

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF VIRGINIA
RICHMOND DIVISION**

SAMSUNG ELECTRONICS CO., LTD., and
SAMSUNG ELECTRONICS AMERICA,
INC.,

Plaintiffs,

v.

NVIDIA CORPORATION, VELOCITY
MICRO, INC. D/B/A VELOCITY MICRO,
AND VELOCITY HOLDINGS, LLC

Defendants.

CIVIL ACTION NO. 3:14-CV-00757-REP

JURY TRIAL DEMANDED

SECOND AMENDED COMPLAINT

Plaintiffs Samsung Electronics Company, Ltd. (“SEC”) and Samsung Electronics America, Inc. (“SEA”) (collectively “Samsung”), by and through their undersigned attorneys, hereby file this Second Amended Complaint¹ against NVIDIA Corporation (“NVIDIA”), Old Micro, Inc. F/K/A Velocity Micro, Inc. and Velocity Holdings, LLC (collectively “Velocity”). All defendants will be referred to collectively as “Defendants.” Samsung states as follows:

THE PARTIES

1. Plaintiff SEC is a multi-national corporation organized under the laws of Korea with its principal place of business located at 416 Maetan-3dong, Yeongtong-gu, Suwon-City, Gyeonggi-do, Korea 443-742.

¹ The Second Amended Complaint amends the First Amended Complaint by changing the name of the defendant originally identified as “Velocity Micro, Inc. D/B/A Velocity Micro” to “Old Micro, Inc. F/K/A Velocity Micro, Inc.” The only portions of the Second Amended Complaint that differ from the First Amended Complaint are the introductory paragraph, paragraph 4, paragraph 5, and footnote 2. Exhibit I, attached hereto, shows Samsung’s changes to those portions of the Second Amended Complaint. The remainder of the Second Amended Complaint is unchanged from the First Amended Complaint.

2. Plaintiff SEA is a corporation organized and existing under the laws of the state of New York with its principal place of business in Ridgefield Park, New Jersey and is a wholly-owned subsidiary of SEC.

3. Defendant NVIDIA is a Delaware corporation with its headquarters located at 2701 San Tomas Expressway, Santa Clara, California 95050. NVIDIA imports into the United States, offers for sale, sells and/or uses in the United States graphics processing units (GPUs), system on a chip (SOC) units, graphics cards, and mobile computing devices such as tablet computers.

4. Defendant Old Micro, Inc., which was formerly known as Velocity Micro, Inc., is a corporation organized and existing under the laws of Virginia. The State Corporation Commission of the Commonwealth of Virginia lists the location of Old Micro, Inc.'s registered agent as 9030 Stony Point Parkway, Suite 400, Richmond, Virginia, 23235. At times relevant to this action, Old Micro, Inc. manufactured computers that it offered for sale, sold and/or used in the United States.

5. Defendant Velocity Holdings, LLC is a limited liability corporation organized and existing under the laws of Virginia, which does business as Velocity Micro. Velocity Holdings, LLC has its principal place of business located at 835 Grove Rd., Midlothian, Virginia 23114. The members of Velocity Holdings, LLC reside in the Eastern District of Virginia. Velocity Holdings, LLC is a manufacturer of computers that it offers for sale, sells and/or uses in the United States.

SAMSUNG

6. From its inception as a small business in Taegu, Korea, Samsung has grown to become one of the world's leading electronics companies, specializing in digital products,

semiconductors, memory, and system integration. Today, Samsung's innovative consumer products are widely recognized and appreciated across the globe. Samsung designs, develops, manufactures, and sells leading consumer electronics, including mobile phones, smartphones, tablet computers, and laptop computers.

7. Samsung has a long history of groundbreaking innovation across a wide range of technologies. During the 1970s and 1980s, Samsung's core technology businesses diversified and expanded globally. For example, Samsung began production of personal computers in 1983 and selected telecommunications and semiconductors as core business lines in 1988. **About Samsung.**² During this period, Samsung additionally challenged itself to restructure old businesses and enter new ones with the aim of becoming one of the world's top five electronics companies.

8. Samsung's commitment to innovation is demonstrated in part by the billions of dollars in research and development expenditures incurred over the years. From 2005 through 2010 alone, Samsung invested more than \$35 billion in research and development. More than a quarter of all Samsung employees—over 55,100 engineers overall—are engaged in cutting-edge research and development projects.

9. Samsung's commitment to innovation and investment in research and development is demonstrated by the fact that SEC has in its portfolio over 40,000 United States utility patents and over 4,000 design patents. Samsung is consistently ranked ahead of other technology companies in terms of the number of issued patents obtained in the United States, with over 4,000 U.S. utility patents issued each year in recent years.

² Bolded terms or references are defined in the text of this Second Amended Complaint or in the section entitled "References."

10. Samsung is the assignee and owner of the patents at issue in this action, which relate to semiconductor technologies and computing devices: U.S. Patent No. 5,860,158 (the “**158 Patent**”), U.S. Patent No. 6,262,938 (the “**938 Patent**”), U.S. Patent No. 6,287,902 (the “**902 Patent**”), U.S. Patent No. 6,819,602 (the “**602 Patent**”), U.S. Patent No. 8,252,675 (the “**675 Patent**”), U.S. Patent No. 6,804,724 (the “**724 Patent**”), U.S. Patent No. 7,073,054 (the “**054 Patent**”), and U.S. Patent No. 5,777,854 (the “**854 Patent**”) (collectively the “**Asserted Patents**”).

11. Samsung has expended significant resources and efforts to become a pioneer in the area of mobile devices including smartphones, tablets, and computers. Samsung’s research and development successes have propelled the company to its status as one of the world’s leading electronics companies.

12. As the Android operating system has grown, Samsung has grown with it to become the leading supplier of Android-based devices, including both Android-based smartphones and Android-based tablet computers.

13. Samsung has also continuously innovated in the design and manufacturing of semiconductor products. Samsung established its first semiconductor fabrication facility in 1984 and began solidifying its position as a leading company in semiconductor technology. Samsung’s semiconductor research and development has led numerous industry innovations. Just one year after its initial semiconductor facility was established, Samsung established the industry’s first 6 inch fabrication facility in 1985. In 1992, Samsung developed the world’s first 64Mb DRAM. In 1998, Samsung began shipping the world’s first 128 Mb flash memory, and in 2006 Samsung launched the world’s first 32GB solid state disk drive.

14. Because of Samsung's commitment to semiconductor research and development, Samsung is now recognized as one of the world's leading semiconductor manufacturers. Samsung's semiconductor foundry labs include state-of-the-art facilities that support volume production of chips using 90 nm, 65 nm, 45/40 nm, and 32/28 nm processes.

15. Samsung's deep experience and comprehensive knowledge of high-volume manufacturing technology has led to numerous innovations in the design and manufacturing of graphics technologies and computing devices.

16. Without the ability to enforce its intellectual property rights, such as those relating to the semiconductor technologies and computing devices at issue in this action, Samsung would not be able to sustain the extensive commitment to research and development that has enabled it to lead the way into numerous improvements across a broad range of technologies.

NVIDIA

17. NVIDIA designs, develops, manufactures, imports, and sells GPUs, SOCs, and graphics cards for a wide range of products including mobile devices and desktop systems. NVIDIA creates and uses reference graphics card designs for incorporating NVIDIA GPUs and SOCs into graphics cards, in addition to designing, manufacturing, using and selling graphics cards.

