned to the threads for multithreaded execution based on an allocation priority among the at least two sample types” and limitation 15[b] requires “each of the threads being allocated the sample types based on a priority assigned to the sample type.” Dr. Fussell does not provide any discussion or explanation of how the disclosure in the November 2000 R400 Architecture Proposal shows the two sample types (*i.e.*, vertices and pixels) are assigned to the threads based on the alleged disclosed priority of vertices over pixels or how each of the threads are allocated vertices and pixels based on the alleged disclosed priority of vertices over pixels. Thus, I do not find Dr. Fussell’s testimony shows conception of limitations 1[d] and 15[b] based on the November 2000, R400 Architecture Proposal.

Second, Dr. Fussell points to Dr. Goldberg’s analysis of the source code for the ATI Unified Shader, arguing that [] (See RX-009C (Fussell) at Q&A 125.)

However, Dr. Fussell himself discounts this theory stating that [

] (*Id.*) Thus, I do not find Respondents have shown prior conception based on the [] source code. Moreover, as the source code was not drafted until [], and Respondents make no attempt to tie the source code back to November 2000, the cited source code cannot possibly support the November 2000 date of conception argued by Respondents.

Third, Dr. Fussell contends that [] in the ATI Unified Shader shows an allocation priority scheme. (RX-09C (Fussell) at Q&A 125.) Dr. Fussell, however, never alleges, or provides any support for, the [] being conceived in

PUBLIC VERSION

[]. (*Id.*; *see also* CX-2130C (Doggett) at Q&A 126-129.) In fact, the earliest document to which Dr. Fussell cites that discusses [] is from May 13, 2002. (RX-009C (Fussell) at Q&A 125.) Moreover, the evidence suggests the brief mention of [] in the May 2002 sequencer document does not provide an enabling disclosure that one of ordinary skill could use to implement the []. (*See* CX-2130C (Doggett) at Q&A 132-136; Tr. (Gruber) at 757:11-22 []) Accordingly, the May 2002 sequencer document cannot support the November 2000, date of conception argued by Respondents.

[] (Tr. at 755:1-757:25, 763:8-764:11; CX-2130C (Doggett) at Q&A 133.) [] (*Id.*) []

[] (*See* CDX-68C-016; RX-3413C; CX-2130C (Doggett) at Q&A 134-135; Tr. (Gruber) at 763:8-764:11; Tr. (Doggett) at 414:13-24.)

With regard to the dynamic load balancing limitations of the Asserted Claims, limitation 1[e] and the corresponding portions of limitations 15[preamble] and 15[b], Dr. Fussell testified the ATI Unified Shader included:

[

]

(RC-0009C (Fussell) at Q&A 128.) While a stated goal of the ATI Unified Shader was to [] between vertex and pixel processing, Dr. Fussell did not show that the ATI Unified Shader achieved this goal by “the thread control unit assigning one of the at least two sample types to the PCU based on the stored thread state data for dynamically balancing the number of samples assigned to the threads on a cycle to cycle basis,” as the claim requires. (CX-2130C (Doggett) at Q&A 160-162.) Moreover, to the extent that Dr. Fussell relies upon Dr. Goldberg’s analysis of thread buffer RTL code, such code was added after May 29, 2002 and thus cannot support a date of conception before the '685 patent.

Thus, I find Respondents have no credible evidence, let alone clear and convincing evidence, that individuals at ATI possessed a “complete and operative invention” or possessed “an operative method of making the invention” prior to June 30, 2003 effective filing date of the '685 patent.

b. Reduction to Practice

In addition to showing prior conception, Respondents must also show that ATI was diligent in reducing the ATI Unified Shader to practice. 35 U.S.C. § 102(g)(2) (“In determining priority of invention under this subsection, there shall be considered not only the respective dates of conception and reduction to practice of the invention, but also the reasonable diligence of one who was first to conceive and last to reduce to practice, from a time prior to conception by the other.”). Reduction to practice may be actual or constructive. *In re Katz*, 687 F.2d 450, 454 (Cust. & Pat.App., 1982) (“This section of the statute embodies the principle that to be entitled to a patent one must be the first to have made the invention. However, prior conception of the invention by another does not defeat one's right. No possible barrier is created by s 102(g)

PUBLIC VERSION

unless another has either actually reduced the invention to practice or has constructively reduced it to practice by filing a patent application.”)

“In order to establish an actual reduction to practice, the inventor must prove that: (1) he constructed an embodiment or performed a process that met all the limitations ... and (2) he determined that the invention would work for its intended purpose.” *Cooper v. Goldfarb*, 154 F.3d 1321, 1327 (Fed.Cir.1998). “Testing is required to demonstrate reduction to practice in some instances because without such testing there cannot be sufficient certainty that the invention will work for its intended purpose.” *Slip Track Sys., Inc. v. Metal-Lite, Inc.*, 304 F.3d 1256, 1267 (Fed.Cir.2002).

Constructive reduction to practice, on the other hand, occurs with the filing of a patent application that discloses the invention. *Tyco Healthcare Group LP v. Ethicon Endo-Surgery, Inc.* 774 F.3d 968, 975 (Fed. Cir. 2014) (“The filing of a patent application is constructive reduction to practice of the invention disclosed therein.”) Complainant argues that to constitute a constructive reduction to practice the patent application must include a claim drawn to the invention, but there is no such requirement under the law. In order to constitute constructive reduction to practice, the law requires only that the patent application disclose the invention in such manner as to meet “the written description and enablement requirements of 35 U.S.C. § 112 ¶ 1.” *Frazer v. Schlegel*, 498 F.3d 1283, 1287 (Fed.Cir.2007); *Bigham v. Godtfredsen*, 857 F.2d 1415, 1417 (Fed.Cir.1988) (“To serve as constructive reduction to practice, the disclosure of the subject matter ... must meet the requirements of 35 U.S.C. § 112, first paragraph.”); *see also* *Rengo Co. Ltd. v. Molins Mach. Co., Inc.*, 657 F.2d 535, 548 (3d Cir.1981) (“[A]n American application will be regarded as a reduction to practice only if it describes the invention with the

PUBLIC VERSION

particularity required by Section 112.”); *Automotive Technologies Intern., Inc. v. Siemens VDO Automotive Corp.*, 669 F.Supp.2d 836, 846 (E.D.Mich. 2009).

Respondents argue that the ATI Unified Shader was constructively reduced to practice through patent applications that matured into U.S. Patent Nos. 7,239,322 and 6,897,871 and actually reduced to practice through ATI’s Xenos chip. (RIB at 110, 118.)

(1) ATI Patents

Respondents rely on U.S. Patent Nos. 7,239,322 and 6,897,871 (“the ATI Patents to show a constructive reduction to practice of the ATI Unified Shader. U.S. Patent No. 7,239,322, is titled, “Multi-thread graphic processing system” and U.S. Patent No. 6,897,871 is titled, “Graphics processing architecture employing a unified shader.” Respondents do not rely on the ATI Patents individually, but rather together to show a reduction to practice. The case law, however, seems to uniformly describe constructive reduction to practice in terms of a single patent application. *See e.g., Ariad Pharms., Inc. v. Eli Lilly & Co.*, 598 F.3d 1336, 1351 (Fed.Cir.2010) (en banc) (The disclosure must “reasonably convey[] to those skilled in the art that the inventor had possession of the claimed subject matter as of the filing date.”); *id.* (Possession means “possession as shown in the disclosure” and “requires an objective inquiry into the four corners of the specification from the perspective of a person of ordinary skill in the art.”). Respondents cite no case law supporting the proposition that two separate patent applications, directed to separate inventions, filed months apart, that mature into two separate patents, can be considered together to constitute a single constructive reduction to practice.

A constructive reduction to practice is a manifestation in writing of the definite and permanent idea required to show conception. Therefore, it too must describe the invention in a definite way. *Ariad Pharm., Inc.* (quoting *Vas-Cath Inc. v. Mahurkar*, 935 F.2d 1555, 1562–63 (Fed.Cir.1991) (The test for sufficiency of a written description is “whether the disclosure clearly

PUBLIC VERSION

‘allow[s] persons of ordinary skill in the art to recognize that [the inventor] invented what is claimed.’”).) Such is also required to satisfy section 112, para 1. *Streck, Inc. v. Research & Diagnostic Systems, Inc.*, 665 F.3d 1269, 1285 (Fed. Cir. 2012) (“A constructive reduction to practice that in a definite way identifies the claimed invention can satisfy the written description requirement.”) (internal quotations omitted). Here, Respondents do not point to any claim, embodiment, or distinct description in either of the ATI Patents to show construction reduction to practice. Rather, as previously stated, Respondents must point to bits and pieces from both ATI Patents to support their argument for constructive reduction to practice. Under such circumstances, I cannot find Respondents have clearly identified the alleged ATI Unified Shader invention in a definite way such that one of ordinary skill in the art would understand the inventors to be in possession of an invention comprising all the limitations of the Asserted Claims of the ‘685 patent. (*See also*, CX-2130 (Doggett) at Q&A 180-186.) Accordingly, for at least the reasons above, I find Respondents have failed to show that the ATI Patents constitute a constructive reduction to practice of the alleged ATI Unified Shader invention or that the ATI Patents disclose all the limitations of the Asserted Claims of the ‘685 patent.

(2) Xenos chip

Respondents contend the ATI Xenos chip constitutes an actual reduction to practice of the ATI Unified Shader. Yet none of the evidence Respondents cite is tied to an actual Xenos chip. Much of the evidence cited by Respondents is the same R400 documentation Respondents rely on to show conception. While there is evidence to suggest the R400 was eventually renamed Xenos, *see* RX-362C at 6 that does not mean the features discussed in R400 development documents were implemented in the Xenos chip as Respondents would like me to believe. Because Respondents failed to tie the documentation on which they rely to an actual

PUBLIC VERSION

Xenos chip, I find Respondents have not shown by clear and convincing evidence the operation and features of Xenos chip as would be required to prove an actual reduction to practice.