18. NVIDIA additionally manufactures, or contracts others to manufacture, mobile processors for smartphones and tablets. NVIDIA sells tablet computers under the name "SHIELD Tablet." NVIDIA's tablet computers use the Android operating system and compete directly with Samsung tablet computers that also incorporate the Android operating system.

19. NVIDIA sells and offers to sell products and services throughout the United States, including in this Judicial District, through its website and through major electronics

retailers in North America. NVIDIA introduces products and services that infringe the **Asserted Patents** knowing that they would be sold in this Judicial District and elsewhere in the United States.

20. NVIDIA conducts a significant amount of business in Virginia through online sales and advertisements directly to consumers and through product sales by NVIDIA's distributors and resellers. NVIDIA targets Virginia residents through training programs offered in this Judicial District, such as "CUDA Defense Workshops" offered for consumers of its GPU platforms and products. *See Cuda Defense Workshops*. CUDA Defense Workshops are multi-day training programs "designed for programmers working in the defense industry, who are looking to develop skills in application design and optimization to fully leverage the parallel computing capabilities of compute GPUs using the CUDA platform." *Id.* For example, NVIDIA sponsored and delivered a free two-day CUDA Defense Workshop on its GPU products on May 23-24, 2013 at 13600 Eds Dr., Herndon, VA 20171. *Id.*

21. NVIDIA is registered as a foreign corporation with the Commonwealth of Virginia and may be served with process through its registered agent in the Commonwealth of Virginia, Corporation Service Company, 1111 E. Main St., Richmond, VA 23219.

VELOCITY

22. Velocity produces custom high-performance desktop and mobile computers (including tablets) that are designed for applications such as gaming, digital graphic design, home theater use, and common home and office use. Velocity also provides personal computers designed for specialized applications, including scientific workstations.

23. As of October 29, 2014, Velocity advertised itself as "the premier high-performance innovator of consumer technology and electronics in North America." **About**

Velocity. Velocity additionally advertised that it is “100% based in the USA in Richmond, VA.” *Id.* Velocity also advertised that its laptop computers are “[a]ssembled by our expert engineers in Richmond, VA” **Velocity Gaming and Enthusiast Laptops.** Velocity also advertises that its products include “[c]areful custom integration and testing, final assembly by hand in Richmond, Virginia, USA.” **Configure Your M17.**

24. Velocity prominently advertises that several of its product lines use NVIDIA GPUs. For example, on February 18, 2014, Velocity announced “the immediate availability of the NVIDIA GeForce GTX 750, GTX 750 Ti, and the amazing GTX TITAN Black in select Gaming/Enthusiast and Workstation desktops.” **Velocity NVIDIA GeForce Press Release.** “‘Every evolution of the GeForce cards presents better and better choices for consumers, and these new products are no exception,’ said Randy Copeland, President and CEO of Velocity. ‘With the best-in-class performance of the TITAN Black and incredible price for performance of the GTX 750 and 750 Ti, NVIDIA has brought two great choices to market. We’re excited to offer them to our enthusiast customers.’” *Id.*

25. In another example, on May 23, 2013, Velocity issued a press release announcing “desktops powered by NVIDIA GeForce GTX 780.” **Velocity Desktop Press Release.** “‘Since our first Editors’ Choice award in 2002, we’ve seen PC hardware come a long way. NVIDIA has always led the way in that innovation,’ said Randy Copeland, President and CEO of Velocity. ‘After spending a few days testing the GeForce GTX 780, all I can say is WOW – they’ve done it again. This is seriously fast hardware and a must have for enthusiasts.’” *Id.*

26. In yet another example, on May 25, 2010, Velocity issued a press release announcing the release of its “Raptor X17 Notebook Featuring NVIDIA® GeForce® GTX 480M Mobile Graphics.” **Raptor X17 Press Release.** “‘With the release of the NVIDIA

GeForce GTX 480M, we're able to provide a new level of mobile gaming performance to our customers who demand the best hardware,' said Chip Lowell, VP of Sales for Velocity. 'This is more than just a desktop replacement. It's the ultimate mobile gaming system.'" *Id.*

27. Velocity has a history of issuing press releases prominently advertising Velocity as the global launch partner of select NVIDIA GPUs. For example, on June 5, 2006, Velocity issued a press release announcing it as a global launch partner for NVIDIA's GeForce 7950 GX2 Graphics Solution. **Velocity NVIDIA GeForce Press Release.** "'We are proud to be a global launch partner with NVIDIA on this exciting new graphics solution. NVIDIA GeForce 7950 GX2 takes gaming performance to extreme levels,'" said Randy Copeland, Velocity President and founder. Our strong relationship with NVIDIA enables us to continue to offer our customers the level of innovation and performance that they've grown to expect from us.'" *Id.* Ujesh Desia, general manager of desktop GPUs at NVIDIA, stated that "[w]ith that kind of graphics horsepower at their disposal, Velocity has the knowledge necessary to configure a PC for customers of BestBuy.com that will deliver an immersive, high-performance gaming experience.'" *Id.*

28. In another example, Velocity issued a press release announcing that "Velocity Micro Partners With NVIDIA To Launch GeForce 6 Series of Graphics." **Velocity NVIDIA Partner Press Release.** "'We're pleased to be able to launch such compelling technology in conjunction with a well-respected company like NVIDIA,' said Velocity founder and president Randall Copeland. 'This partnership is just one of the many reasons why Velocity Micro is able to offer premium products to meet the needs of all its customers.'" *Id.*

29. Velocity sells desktop and laptop computers that incorporate NVIDIA GPUs. Velocity competes with Samsung in the sale of computers.

30. In 2010, Velocity introduced a tablet computer under the name “Cruz Tablet” that uses the Android operating system and competes with Samsung tablet computers that incorporate the Android operating system.

31. Velocity sells and offers to sell products and services throughout the United States, including in this Judicial District, through its website and through major electronics retailers in North America. Velocity advertises that it “competes on a national level with a full product line that includes PCs, notebooks, and peripherals. Velocity products have been sold at nearly every major electronics retailer in North America including Best Buy, Circuit City, RadioShack, Costco, Sears, Target, and many others.” **Velocity Company History.** Velocity introduces products and services that infringe the **Asserted Patents** knowing that they would be sold in this Judicial District and elsewhere in the United States.

JURISDICTION AND VENUE

32. This action arises under the patent laws of the United States, Title 35 of the United States Code. This Court has subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

33. Pursuant to 28 U.S.C. § 1367(a), this Court also has supplemental jurisdiction over the state law claim of false advertising under Va. Code §§ 18.2-216 and 59.1-68.3. This state law claim is so related to the claims under which the Court has original subject matter jurisdiction that it forms part of the same case and controversy under Article III of the United States Constitution.

34. Pursuant to 28 U.S.C. § 1332, this Court also has subject matter jurisdiction over the state law claim asserted in this action because there is complete diversity between the parties

to the state law claim and because the amount in controversy against defendant NVIDIA separately exceeds the value of \$75,000, exclusive of interests and costs.

35. This Court has personal jurisdiction over NVIDIA and Velocity by virtue of the business activities Defendants conduct within the Commonwealth of Virginia, resulting in sufficient minimum contacts with this forum. NVIDIA and Velocity conduct substantial business activities in the state, including acts of patent infringement that have injured Samsung, and therefore Defendants have purposefully availed themselves of the laws of the Commonwealth of Virginia.