Moreover, even if Respondents reliance on the R400 development documents to show the operation of a Xenos chip was permissible, which it is not, for the reasons discussed, *supra*, with regard to conception, Respondents have failed to show by clear and convincing evidence that the ATI Unified Shader includes all the limitations of the Asserted Claims of the '685 patent.

c. Conclusion

For at least the reasons discussed above, I find Respondents have failed to prove by clear and convincing evidence that the ATI Unified Shader was conceived prior to the invention of the '685 patent, that the ATI Unified Shader was reduced to practice, either constructively or actually, or that the ATI Unified Shader includes all the limitations of the Asserted Claims of the '685 patent. Accordingly, I find the ATI Unified Shader does not anticipate the Asserted Claims of the '685 patent.

2. Obviousness – 35 U.S.C. 103

The Parties' Positions

NVIDIA argues that as a threshold issue, Respondents do not even *contend* that Selzer and Amanatides are graphics processors and thus, any such contention has been waived. (CRB at 112.) NVIDIA contends Respondents only argue that Van Hook is a graphics processor and that the *combination* of Van Hook with Selzer or Amanatides would result in a graphics processor. (*Id.*) NVIDIA argues a person of ordinary skill would not look to combine single-threaded general purpose processors with a multithreaded graphics processor from two different fields of art. (*Id.*)

NVIDIA also argues that the asserted references fail, alone or in combination, to suggest or disclose a multithreaded processor that processes two sample types simultaneously or that

PUBLIC VERSION

dynamically balances the sample types. (*Id.*) NVIDIA contends that Respondents admit that each processor in Amanatides and Selzer is only the equivalent of a single thread and therefore each reference lacks a “multithreaded processing unit,” the first and most basic element of the ’685 patent claims. (*Id.*) NVIDIA argues Van Hook discloses a multithreaded processor that is able to operate only on one sample type and that a person of ordinary skill in the art would not seek to combine a single-threaded general purpose processor able to operate on one sample type with a multithreaded graphics processor able to operate on one sample type, to create a multithreaded graphics processor that is able to concurrently operate on multiple sample types. (*Id.* at 112-113.) NVIDIA argues these systems are fundamentally different and incompatible systems. (*Id.* at 113.) Accordingly, NVIDIA argues, none of Van Hook, Selzer and Amanatides discloses a multithreaded processor that can process multiple sample types simultaneously as required by at least claim limitations 1[preamble], 1[d], 1[e], 15[preamble] and 15[a]. (*Id.*) Thus, NVIDIA argues, the processors also cannot balance or adjust the number of threads assigned to each sample type (required by claim limitations 1[e] and 15[preamble]). (*Id.*)

NVIDIA contends that Dr. Dally explained in significant detail the vast differences between multithreaded processors, such as Van Hook, and the multiprocessor systems disclosed in Selzer and Amanatides, and why a person of ordinary skill would not seek to combine them. (*Id.*) NVIDIA asserts that Respondents chose not to challenge Dr. Dally on his expert opinion. (*Id.*) NVIDIA argues that a single-threaded processor goes through a time-consuming process to switch threads. (*Id.*) NVIDIA contends that it must save thread state data for an executing thread, retrieve state data for a different thread, and eventually begin processing the new thread – a process that takes multiple clock cycles. (*Id.*) NVIDIA argues a multithreaded processor

PUBLIC VERSION

stores state data for multiple threads at the same time, so it can quickly switch between processing different threads on every cycle. (*Id.*)

NVIDIA argues that Selzer and Amanatides disclose a system with multiple single-threaded processors each capable of operating on one sample type. (*Id.*) NVIDIA argues Van Hook discloses a multithreaded processor only capable of operating on one sample type. (*Id.*) Thus, NVIDIA argues, as Dr. Dally opined, one of ordinary skill would not try to combine such incompatible systems, and even if someone did, it still would not result in a multithreaded graphics processor that can process vertex and pixel operations simultaneously and perform dynamic load balancing between sample types as claimed in the '685 Patent. (*Id.*)

NVIDIA also argues that the asserted references fail, alone or in combination, to suggest or disclose the assignment of samples to threads in a multithreaded processing unit or the assignment of samples to threads based on a priority. (*Id.* at 114.) NVIDIA contends that claims 1 and 15 of the '685 patent require the assignment of samples to threads in a multithreaded processor (claim limitations 1[a], 1[d], 15[preamble] and 15[b]), and basing the assignment on a priority assigned to a sample type. (*Id.*) NVIDIA argues Van Hook, Selzer and Amanatides all fail to disclose the assignment of multiple sample types to threads in a multithreaded processor. (*Id.*) NVIDIA asserts that absent this basic concept, even if one of skill in the art tried to combine the references, they still would not achieve a multithreaded processor that allocates samples to threads based on an allocation priority. (*Id.*)

NVIDIA argues that Van Hook discloses a multithreaded processor capable of processing only one sample type, either vertices or pixels. (*Id.*) Thus, NVIDIA argues, Van Hook does not – and cannot – disclose any assignment or allocation of threads based on priority given to a sample type. (*Id.*) Yet NVIDIA contends, Respondents argued that Van Hook suggests how to

PUBLIC VERSION

“implement a priority assigned to classes of work tasks in a multithreaded system.” (*Id.*)

NVIDIA argues Van Hook does not disclose how to assign samples to threads and since Van Hook does not disclose processing multiple sample types in the same execution unit simultaneously, it cannot even suggest assigning threads of different sample types based on a priority. (*Id.*) In fact, NVIDIA contends, Van Hook teaches away from the invention by suggesting that vertex and pixel processing should be performed by different arithmetic units. (*Id.*)

NVIDIA argues that nothing in Selzer discloses the assignment of sample types to hardware thread contexts based on a priority given to a particular sample type – not even in a single-threaded processor, much less a multithreaded processor. (*Id.* at 115.) NVIDIA asserts that Selzer describes a multiprocessor system comprised of multiple single-threaded digital signal processors (“DSPs”). (*Id.*) According to NVIDIA, when no rendering work is available, a pixel rendering module may switch its function to become a “geometry module” to process vertices, but that this is essentially “work-stealing.” (*Id.*) NVIDIA contends that the use of idle rendering modules to process vertices is not equivalent to using a priority to assign sample types to the multiple hardware thread contexts of a multithreaded processor. (*Id.*)

NVIDIA argues that Amanatides also fails to disclose the allocation or assignment of sample types based on a priority given to a sample type, even in the context of single-threaded processors. (*Id.*) NVIDIA argues that Amanatides only describes a multiprocessor with multiple single-threaded processors, which gives priority over a shared bus to render messages (“RMs”), or pixel tasks. (*Id.*) NVIDIA argues that this priority is implemented to avoid deadlock of the processors and not to assign a sample to a thread. (*Id.*) Moreover, NVIDIA argues the RMs are broadcast to all of the processors at once rather than allocated or assigned to a particular

PUBLIC VERSION

processor or thread based on a priority, as required by the asserted claims of the '685 Patent.

(Id.)

NVIDIA asserts that recognizing that all three asserted references fail to suggest the use of a priority assigned to a sample during thread allocation – even in the context of single-threaded processors – Respondents resort to introducing a fourth reference (“Fiske”) to imply that all multithreaded processors must use a priority in order to allocate or assign threads. *(Id.)*

NVIDIA contends this is incorrect, because Respondents rely only on the combination of Van Hook with Selzer or Amanatides to support their claim of obviousness and Respondents cannot now introduce yet another reference to support their claim of obviousness. *(Id. at 115-116.)*

Also, NVIDIA contends Fiske only teaches a way in which threads can be scheduled, not assigned to sample types. *(Id.)* According to NVIDIA, the threads have already been assigned to samples when they are scheduled for execution in Fiske. *(Id.)*

Respondents argue Selzer in combination with Van Hook render obvious claims 1 and 15 of the '685 patent. Respondents contend that the combination of Selzer with Van Hook is straight-forward with each rendering module (*i.e.*, physical processor) in Selzer mapping to a thread (*i.e.*, logical processor) in Van Hook with the execution pipeline of Van Hook taking on the capability of executing instructions for both vertex processing and pixel processing, and the priority rule in Selzer of pixel over vertex serving as a thread allocation priority rule in the combined multithreaded system. (RIB at 129.) Respondents maintain that implementing the priority rule of Selzer as a thread allocation priority rule would be a straightforward application of using what Dr. Dally described in 1995 as “[t]hread prioritization,” “a simple means of guiding the scheduling of threads.” *(Id. at 129-30.)* Respondents argue that the combined system would be a graphics processor for simultaneous multithreaded execution as recited in the

PUBLIC VERSION

preamble of claim 1, with the modified PGIP and interleaver illustrated in Figures 2 and 3 of Van Hook serving as the “multithreaded processing unit that includes a thread control unit” as recited in claim elements 1[a] and 1[b], and Program Counters 202A through 202N and Registers 205A through 205N serving as the thread storage resource recited in claim element 1[c]. (*Id.* at 130.) Respondents allege the combined system would use Selzer’s pixel over vertex priority for thread allocation, meeting claim element 1[d], which would result in load balancing under claim element 1[e], at least under Dr. Doggett’s infringement interpretation. (*Id.*) Respondents also argues that should claim element 1[e] be interpreted to require an execution priority, including such a priority would be an obvious design choice in view of Fiske. (*Id.*) Respondents contend it was known to those of ordinary skill in the art, as confirmed by the Fiske paper, that all multithreaded architectures must make decisions at both the thread allocation stage and the execution stage and “both decisions are important.” (*Id.*) Respondents claim that given the recognition that both types of priorities are important, whether to use an execution priority, an allocation priority, or both was a matter of design choice. (*Id.*)

Respondents next argue the combination of Selzer with Van Hook includes the graphics processor of the preamble of Claim 15 for the reasons discussed with regard to the preamble and claim elements 1[a] and [b]. (*Id.*) Respondents assert that the thread control unit configured to store pointers to program instructions of claim element 15[a] is met by the program counters of Van Hook, the PCU of claim element 15[b] is met by the execution pipeline of Van Hook modified, like the rendering processors of Selzer, to be a unified shader, and the “each of the threads being allocated...” is met for the same reasons as claim element 1[d], at least under Nvidia’s interpretation. (*Id.* at 130-31.)