36. NVIDIA and Velocity directly and/or through intermediaries make, offer for sale, sell, and/or advertise (including through a website) products and services in the Commonwealth of Virginia and in this Judicial District. Defendants NVIDIA and Velocity purposefully and voluntarily place their infringing products into the stream of commerce with the expectation that they will be purchased by consumers in the Commonwealth of Virginia and in this Judicial District. These infringing products have been and continue to be purchased by consumers in the Commonwealth of Virginia and in this Judicial District.

37. NVIDIA has been registered as a foreign corporation with the Commonwealth of Virginia since at least May 12, 2010. NVIDIA may be served with process through its registered agent in the Commonwealth of Virginia, Corporation Service Company, 1111 E. Main St., Richmond, VA 23219.

38. Velocity advertised that it is “100% based in the USA in Richmond, VA.” **About Velocity**. Velocity additionally advertised that its laptop computers are “[a]ssembled by our expert engineers in Richmond, VA,” **Gaming and Enthusiast Laptops**, and that its products

include “[c]areful custom integration and testing, final assembly by hand in Richmond, Virginia, USA.” **Configure Your M17.**

39. Velocity has been registered as a corporation with the Commonwealth of Virginia since at least January 7, 1998. Velocity may be served with process through their registered agents in the Commonwealth of Virginia, Freed & Shepherd, P.C., 9030 Stony Point Pkwy Suite 400, Richmond, VA 2325 and Gary S Cook, 4551 Cox Rd Suite 210, Glen Allen, VA 23060.

40. Venue is proper in this Judicial District pursuant to 28 U.S.C. §§ 1391(b)-(c) and 1400(b) because Defendants have conducted business in this Judicial District, have offices and facilities in this District, have committed specific acts of infringement and/or induced or contributed to acts of infringement in this District, and continue to commit and/or induce or contribute to acts of infringement in this District, entitling Samsung to relief.

ACCUSED PRODUCTS

41. The “**Accused Products**” means all **Accused GPUs** (graphics processing units) and **Accused SOCs** (systems-on-chip) and all products that contain an **Accused GPU** or **Accused SOC** that have been made, sold, or offered for sale, or imported into the United States at any time since November 4, 2008, including all products listed below.

42. The “**Accused GPUs**” means all NVIDIA GPUs and all products that contain a NVIDIA GPU that have been made, used, sold, offered for sale, or imported into the United States since November 4, 2008, including the products identified in paragraphs 43 to 55, below.

43. The “**938 Accused GPUs**” means all NVIDIA GPUs and all products that contain NVIDIA GPUs that are designed to operate in conjunction with JEDEC standard SDRAM or SGRAM that includes posted CAS latency functionality or any equivalent thereof, and includes, without limitation, the GeForce 8100 mGPU, GeForce 8300 mGPU, GeForce 8300

GS, GeForce 8400 GS, GeForce 8500 GT, GeForce 8600 GT, GeForce 9300 mGPU, GeForce 9300 GS, GeForce 9400 GT, GeForce 9500 GT, GeForce 9600 GS, GeForce 9600 GS, GeForce G100, GeForce GT120, GeForce GT130, GeForce 205, GeForce 210, GeForce GT 220, GeForce GT 230, GeForce GT 240, GeForce 310, GeForce 315, GeForce GT 330, GeForce 405, GeForce GT 510, GeForce GT 520, GeForce GT 530, GeForce GT 545 DDR3, GeForce 605, GeForce GT 610, GeForce GT 620, GeForce GT 625, GeForce GT 630, GeForce GT 630 (DDR3), GeForce GT 635, GeForce GT 640 (DDR3), GeForce GT 705, GeForce GT 710, GeForce GT 720, GeForce GT 730 (64b DDR3), GeForce GT 730 (128b DDR3), GeForce GT 740 (DDR3), GeForce GTX 745, GeForce 8200M G, GeForce 8400M G, GeForce 8400M GS, GeForce 8400M GT, GeForce 8600M GS, GeForce 8600M GT, GeForce 9100M G mGPU, GeForce 9200M GS, GeForce 9300M G, GeForce 9300M GS, GeForce 9400M G, GeForce 9500M G, GeForce 9500M GS, GeForce 9600M GS, GeForce 9600M GT, GeForce G 102M, GeForce G 103M, GeForce G 105M, GeForce G 110M, GeForce GT 120M, GeForce GT 130M, GeForce GT 220M, GeForce 305M, GeForce 310M, GeForce 315M, GeForce 320M, GeForce GT 320M, GeForce GT G25M, GeForce GT 330M, GeForce GT 335M, GeForce GT 350M, GeForce GTS 350M, GeForce GTS 360M, GeForce 510M, GeForce GT 415M, GeForce GT 420M, GeForce GT 425M, GeForce GT 445M, GeForce GT 520M, GeForce GT 520MX, GeForce GT 525M, GeForce GT 540M, GeForce GT 550M, GeForce GT 555M, GeForce 610M, GeForce GT 620M, GeForce GT 625M, GeForce GT 630M, GeForce GT 635M, GeForce 640M LE, GeForce GT 640M, GeForce GT 645M, GeForce GT 650M, GeForce 710M, GeForce GT 720M, GeForce GT 730M, GeForce GT 735M, GeForce GT 740M, GeForce GT 745M, GeForce GT 750M, GeForce 820M, GeForce 825M, GeForce 830M, GeForce 840M, GeForce GTX 850M, Quadro FX 1000, Quadro FX 1100, Quadro FX 2000, Quadro FX 350, Quadro FX 370, Quadro FX 370 LP,

Quadro FX 740, Quadro FX 50, Quadro FX 1700, Quadro 400, Quadro 600, Quadro 410, Quadro K600, Quadro K420, Quadro K620, NVS 300, NVS 310, NVS 5400M, Quadro FX 360M, Quadro 500M, Quadro 1000M, Quadro 2000M, Quadro K5000M, Quadro K1000M, Quadro K2000M, Quadro NVS 120M, Quadro NVS 130M, and Quadro NVS 300M.

44. Each of the products in the preceding paragraph is or contains an NVIDIA GPU designed to operate in conjunction with JEDEC standard SDRAM or SGRAM that includes posted CAS latency functionality or any equivalent thereof.

45. The “**602 Accused GPUs**” means all NVIDIA GPUs and all products that contain NVIDIA GPUs that are designed to support any SDRAM or SGRAM that includes an input data strobe feature supporting single-ended and differential signaling or any equivalent thereof, and includes, without limitation, the GeForce 8100 mGPU, GeForce 8300 mGPU, GeForce 8300 GS, GeForce 8400 GS, GeForce 8500 GT, GeForce 8600 GT, GeForce 9300 mGPU, GeForce 9300 GS, GeForce 9400 GT, GeForce 9500 GT, GeForce 9600 GS, GeForce 9600 GS, GeForce G100, GeForce GT120, GeForce GT130, GeForce 205, GeForce 210, GeForce GT 220, GeForce GT 230, GeForce GT 240, GeForce 310, GeForce 315, GeForce GT 330, GeForce 405, GeForce GT 510, GeForce GT 520, GeForce GT 530, GeForce GT 545 DDR3, GeForce 605, GeForce GT 610, GeForce GT 620, GeForce GT 625, GeForce GT 630, GeForce GT 630 (DDR3), GeForce GT 635, GeForce GT 640 (DDR3), GeForce GT 705, GeForce GT 710, GeForce GT 720, GeForce GT 730 (64b DDR3), GeForce GT 730 (128b DDR3), GeForce GT 740 (DDR3), GeForce GTX 745, GeForce 8200M G, GeForce 8400M G, GeForce 8400M GS, GeForce 8400M GT, GeForce 8600M GS, GeForce 8600M GT, GeForce 9100M G mGPU, GeForce 9200M GS, GeForce 9300M G, GeForce 9300M GS, GeForce 9400M G, GeForce 9500M G, GeForce 9500M GS, GeForce 9600M GS, GeForce 9600M GT,