PUBLIC VERSION

Respondents also say Amanatides in combination with Van Hook render obvious claims 1 and 15 of the '685 patent. Respondents contend the combination of Amanatides with Van Hook is the same as Selzer with Van Hook, with each unified shader processor in Amanatides, labeled "G/R," (*i.e.*, physical processor) in Amanatides, mapping to a thread (*i.e.*, logical processor) in Van Hook. (*Id.* at 131.) Respondents argue that the execution pipeline of Van Hook takes on the capability executing instructions for both vertex processing and pixel processing as disclosed in Amanatides, and the priority rules in Amanatides of pixel = high priority and vertex = low priority serves as a thread allocation priority rule in the combined multithreaded system, at least under Nvidia's infringement assertions. (*Id.*) Therefore, Respondents assert, each element of claims 1 and 15 are met for the same reasons discussed with regard to the Selzer combination with Van Hook. (*Id.*)

The Staff argues that while Respondents have demonstrated that multithreaded architectures and dynamic load balancing were well known in the prior art, and were used in graphics processors, Respondents have not shown that the asserted prior art references, either alone or in combination, teach a person of ordinary skill in the art at the time of the invention the following limitations of claims 1 and 15 of the '685 patent: (1) "the at least two sample types are assigned to the threads for multithreaded execution based on an allocation priority among the at least two sample types" (claim 1[d]); (2) "the thread control unit assigning one of the at least two sample types to the PCU based on the stored thread state data for dynamically balancing the number of samples assigned to the threads on a cycle to cycle basis" (claim 1[e]); and (3) "each of the threads being allocated the sample types based on a priority assigned to the sample type" (claim 15[b]). (SIB at 47.)

PUBLIC VERSION

Discussion

Respondents argue that the Asserted Claims of the '685 patent are obvious in light of an article titled "*Dynamic Load Balancing within a High Performance Graphics System*" authored by Selzer (hereinafter "Selzer") in view of U.S. Patent 7,847,803 to Van Hook (hereinafter "Van Hook"). Respondents also argue that the Asserted Claims of the '685 patent are obvious in light of an article titled, "*A simple, flexible, parallel graphics architecture*" authored by Amanatides et al. (hereinafter "Amanatides") in view of Van Hook.

NVIDIA argues at the outset that one of ordinary skill in the art with experience in graphics processing would not consider single-threaded general purpose processors as relevant prior art to the '685 patent, which according to NVIDIA "explicitly claims features of a multithreaded graphics processor." (CRB at 111-112.) NVIDIA argues neither Selzer nor Amanatides are graphics processors. I disagree.

I have found herein that the term "graphics processor" in the preambles of the Asserted Claims is not limiting and even if it were that its proper construction would be "hardware capable of processing graphics data." (*See supra*, at VI.B.2.) Thus, under this construction, general purpose processors that perform graphics operations such as those disclosed in Selzer and Amanatides would be relevant art. Moreover, the legal standard for determining "analogous art" includes not only "whether the art is from the same field of endeavor, regardless of the problem addressed," but "if the reference is not within the field of the inventor's endeavor, whether the reference still is reasonably pertinent to the particular problem with which the inventor is involved." *Scientific Plastic Prods., Inc. v. Biotage Ab*, 766 F.3d 1355, 1358 (Fed. Cir. 2014). Thus, even if Selzer or Amanatides was outside of the field of endeavor, which they

PUBLIC VERSION

are not, I would still find them reasonably pertinent to the particular problem to be solved, namely increasing processor utilization in a graphics pipeline.

Although I disagree with NVIDIA's argument above, I ultimately agree with NVIDIA and the Staff and find for the reasons discussed below that Selzer and Van Hook, individually or in combination, fail to disclose, or even suggest, a number of limitations of the Asserted Claims. Likewise, I find for the reasons discussed below that Amanatides and Van Hook, individually or in combination, fail to disclose, or even suggest, a number of limitations of the Asserted Claims. To prove obviousness, Respondents must show by clear and convincing evidence that the alleged combination discloses all the limitations of the Asserted Claims of the '685 patent and that one of ordinary skill in the art would have reason to combine the prior art elements. *See, e.g., Unigene Labs., Inc. v. Apotex, Inc.*, 655 F.3d 1352, 1360 (Fed.Cir.2011) (citing *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 418, 421 (2007)) ("Obviousness requires more than a mere showing that the prior art includes separate references covering each separate limitation in a claim under examination. Rather, obviousness requires the additional showing that a person of ordinary skill at the time of the invention would have selected and combined those prior art elements....").

a. Selzer in view of Van Hook

Discussion

Selzer was published in 1993, more than a year before the effective filing date of the '685 patent, and is thus prior art under 35 U.S.C. § 102(b). (RX-536.) Van Hook issued on December 10, 2010 on an application filed on July 26, 2000. (RX-350) Because the filing date of Van Hook is before the effective filing date of the '685 patent it is prior art under 35 U.S.C. § 102(e).

The Asserted Claims require a multithreaded processor with threads that process at least two sample types. The evidence shows Selzer describes a collection of single threaded, general purpose DSPs arranged in parallel with each individual processor functioning as a geometry

PUBLIC VERSION

module or a rendering module. (CX-2128C (Dally) at Q&A 72, 93, 102, 106.) In Selzer, each single-threaded processor is capable of operating on one sample type. (*Id.* at Q&A 93, 102, 127.) Nothing in Selzer discloses the multithreaded execution of at least two sample types. (*Id.*)

The evidence shows Van Hook's disclosure is limited to only a portion of a multithreaded graphics processor. (*Id.* at Q&A 63-65, 86, 100, 125.) That portion is an execution unit that can process only a single sample type, either vertices or pixels, but not both. (*Id.*) Because the execution unit disclosed in Van Hook can only process one sample type, Van Hook does not disclose or suggest a multithreaded processor that can assign or allocate threads to process different sample types. (*Id.* at Q&A 65, 100, 125.) Thus, even as a combination, Selzer and Van Hook fail to disclose or suggest a multithreaded graphics processor with threads that can process at least two sample types as claimed in the '685 Patent. (*Id.* at Q&A 92, 95, 99, 100, 102, 103, 123-25, 127, 128.) Thus, I find Respondents have failed to show by clear and convincing evidence that Selzer in light of Van Hook discloses "[a] graphics processor for simultaneous multithreaded execution of program instructions associated with threads to process at least two sample types."

The Asserted Claims also require the assignment/allocation of the at least two sample types to threads for execution based on a priority among the at least two sample types. Van Hook makes no mention of assigning sample types to threads, disclosing only that instructions for different programs are interleaved based on instruction dependencies. (CX-2128C (Dally) at Q&A 105, 133.) And, because Van Hook fails to disclose the assignment of samples to threads, it cannot disclose or suggest a priority among at least two sample types with respect to thread assignment. (*Id.*) In fact the word "priority" does not occur in the specification of Van Hook.

PUBLIC VERSION

(RX-350.) Van Hook also does not mention thread scheduling. (CX-2128C (Dally) at Q&A 105.)

Selzer describes multiple single-threaded DSPs each of which used as either a geometry module or as a rendering module. (*Id.* at Q&A 70.) The evidence shows that in Selzer the rendering module DSP is coupled to a portion of the frame buffer and primarily operates to process the pixels associated with that portion of the frame buffer. (*Id.* at Q&A 106, 135.) The evidence shows that when there is no rendering work available, a rendering module may switch its function and become a geometry module and process vertices. (*Id.* at Q&A 71, 106, 135.) However, the evidence is clear that this use of idle rendering modules to process vertices is not at all equivalent to using a priority to assign sample types to the multiple hardware thread contexts of a multi-threaded processor. (*Id.* at Q&A 106, 135.)

Thus, for at least the reasons above, even as a combination, Selzer and Van Hook fail to disclose or suggest assigning threads based on an allocation priority among sample types as claimed in the '685 Patent. (*Id.* at Q&A 104-106, 132, 135.) Accordingly, I find Respondents have failed to show by clear and convincing evidence that Selzer in light of Van Hook discloses "at least two sample types are assigned to the threads for multithreaded execution based on an allocation priority among the at least two sample types."

In light of my findings above that Selzer and Van Hook do not disclose, either alone or in combination, each of the elements of claims 1 and 15 of the '685 patent, I find Respondents have failed to prove by clear and convincing evidence that claim 1 or claim 15 of the '685 patent is obvious in view of Selzer in light of Van Hook.

PUBLIC VERSION

b. Amanatides in view of Van Hook

Discussion

Amanatides et al., was published in 1993, more than a year before the effective filing date of the '685 patent, and is thus prior art under 35 U.S.C. § 102(b). (RX-354.) Van Hook issued on December 10, 2010 on an application filed on July 26, 2000. (RX-350) Because the filing date of Van Hook is before the effective filing date of the '685 patent it is prior art under 35 U.S.C. § 102(e).

The Asserted Claims require a multithreaded processor with threads that process at least two sample types. The evidence shows Amanatides describes a graphics system that uses general purpose single-threaded microprocessors to process both vertices and pixels. (CX-2128C (Dally) at 77, 157, 176.) Amanatides uses Intel i860 as the general purpose single-threaded microprocessor. (*Id.*) The evidence shows each processor is fed by a single FIFO that holds both vertex tasks and pixel tasks. (*Id.*) Vertex tasks are represented in the form of geometry messages (“GMs”) and pixel tasks are represented in the form of render messages (“RMs”). (*Id.*) Nothing in Amanatides discloses the multithreaded execution of at least two sample types. (*Id.*)

The evidence shows Van Hook’s disclosure is limited to only a portion of a multithreaded graphics processor. (*Id.* at Q&A 63-65, 86, 155, 174.) That portion is an execution unit that can process only a single sample type, either vertices or pixels, but not both. (*Id.*) Because the execution unit disclosed in Van Hook can only process one sample type, Van Hook does not disclose or suggest a multithreaded processor that can assign or allocate threads to process different sample types. (*Id.* at Q&A 65, 155, 174.)