GeForce G 102M, GeForce G 103M, GeForce G 105M, GeForce G 110M, GeForce GT 120M, GeForce GT 130M, GeForce GT 220M, GeForce 305M, GeForce 310M, GeForce 315M, GeForce 320M, GeForce GT 320M, GeForce GT G25M, GeForce GT 330M, GeForce GT 335M, GeForce GT 350M, GeForce GTS 350M, GeForce GTS 360M, GeForce 510M, GeForce GT 415M, GeForce GT 420M, GeForce GT 425M, GeForce GT 445M, GeForce GT 520M, GeForce GT 520MX, GeForce GT 525M, GeForce GT 540M, GeForce GT 550M, GeForce GT 555M, GeForce 610M, GeForce GT 620M, GeForce GT 625M, GeForce GT 630M, GeForce GT 635M, GeForce 640M LE, GeForce GT 640M, GeForce GT 645M, GeForce GT 650M, GeForce 710M, GeForce GT 720M, GeForce GT 730M, GeForce GT 735M, GeForce GT 740M, GeForce GT 745M, GeForce GT 750M, GeForce 820M, GeForce 825M, GeForce 830M, GeForce 840M, GeForce GTX 850M, Quadro FX 1000, Quadro FX 1100, Quadro FX 2000, Quadro FX 350, Quadro FX 370, Quadro FX 370 LP, Quadro FX 740, Quadro FX 50, Quadro FX 1700, Quadro 400, Quadro 600, Quadro 410, Quadro K600, Quadro K420, Quadro K620, NVS 300, NVS 310, NVS 5400M, Quadro FX 360M, Quadro 500M, Quadro 1000M, Quadro 2000M, Quadro K5000M, Quadro K1000M, Quadro K2000M, Quadro NVS 120M, Quadro NVS 130M, Quadro NVS 300M, and Jetson TK1.

46. Each of the products in the preceding paragraph is or contains an NVIDIA GPU designed to support any SDRAM or SGRAM that includes an input data strobe feature supporting single-ended and differential signaling or any equivalent thereof.

47. The “**Accused 28 nm GPUs**” means all NVIDIA GPUs utilizing 28 nanometer fabrication processing and all products that contain such a GPU, and includes, without limitation, the GeForce 710M, GeForce 820M, GeForce 825M, GeForce 830M, GeForce 840M, GeForce GT 620M, GeForce GT 625M, GeForce GT 630, GeForce GT 630M, GeForce GT 635, GeForce

GT 640, GeForce GT 640M, GeForce GT 640M LE, GeForce GT 645M, GeForce GT 650M, GeForce GT 710, GeForce GT 720, GeForce GT 720M, GeForce GT 730, GeForce GT 730M, GeForce GT 735M, GeForce GT 740, GeForce GT 740M, GeForce GT 745M, GeForce GT 750M, GeForce GT 755M, GeForce GTX 645, GeForce GTX 650, GeForce GTX 650 Ti, GeForce GTX 650 Ti Boost, GeForce GTX 660, GeForce GTX 660 Ti, GeForce GTX 660M, GeForce GTX 670, GeForce GTX 670MX, GeForce GTX 675MX, GeForce GTX 680, GeForce GTX 680M, GeForce GTX 680MX, GeForce GTX 690, GeForce GTX 745, GeForce GTX 750, GeForce GTX 750 Ti, GeForce GTX 760, GeForce GTX 760 Ti, GeForce GTX 760M, GeForce GTX 765M, GeForce GTX 770, GeForce GTX 770M, GeForce GTX 780, GeForce GTX 780 Ti, GeForce GTX 780M, GeForce GTX 850M, GeForce GTX 860M, GeForce GTX 870M, GeForce GTX 880M, GeForce GTX 960, GeForce GTX 970, GeForce GTX 970M, GeForce GTX 980, GeForce GTX 980M, GeForce GTX 980M, GeForce GTX Titan, GeForce GTX Titan Black, GeForce GTX Titan Z, GF108, GF117, GK104, GK106, GK107, GK110, GK208, GK20A, GK210, GM107, GM108, GM204, GRID K1, GRID K2, NVS 510, NVS 5200M, NVS 5400M, Quadro 410, Quadro K1000M, Quadro K1100M, Quadro K2000, Quadro K2000D, Quadro K2000M, Quadro K2100M, Quadro K2200, Quadro K3000M, Quadro K3100M, Quadro K4000, Quadro K4000M, Quadro K4100M, Quadro K420, Quadro K4200, Quadro K5000, Quadro K5000M, Quadro K500M, Quadro K5100M, Quadro K510M, Quadro K5200, Quadro K600, Quadro K6000, Quadro K610M, and Quadro K620.

48. Each of the products in the preceding paragraph is or contains an NVIDIA GPU fabricated utilizing 28 nanometer fabrication processing.

49. The “**Accused 40 nm and Other GPUs**” means all NVIDIA GPUs utilizing 40 nanometer, 55 nanometer, 65 nanometer, 80 nanometer, or 90 nanometer fabrication processing

and all products that contain such a GPU, and includes, without limitation, the G71, G71GL, G71GLM, G72GL, G72GLM, G72M, G73GL, G73GLM, G80, G84, G84GL, G84M, G86, G86M, G92, G92428B1, G92a, G92b, G92M, G94, G94300A1, G94a, G94b, G94M, G96, G96b, G96M, G98, G98M, GeForce 205, GeForce 210, GeForce 305M, GeForce 310, GeForce 310M, GeForce 315, GeForce 315M, GeForce 320M, GeForce 405, GeForce 410M, GeForce 510, GeForce 6052, GeForce 610M, GeForce 8100 mGPU, GeForce 8200 mGPU, GeForce 8200M G, GeForce 8300 GS, GeForce 8300 mGPU, GeForce 8400 GS, GeForce 8400 GS rev.3, GeForce 8400M G, GeForce 8400M GS, GeForce 8400M GT, GeForce 8500 GT, GeForce 8600 GS, GeForce 8600 GT, GeForce 8600 GTS, GeForce 8600M GS, GeForce 8600M GT, GeForce 8700M GT, GeForce 8800 GS, GeForce 8800 GT, GeForce 8800 GTS, GeForce 8800 GTS 112 , GeForce 8800 GTX, GeForce 8800 Ultra, GeForce 8800M GTS, GeForce 8800M GTX, GeForce 9100M G mGPU, GeForce 9200M GS, GeForce 9300 GE, GeForce 9300 GS, GeForce 9300 mGPU, GeForce 9300M G, GeForce 9300M GS, GeForce 9400 GT, GeForce 9400 mGPU, GeForce 9400M G, GeForce 9500 GT, GeForce 9500M G, GeForce 9500M GS, GeForce 9600 GS, GeForce 9600 GSO, GeForce 9600 GSO 512, GeForce 9600 GT, GeForce 9600 GT Green Edition, GeForce 9600M GS, GeForce 9600M GT, GeForce 9650M GS, GeForce 9650M GT, GeForce 9700M GT, GeForce 9700M GTS, GeForce 9800 GT, GeForce 9800 GT Green Edition, GeForce 9800 GTX, GeForce 9800 GTX+, GeForce 9800 GX2, GeForce 9800M GS, GeForce 9800M GT, GeForce 9800M GTS, GeForce 9800M GTX, GeForce G 100, GeForce G 102M, GeForce G 103M, GeForce G 105M, GeForce G 110M, GeForce G210M, GeForce GT 120, GeForce GT 120M, GeForce GT 130, GeForce GT 130M, GeForce GT 140, GeForce GT 220, GeForce GT 220M, GeForce GT 230, GeForce GT 230M, GeForce GT 240, GeForce GT 240M, GeForce GT 320, GeForce GT 320M, GeForce GT 325M, GeForce GT 330 , GeForce