Thus, even as a combination, Amanatides and Van Hook fail to disclose or suggest a multithreaded graphics processor with threads that can process at least two sample types as

PUBLIC VERSION

claimed in the '685 Patent. (*Id.* at Q&A 149, 152, 154, 155, 173, 174.) Accordingly, I find Respondents have failed to show by clear and convincing evidence that Amanatides in light of Van Hook discloses “[a] graphics processor for simultaneous multithreaded execution of program instructions associated with threads to process at least two sample types.”

The Asserted Claims also require the assignment/allocation of the at least two sample types to threads for execution based on a priority among the at least two sample types. Van Hook makes no mention of assigning sample types to threads, disclosing only that instructions for different programs are interleaved based on instruction dependencies. (CX-2128C (Dally) at Q&A 160, 182.) And because Van Hook fails to disclose the assignment of samples to threads, it cannot disclose or suggest a priority among at least two sample types with respect to thread assignment. (*Id.*) In fact the word “priority” does not occur in the specification of Van Hook. (RX-350.) Van Hook also does not mention thread scheduling. (CX-2128C (Dally) at Q&A 160.)

Amanatides only describes a multiprocessor with multiple single-threaded processors, which gives priority over a shared bus to render messages (“RMs”), or pixel tasks. (CX-2128C (Dally) at Q&A 77, 161, 177.) This priority is implemented to avoid deadlock of the processors and not to assign a sample to a thread – whether based on a sample type priority or any other criteria. (*Id.*) The evidence shows that priority to enter a bus is not at all equivalent to using a priority to assign sample types to the multiple hardware thread contexts of a multi-threaded processor. (*Id.*) Moreover, the RMs are broadcast to all of the processors at once, not allocated or assigned to a particular processor or thread based on a priority, as required by the asserted claims of the '685 patent. (*Id.*)

PUBLIC VERSION

Thus, for at least the reasons above, even as a combination, Amanatides and Van Hook fail to disclose or suggest assigning threads based on an allocation priority among sample types as claimed in the '685 Patent. (*Id.* at Q&A 149, 159-61,181-184.) Accordingly, I find Respondents have failed to show by clear and convincing evidence that Amanatides in light of Van Hook discloses “at least two sample types are assigned to the threads for multithreaded execution based on an allocation priority among the at least two sample types.”

In light of my findings above that Amanatides and Van Hook do not disclose, either alone or in combination, each of the elements of claims 1 and 15 of the '685 patent, I find Respondents have failed to prove by clear and convincing evidence that claim 1 or claim 15 of the '685 patent is obvious in view of Amanatides in light of Van Hook.

VII. DOMESTIC INDUSTRY

Section 337(a)(2) provides that “Subparagraph(B) ... of paragraph (1) apply only if an industry in the United States, relating to the articles protected by the patent, copyright, trademark, mask work, or design concerned, exists or is in the process of being established.” Here, the Commission has determined that the economic prong of the domestic industry has been established, and NVIDIA seeks to satisfy the technical prong of the domestic industry requirement with respect to claim 14 of the '140 patent, claims 21-23 of the '372 patent, and claim 15 of the '685 patent. However, I have found herein that some of these claims are invalid. I note this because even if the technical and economic prongs are both met, a violation of Section 337 cannot be found based on the Complainant’s practice of invalid claims.

To establish a violation of Section 337 based on patent infringement, a Complainant must show that a domestic industry exists (or is in the process of being established) with respect to “articles protected by the patent.” 19 U.S.C. § 1337(a)(2) (emphasis added). A violation cannot be found where Complainant has not shown DI articles practicing a valid and enforceable claim.

PUBLIC VERSION

Certain Audiovisual Components and Products Containing the Same, Inv. No. 337-TA-837, Comm'n Op. at 33 (Mar. 10, 2014) ("Because invalid claims cannot protect articles, [Complainant] has not proven that a valid patent claim protects [its DI] products. For this additional reason, [Complainant] has not proven a violation of section 337.") (citations omitted);

Certain Integrated Circuit Chips and Products Containing the Same, Inv. No. 337-TA-859, Comm'n Op. at 52, n.27 (Aug. 22, 2014) ("The Commission notes that the correct finding when the DI products practice only invalid claims is that there is no violation of section 337, not that there is no domestic industry").

Here, NVIDIA has only met the technical prong of the domestic industry requirement with respect to invalid claims of both the '372 patent and the '140 patent. Thus, there can be no violation of Section 337 as to those patents.

VIII. CONCLUSIONS OF LAW

1. The Commission has personal jurisdiction over the parties and subject-matter jurisdiction over the Accused Products.
2. The importation or sale requirement of Section 337 is satisfied.
3. The Accused Products with Adreno, PowerVR, or Mali GPUs do not infringe claim 14 of U.S. Patent No. 7,209,140.
4. The NVIDIA DI Products do not practice claim 14 of the '140 patent.
5. Claim 14 of the '174 patent is invalid as obvious pursuant to 35 U.S.C. § 103.
6. The domestic industry requirement is not met with regard to the '140 patent.
7. There has been no violation of Section 337 with regard to the '140 patent.
8. The Accused Products (except those listed in RDX-408C) with Adreno, PowerVR, or Mali GPUs infringe claim 23 of U.S. Patent No. 6,690,372.
9. The NVIDIA DI Products practice at least claims 21 and 22 of the '372 patent.
10. Claims 23 and 24 of the '372 patent are invalid as anticipated under 35 U.S.C. § 102.
11. The domestic industry requirement is not met with regard to the '372 patent.
12. There has been no violation of Section 337 with regard to the '372 patent.
13. The Accused Products with Adreno or Mali GPUs do not infringe claims 1 or 15 of U.S. Patent No. 7,038,685.
14. The NVIDIA DI Products practice claim 15 of the '685 patent.
15. Claims 1 or 15 of the '685 patent have not been shown to be invalid under 35 U.S.C. §§ 102, 103.
16. The domestic industry requirement is met with regard to the '685 patent.
17. There has been no violation of Section 337 with regard to the '685 patent.

PUBLIC VERSION

IX. INITIAL DETERMINATION AND ORDER

Based on the foregoing,^{13 14} it is my Initial Determination that there is no violation of Section 337 of the Tariff Act of 1930, as amended, in the importation into the United States, the sale for importation, or the sale within the United States after importation of Certain Consumer Electronics And Display Devices With Graphics Processing And Graphics Processing Units Therein, in connection with the Asserted Claims of U.S. Patent No. 7,209,140; U.S. Patent No. 6,690,372; or U.S. Patent No. 7,038,685. Furthermore, it is the determination of this Administrative Law Judge that a domestic industry in the United States does not exist that practices or exploits U.S. Patent No. 7,209,140 or U.S. Patent No. 6,690,372, but that a domestic industry in the United States does exist that practices or exploits U.S. Patent No. 7,038,685.

The undersigned hereby CERTIFIES to the Commission this Initial Determination AND Recommended Determination on Remedy and Bond, 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 as listed in the appendices hereto.¹⁵

Pursuant to 19 C.F.R. § 210.42(h), this Initial Determination shall become the determination of the Commission unless a party files a petition for review pursuant to 19 C.F.R.

¹³ 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.

¹⁴ To the extent any party cited to their pre-hearing brief in support of an argument they made in any of their post-hearing briefs, those citations to the pre-hearing brief are hereby stricken as a violation of my Ground Rules. (See Order No. 2, Ground Rule 15.1.1.)

¹⁵ 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.

PUBLIC VERSION

§ 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.

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 and Recommended Determination, the parties shall jointly submit: (1) a proposed public version of these opinions 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.¹⁶

SO ORDERED.


Thomas B. Pender
Administrative Law Judge

¹⁶ 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.

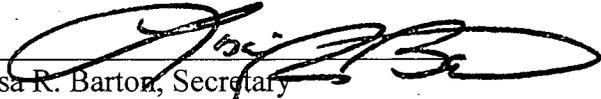
See 19 C.F.R. § 201.6(a). Thus, to constitute confidential business information the disclosure of the information sought to be designated confidential must *likely have the effect of* either: (1) impairing the Commission's ability to obtain such information as is necessary to perform its statutory functions; or (2) *causing substantial harm* to the competitive position of the person, firm, partnership, corporation, or other organization from which the information was obtained.

**IN THE MATTER OF CERTAIN CONSUMER
ELECTRONICS AND DISPLAY DEVICES WITH GRAPHICS
PROCESSING AND GRAPHICS PROCESSING UNIT THEREIN**

337-TA-932

CERTIFICATE OF SERVICE

I, Lisa R. Barton, hereby certify that the attached **PUBLIC VERSION INITIAL DETERMINATION** has been served upon the **Commission Investigative Attorney, Whitney Winston, Esq.**, and the following parties as indicated on NOV 10 2015



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

FOR COMPLAINANTS NVIDIA CORPORATION:

I. Neel Chatterjee, Esq.
ORRICK, HERRINGTON & SUTCLIFFE LLP
1000 Marsh Road
Menlo Park, California 94025

() Via Hand Delivery
 Via Express Delivery
() Via First Class Mail
() Other: _____

FOR RESPONDENT QUALCOMM, INC:

Deanna Tanner Okun, Esq.
ADDUCI, MASTRIANI & SCHAUMBERG, LLP
1133 Connecticut Avenue, N.W., 12th Floor
Washington, DC 20036

() Via Hand Delivery
 Via Express Delivery
() Via First Class Mail
() Other: _____

**FOR RESPONDENTS SAMSUNG ELECTRONICS CO., LTD.,
SAMSUNG ELECTRONICS AMERICA, INC., SAMSUNG
TELECOMMUNICATIONS AMERICA, LLC, AND SAMSUNG
SEMICONDUCTOR, INC. (Collectively, "Samsung"):**

D. Sean Trainer, Esq.
KIRKLAND & ELLIS LLP
655 Fifteenth Street, N.W.
Washington, DC 20005

() Via Hand Delivery
 Via Express Delivery
() Via First Class Mail
() Other: _____