GT 330M, GeForce GT 335M, GeForce GT 340, GeForce GT 415M, GeForce GT 420, GeForce GT 420M, GeForce GT 425M, GeForce GT 430, GeForce GT 435M, GeForce GT 440, GeForce GT 445M, GeForce GT 520, GeForce GT 520M, GeForce GT 520MX, GeForce GT 525M, GeForce GT 530, GeForce GT 540M, GeForce GT 545, GeForce GT 550M, GeForce GT 555M, GeForce GT 610, GeForce GT 620, GeForce GT 625, GeForce GT 630M, GeForce GT 635M, GeForce GT 6408, GeForce GT 640M LE, GeForce GT 645, GeForce GT 705, GeForce GT 730, GeForce GTS 150, GeForce GTS 150M, GeForce GTS 160M, GeForce GTS 240, GeForce GTS 250, GeForce GTS 250 Green, GeForce GTS 250M, GeForce GTS 260M, GeForce GTS 350M, GeForce GTS 360M, GeForce GTS 450, GeForce GTX 260, GeForce GTX 260 Core 216, GeForce GTX 260M, GeForce GTX 275, GeForce GTX 280, GeForce GTX 280M, GeForce GTX 285, GeForce GTX 285M, GeForce GTX 295, GeForce GTX 460, GeForce GTX 460 SE, GeForce GTX 460M, GeForce GTX 465, GeForce GTX 470, GeForce GTX 470M, GeForce GTX 480, GeForce GTX 480M, GeForce GTX 485M, GeForce GTX 550 Ti, GeForce GTX 555, GeForce GTX 560, GeForce GTX 560 SE, GeForce GTX 560 Ti, GeForce GTX 560 Ti 448 Cores Limited Edition, GeForce GTX 560M, GeForce GTX 570, GeForce GTX 570M, GeForce GTX 580, GeForce GTX 580M, GeForce GTX 590, GeForce GTX 670M, GeForce GTX 675M, GF100, GF104, GF104-225, GF104-300-KB, GF106, GF106GL, GF108, GF108GL, GF110, GF1105, GF114, GF116, GF118, GF119, GT200, GT-200, GT200GL, GT215, GT215M, GT216, GT216GL, GT216M, GT218, GT218GL, GT218M, MCP51, MCP77MH, MCP77MV, MCP78, MCP79MH, MCP79MVL, MCP79MX, MCP79XT, MCP7AS, MCP7AU, MCP89, NV43GL, NVS 2100M, NVS 210S, NVS 295, NVS 300, NVS 310, NVS 3100M, NVS 315, NVS 420, NVS 4200M, NVS 5100M, NVS 5200M, NVS 5400M, Quadro 1000M, Quadro 2000, Quadro 2000M, Quadro 3000M, Quadro 400, Quadro 4000, Quadro 4000M, Quadro 5000,

Quadro 5000M, Quadro 500M, Quadro 5010M, Quadro 600, Quadro 6000, Quadro 7000, Quadro CX, Quadro FX 1500, Quadro FX 1500M, Quadro FX 1600M, Quadro FX 1700, Quadro FX 1700M, Quadro FX 1800, Quadro FX 1800M, Quadro FX 2500M, Quadro FX 2700M, Quadro FX 2800M, Quadro FX 350, Quadro FX 3500, Quadro FX 3500M, Quadro FX 350M, Quadro FX 3600M, Quadro FX 360M, Quadro FX 370, Quadro FX 370 LP, Quadro FX 3700, Quadro FX 3700M, Quadro FX 370M, Quadro FX 380, Quadro FX 380 LP, Quadro FX 3800, Quadro FX 3800M, Quadro FX 380M, Quadro FX 4600, Quadro FX 4600 SDI, Quadro FX 470, Quadro FX 4700X2, Quadro FX 4800, Quadro FX 540, Quadro FX 550, Quadro FX 5500, Quadro FX 5500 SDI, Quadro FX 560, Quadro FX 5600, Quadro FX 560M, Quadro FX 570, Quadro FX 570M, Quadro FX 580, Quadro FX 5800, Quadro FX 770M, Quadro FX 880M, Quadro NVS 110M, Quadro NVS 120M, Quadro NVS 130M, Quadro NVS 135M, Quadro NVS 140M, Quadro NVS 150M, Quadro NVS 160M, Quadro NVS 300M, Quadro NVS 320M, Quadro NVS 510M, Quadro VX 200, and RSX.

50. Each of the products in the preceding paragraph is or contains an NVIDIA GPU fabricated utilizing 40 nanometer, 55 nanometer, 65 nanometer, 80 nanometer, or 90 nanometer fabrication processing.

51. The “**724 Accused Mobile GPUs**” means all NVIDIA GPUs that support DisplayPort and that are used, or intended for use, in laptop computers and includes, without limitation, the GeForce 9800M GTX, GeForce 9800M GTS, GeForce 9800M GS, GeForce 9700M GTS, GeForce 9700M GT, GeForce 9650M GT, GeForce 9650M GS, GeForce 9600M GT, GeForce 9600M GS, GeForce 9500M GS, GeForce 9500M G, GeForce GTX 980M, GeForce GTX 970M, GeForce GTX 880M, GeForce GTX 870M, GeForce GTX 860M, GeForce GTX 850M, GeForce 840M, GeForce 830M, GeForce 820M, GeForce GTX 780M, GeForce

GTX 770M, GeForce GTX 765M, GeForce GTX 760M, GeForce GT 755M, GeForce GT 750M, GeForce GT 745M, GeForce GT 740M, GeForce GT 735M, GeForce GT 730M, GeForce GT 720M, GeForce 710M, GeForce GTX 285M, GeForce GTX 280M, GeForce GTX 260M, GeForce GTS 360M, GeForce GTS 350M, GeForce GTS 260M, GeForce GTS 250M, GeForce GTS 160M, GeForce GTS 150M, GeForce GT 335M, GeForce GT 330M, GeForce GT 325M, GeForce GT 240M, GeForce GT 230M, GeForce GT 130M, GeForce G210M, GeForce G110M, GeForce G105M, GeForce G102M, GeForce 410M, GeForce 410M, GeForce 405M, GeForce 315M, GeForce 310M, and GeForce 305M.

52. Each of the products in the preceding paragraph is or contains an NVIDIA GPU that supports DisplayPort and that is used, or intended for use, in a laptop computer.

53. The “**724 Accused Mobile GPUs with Analog Output**” means all NVIDIA GPUs that support DisplayPort and an analog output, such as DVI or VGA, and that are intended for use or used in laptop computers and includes, without limitation, the GeForce 305M, GeForce 310M, GeForce 315M, GeForce 405M, GeForce 410M, GeForce 410M, GeForce 710M, GeForce 820M, GeForce 830M, GeForce 840M, GeForce 9500M G, GeForce 9500M GS, GeForce 9600M GS, GeForce 9600M GT, GeForce 9650M GS, GeForce 9650M GT, GeForce 9700M GT, GeForce 9700M GTS, GeForce 9800M GS, GeForce 9800M GTS, GeForce 9800M GTX, GeForce G102M, GeForce G105M, GeForce G110M, GeForce G210M, GeForce GT 130M, GeForce GT 230M, GeForce GT 240M, GeForce GT 325M, GeForce GT 330M, GeForce GT 335M, GeForce GT 720M, GeForce GT 730M, GeForce GT 735M, GeForce GT 740M, GeForce GT 745M, GeForce GT 750M, GeForce GT 755M, GeForce GTS 150M, GeForce GTS 160M, GeForce GTS 250M, GeForce GTS 260M, GeForce GTS 350M, GeForce GTS 360M, GeForce GTX 260M, GeForce GTX 280M, GeForce GTX 285M, GeForce GTX

760M, GeForce GTX 765M, GeForce GTX 770M, GeForce GTX 780M, GeForce GTX 850M, GeForce GTX 860M, GeForce GTX 870M, GeForce GTX 880M, GeForce GTX 970M, and GeForce GTX 980M.

54. Each of the products in the preceding paragraph is or contains an NVIDIA GPU that supports DisplayPort and an analog output, such as DVI or VGA, and that is used, or intended for use, in a laptop computer.

55. The “**Unified Cache GPUs**” means all NVIDIA GPUs that include a unified L2 cache, as well as all products that contain such a GPU, and includes, without limitation, all NVIDIA GPUs using the Fermi, Kepler, or Maxwell architectures, the GF100, GF104, GF104-225-A1, GF104-300-KB-A1, GF106, GF106 (N12E-GE2), GF106-250, GF106GL (GF106-875), GF108, GF108 (N13P-GL), GF108GL, GF110, GF110-270-A15, GF110-275-A15, GF110-351-A1, GF110-375-A15, GF1105, GF114, GF114 (N13E-GS1), GF114 (N13E-GS1-LP), GF1144, GF114-400-A1, GF116, GF116-200, GF116-400, GF117, GF117 (N13M-GS), GF118, GF119, GF119 (N13M-GE), GF119-300-A1, GK104, GK104 (N13E-GR), GK104 (N13E-GSR), GK104 (N13E-GTX), GK104-200-KD-A2, GK104-225-A2, GK104-300-KD-A2, GK104-325-A2, GK104-355-A2, GK104-400-A2, GK104-425-A2, GK104-895, GK106, GK106-220-A1, GK106-240-A1, GK106-400-A1, GK107, GK107 (N13E-GE), GK107 (N13P-GS), GK107 (N13P-GT), GK107-425-A2, GK107-450-A2, GK110, GK110-300-A1, GK110-400-A1, GK110-425-B1, GK110-430, GK208, GK208-301-A1, GM107, GM107-300-A2, GM107-400-A2, GM108, GM204, GeForce 405, GeForce GT 420, GeForce GT 430, GeForce GT 430, GeForce GT 440, GeForce GT 440, GeForce GTS 450, GeForce GTS 450, GeForce GTX 460 SE, GeForce GTX 460, GeForce GTX 460, GeForce GTX 460 v2, GeForce GTX 465, GeForce GTX 470, GeForce GTX 480, GeForce 410M, GeForce GT 415M, GeForce GT 420M, GeForce

GT 425M, GeForce GT 435M, GeForce GT 445M, GeForce GTX 460M, GeForce GTX 470M, GeForce GTX 480M, GeForce GTX 485M, GeForce 510, GeForce GT 520, GeForce GT 530, GeForce GT 545 DDR3, GeForce GT 545 GDDR5, GeForce GTX 550 Ti, GeForce GTX 555, GeForce GTX 560 SE, GeForce GTX 560, GeForce GTX 560 Ti, GeForce GTX 560 Ti OEM, GeForce GTX 560 Ti 448 Cores Limited Edition, GeForce GTX 570, GeForce GTX 580, GeForce GTX 590, GeForce GT 520M, GeForce GT 520MX, GeForce GT 525M, GeForce GT 540M, GeForce GT 550M, GeForce GT 555M, GeForce GTX 560M, GeForce GTX 570M, GeForce GTX 580M, GeForce 6052, GeForce GT 610, GeForce GT 620, GeForce GT 625, GeForce GT 630, GeForce GT 630 (Rev. 2), GeForce GT 635, GeForce GT 6408, GeForce GT 640 (DDR3), GeForce GT 640 Rev. 2, GeForce GT 645, GeForce GTX 645, GeForce GTX 650, GeForce GTX 650 Ti, GeForce GTX 650 Ti Boost, GeForce GTX 660, GeForce GTX 660, GeForce GTX 660 Ti, GeForce GTX 670, GeForce GTX 680, GeForce GTX 690, GeForce 610M, GeForce GT 620M, GeForce GT 625M, GeForce GT 630M, GeForce GT 635M, GeForce GT 640M LE, GeForce GT 640M, GeForce GT 645M, GeForce GT 650M, GeForce GTX 660M, GeForce GTX 670M, GeForce GTX 670MX, GeForce GTX 675M, GeForce GTX 675MX, GeForce GTX 680M, GeForce GTX 680MX, GeForce GT 705, GeForce GT 710, GeForce GT 720, GeForce GT 730, GeForce GT 740, GeForce GTX 745, GeForce GTX 750, GeForce GTX 750 Ti, GeForce GTX 760 192-bit, GeForce GTX 760, GeForce GTX 760 Ti, GeForce GTX 770, GeForce GTX 780, GeForce GTX 780 Ti, GeForce GTX Titan, GeForce GTX Titan Black, GeForce GTX Titan Z, GeForce 710M, GeForce GT 720M, GeForce GT 730M, GeForce GT 735M, GeForce GT 740M, GeForce GT 740M, GeForce GT 745M, GeForce GT 750M, GeForce GT 755M, GeForce GTX 760M, GeForce GTX 765M, GeForce GTX 770M, GeForce GTX 780M, GeForce 820M, GeForce 825M, GeForce 830M, GeForce 840M,

GeForce GTX 850M, GeForce GTX 860M, GeForce GTX 870M, GeForce GTX 880M, GeForce GTX 960, GeForce GTX 970, GeForce GTX 980, GeForce GTX 970M, GeForce GTX980M, Quadro 600, Quadro 2000, Quadro 4000, Quadro 5000, Quadro 6000, Quadro 7000, Quadro 410, Quadro K600, Quadro K2000, Quadro K2000D, Quadro K4000, Quadro K5000, Quadro K6000, Quadro K420, Quadro K620, Quadro K2200, Quadro K4200, Quadro K5200, Quadro NVS 295, Quadro NVS 300, Quadro NVS 310, Quadro NVS 315, Quadro NVS 510, Tesla C2050, Tesla C2070, Tesla C2075, Tesla M2050, Tesla M2070, Tesla M2070, Tesla M2070Q, Tesla M2090, Tesla S2050, Tesla K10, Tesla K20, Tesla K20X, Tesla K40, Tesla K80, Quadro 500M, Quadro 1000M, Quadro 2000M, Quadro 3000M, Quadro 4000M, Quadro 5000M, Quadro 5010M, Quadro K500M, Quadro K1000M, Quadro K2000M, Quadro K3000M, Quadro K4000M, Quadro K5000M, Quadro K510M, Quadro K610M, Quadro K1100M, Quadro K2100M, Quadro K3100M, Quadro K4100M, Quadro K5100M, Quadro NVS 4200M, Quadro NVS 5200M, Quadro NVS 5400M, GRID K1, and GRID K2.

56. Each of the products in the preceding paragraph is or contains an NVIDIA GPU with a unified L2 cache. Each of the products in the preceding paragraph is or contains an NVIDIA GPU using the Fermi, Kelper, or Maxwell architectures.

57. The “**Accused SOCs**” means all NVIDIA SOCs and all products that contain a NVIDIA SOC that have been made, used, sold, offered for sale, or imported into the United States at any time since November 4, 2008, including the products identified in paragraphs 60 to 64, below.

58. The “**Accused 28 nm SOCs**” means all NVIDIA SOCs utilizing 28 nanometer fabrication processing and all products that contain such an SOC, and includes, without limitation, the Tegra 4 114, Tegra 4i, and the Tegra K1.

59. Each of the SOCs in the preceding paragraph are fabricated utilizing 28 nanometer fabrication processing.

60. The “**Accused 40 nm and Other SOCs**” means all NVIDIA SOCs utilizing 40 nanometer, 55 nanometer, 65 nanometer, 80 nanometer, or 90 nanometer fabrication processing and all products that contain such a SOC, and includes, without limitation, the Tegra 250, Tegra 250 3D AP25, Tegra 250 3D T25, Tegra 250 AP20H, Tegra 250 T20, Tegra 3 AP33, Tegra 3 T30L, Tegra 3 T33, Tegra 600, Tegra 650, Tegra APX 2500, and the Tegra APX 2600.

61. Each of the SOCs in the preceding paragraph are fabricated utilizing 40 nanometer, 55 nanometer, 65 nanometer, 80 nanometer, or 90 nanometer fabrication processing.

62. The “**Cortex-A9 SOCs**” means all NVIDIA SOCs based on the ARM Cortex-A9 design and all products that contain such an SOC, and includes, without limitation, all SOCs in the Tegra 2 and Tegra 3 families, the Tegra 250, Tegra 250 AP20H, Tegra 250 T20, Tegra 250 3D AP25, Tegra 250 3D T25, Tegra 3 T30L, Tegra 3 T30L, Tegra 3 AP33, and Tegra 3 T33 SOCs.

63. Each of the SOCs in the preceding paragraph is based on the ARM Cortex-A9 design.

64. The “**Cortex-A15 SOCs**” means all NVIDIA SOCs based on the ARM Cortex-A15 design and all products that contain such an SOC, and includes, without limitation, all SOCs in the Tegra 4 and Tegra K1 families, the Tegra 4, Tegra 4 114, Tegra 4i, and Tegra K1 SOCs.

65. Each of the SOCs in the preceding paragraph is based on the ARM Cortex-A15 design.

66. The “**158 Accused Products**” include all **Accused SOCs** (including all **Cortex-A9 SOCs** and all **Cortex-A15 SOCs**) and all **Unified Cache GPUs**.

67. The “**’938 Accused Products**” include all **Accused SOCs** and **’938 Accused GPUs** and/or all systems containing any of the **Accused SOCs** or **’938 Accused GPUs** and JEDEC-standard SDRAM or SGRAM.

68. The “**’902 Accused Products**” include all **Accused Products** (including all **Accused GPUs** and all **Accused SOCs**).

69. The “**’602 Accused Products**” include all **’602 Accused GPUs** and all **Accused SOCs**.

70. The “**’675 Accused Products**” include all **Accused 28 nm GPUs** and all **Accused 28 nm SOCs**.

71. The “**’054 Accused Products**” means all Velocity computer products that contain a hybrid hard drive—that is, a hard drive with both a spinning platter and separate, solid-state non-volatile storage capacity—that have been made, used, sold, offered for sale, or imported into the United States at any time since November 4, 2008. The **’054 Accused Products** include but are not limited to, the NoteMagix M15 laptop computer having a hybrid hard drive (hereinafter the “**NoteMagix M15 HHD**”), which was available at least as of November 4, 2014.

72. The “**’854 Accused Products**” include all Velocity computers and computer cases that include flexible contacts on the computer case, including without limitation the Velocity GX-2 Computer Case.

73. The “**’724 Accused Velocity Laptops**” are all Velocity laptop computers with a DisplayPort port, including without limitation, the NoteMagix M15, NoteMagix M17, NoteMagix M15 Ultra, and NoteMagix M17 Ultra.

74. Each of the computers in the preceding paragraph is a Velocity laptop computer with a DisplayPort port.

75. The “**724 Accused Products**” means all **724 Accused Velocity Laptops**, all **724 Accused Mobile GPUs** (including all **724 Accused Mobile GPUs with Analog Output**), and all laptop computers containing **724 Accused Mobile GPUs**.

REFERENCES

76. References to the document “**About Samsung**” refer to the document contemporaneously produced as SAMS-NVD-0002095 to 2098, containing a Samsung webpage describing Samsung’s History.

77. References to the document “**Cuda Defense Workshops**” refer to the document contemporaneously produced as SAMS-NVD-0002105 to 2106, containing an Acceleware website, available at <http://www.acceleware.com/cuda-defense-workshops>.

78. References to the document “**Velocity NVIDIA GeForce Press Release**” refer to the document contemporaneously produced as SAMS-NVD-0002135 to 2136, containing a Velocity press release entitled “Velocity Micro® Announces the Immediate Availability of NVIDIA GeForce GTX 750, 750 Ti, and TITAN Black Graphics.”

79. **Velocity NVIDIA GeForce Press Release** is a true and accurate copy of a Velocity press release.

80. References to the document “**Velocity Desktop Press Release**” refer to the document contemporaneously produced as SAMS-NVD-0002132 to 2133, containing a Velocity press release entitled “Velocity Micro announces desktops powered by NVIDIA GeForce GTX 780.”

81. **Velocity Desktop Press Release** is a true and accurate copy of a Velocity press release.

82. References to the document “**Raptor X17 Press Release**” refer to the document contemporaneously produced as SAMS-NVD-0002131, containing a Velocity press release entitled “Velocity Micro Announces Raptor X17 Notebook Featuring NVIDIA® GeForce® GTX 480M Mobile Graphics.”

83. **Raptor X17 Press Release** is a true and accurate copy of a Velocity press release.

84. References to the document “**Velocity NVIDIA Partner Press Release**” refer to the document contemporaneously produced as SAMS-NVD-0002137, containing a Velocity press release entitled “Velocity Micro Partners With NVIDIA® To Launch GeForce 6 Series of Graphics.”

85. **Velocity NVIDIA Partner Press Release** is a true and accurate copy of a Velocity press release.

86. References to the document “**NVIDIA Shield**” refer to the document contemporaneously produced as SAMS-NVD-0002110 to 2130, containing an NVIDIA document entitled “The Ultimate Tablet for Gamers.”

87. **NVIDIA Shield** is a true and accurate copy of an NVIDIA webpage.

88. References to the document “**NVIDIA GF100 Whitepaper**” refer to the document contemporaneously produced as SAMS-NVD-0001102 to 1132, entitled “Whitepaper – NVIDIA GF100 – World’s Fastest GPU Delivering Great Gaming Performance with True Geometric Realism V1.5,” NVIDIA (2010).

89. The **NVIDIA GF100 Whitepaper** is a true and correct copy of a document published by NVIDIA in or around 2010.

90. The **NVIDIA GF100 Whitepaper** accurately describes the NVIDIA GF100 GPU.

91. The NVIDIA GF100 GPU uses the Fermi architecture.

92. The **NVIDIA GF100 Whitepaper** also accurately describes other GPUs using the Fermi architecture.

93. References to the document “**NVIDIA Kepler GK110 Whitepaper**” refer to the document contemporaneously produced as SAMS-NVD-0001133 to 1156, entitled “Whitepaper – NVIDIA’s Next Generation CUDA Compute Architecture: Kepler GK110 V1.0,” NVIDIA (2012).

94. The **NVIDIA Kepler GK110 Whitepaper** is a true and correct copy of a document published by NVIDIA in or around 2012.

95. The **NVIDIA Kepler GK110 Whitepaper** accurately describes the NVIDIA GK110 GPU.

96. The NVIDIA GK110 GPU uses the Kepler architecture.

97. The **NVIDIA Kepler GK110 Whitepaper** also accurately describes other GPUs using the Kepler architecture.

98. References to the document “**NVIDIA GeForce GTX 750 Ti Whitepaper**” refer to the document contemporaneously produced as SAMS-NVD-0001091 to 1101, entitled “Whitepaper – NVIDIA GeForce GTX 750 Ti – Featuring First-Generation Maxwell GPU Technology, Designed for Extreme Performance per Watt V1.1,” NVIDIA (2014).

99. The **NVIDIA GeForce GTX 750 Ti Whitepaper** is a true and correct copy of a document published by NVIDIA in or around 2014.

100. The **NVIDIA GeForce GTX 750 Ti Whitepaper** accurately describes the NVIDIA GM107 GPU.

101. The NVIDIA GM107 GPU uses the Maxwell architecture.

102. The **NVIDIA GeForce GTX 750 Ti Whitepaper** also accurately describes other GPUs using the Maxwell architecture.

103. References to the document “**PL310 TRM**” refer to the document contemporaneously produced as SAMS-NVD-0001157 to 1334, entitled “PrimeCell Level 2 Cache Controller (PL310) Revision r2p0 Technical Reference Manual,” ARM (2008).

104. The **PL310 TRM** is a true and correct copy of a document published by ARM in or around 2008.

105. The **PL310 TRM** accurately describes the PL310 cache controller.

106. The **PL310 TRM** accurately describes the PL310 cache controllers used in SOCs in the Tegra 2 and Tegra 3 families.

107. NVIDIA relied on the **PL310 TRM** in the design of the SOCs in the Tegra 2 and Tegra 3 families.

108. References to the document “**AXI Spec**” refer to the document contemporaneously produced as SAMS-NVD-0000158 to 485, entitled “AMBA AXI and ACE Protocol Specification,” ARM (2013).

109. The **AXI Spec** is a true and correct copy of a document published by ARM in or around 2013.

110. The **AXI Spec** accurately describes communication protocols used by the PL310 cache controllers.

111. References to the document “**Cortex-A9 TRM**” refer to the document contemporaneously produced as SAMS-NVD-0000878 to 1090, entitled “Cortex-A9 Technical Reference Manual Revision r4p1,” ARM (2012).

112. The **Cortex-A9 TRM** is a true and correct copy of a document published by ARM in or around 2012.

113. The **Cortex-A9 TRM** accurately describes the ARM Cortex-A9 design.

114. The **Cortex-A9 TRM** accurately describes certain technical characteristics of the **Cortex-A9 SOCs**.

115. NVIDIA relied on the **Cortex-A9 TRM** in the design of the **Cortex-A9 SOCs**.

116. References to the document "**Cortex-A15 TRM**" refer to the document contemporaneously produced as SAMS-NVD-0000486 to 877, entitled "ARM Cortex-A15 MPCore Processor Technical Reference Manual Revision r4p0," ARM (2013).

117. The **Cortex-A15 TRM** is a true and correct copy of a document published by ARM in or around 2013.

118. The **Cortex-A15 TRM** accurately describes the ARM Cortex-A15 design.

119. The **Cortex-A15 TRM** accurately describes certain technical characteristics of the **Cortex-A15 SOCs**.

120. NVIDIA relied on the **Cortex-A15 TRM** in the design of the **Cortex-A15 SOCs**.

121. References to the document "**GK107 Report**" refer to the document contemporaneously produced as SAMS-NVD-0001752 to 1902, entitled "Logic detailed structural Analysis of the nVidia GK107 28nm Graphic Processor," TechInsights (November 29, 2012).

122. The **GK107 Report** describes the NVIDIA GK107 GPU.

123. The **GK107 Report** contains true and accurate scanning electron microscopy ("SEM") and transmission electron microscopy ("TEM") images of the GK107 GPU.

124. The GK107 GPU is used in the NVIDIA GeForce GTX 660 graphics card.

125. The GK107 GPU uses NVIDIA's Kepler GPU architecture.

126. References to the document "**Tegra 250 Report**" refer to the document contemporaneously produced as SAMS-NVD-0001935 to 2084, entitled "Logic detailed structural Analysis of the nVidia Tegra 250 AP," TechInsights (November 17, 2011).

127. The **Tegra 250 Report** contains true and accurate SEM and TEM images of the SOCs in the Tegra 2 family.

128. References to the document "**Process Images Report**" refer to the document contemporaneously produced as SAMS-NVD-0001911 to 1931, entitled "Custom Process Images Tegra 250 AP & GK107 GP," TechInsights (May, 2013).

129. The **Process Images Report** contains true and accurate SEM and TEM images of the NVIDIA GK107 GPU.

130. The **Process Images Report** contains true and accurate SEM and TEM images of SOCs in the Tegra 2 family.

131. References to the document "**90-nm CMOS Device Technology**" refer to the document contemporaneously produced as SAMS-NVD-0001746 to 1749, entitled "A 90-nm CMOS Device Technology with High-speed, General-purpose, and Low-leakage Transistors for System on Chip Applications" by Wu et al., 2002 International Electron Devices meeting (December 8-11, 2002).

132. References to the document "**TSMC 28nm Technology**" refer to the document contemporaneously produced as SAMS-NVD-0002087 to 2088, containing a TSMC webpage entitled "28nm Technology."

133. **TSMC 28nm Technology** is a true and accurate copy of the original document available at <http://www.tsmc.com/english/dedicatedFoundry/technology/28nm.htm>.

134. References to the document “**TSMC 55nm Technology**” refer to the document contemporaneously produced as SAMS-NVD-0002089, containing a TSMC webpage entitled “55nm Technology.”

135. **TSMC 55nm Technology** is a true and accurate copy of the original document available at <http://www.tsmc.com/english/dedicatedFoundry/technology/55nm.htm>.

136. References to the document “**TSMC 28HPC Process**” refer to the document contemporaneously produced as SAMS-NVD-0002085 to 2086, entitled “TSMC 28HPC Process in Volume production,” TSMC Press Release (September 12, 2014).

137. **TSMC 28HPC Process** is a true and accurate copy of a document published by TSMC on or about September 12, 2014, available at <http://www.tsmc.com/tsmcdotcom/PRListingNewsAction.do?action=detail&language=E&newsid=8761>.

138. References to the document “**Half-Node Process**” refer to the document contemporaneously produced as SAMS-NVD-0001903, entitled “ATI, Nvidia back TSMC 80-nm ‘half-node’ process,” Peter Clarke, EE Times (January 17, 2006).

139. **Half-Node Process** is a true and accurate copy of an article published by EE Times on or about January 17, 2006, available at http://www.eetimes.com/document.asp?doc_id=1158740&print=yes.

140. References to the document “**GeForce GTX 660 Features**” refer to the document contemporaneously produced as SAMS-NVD-0001750 to 1751, containing an NVIDIA webpage for the GeForce GTX 660.

141. The **GeForce GTX 660 Features** is a true and correct copy of an NVIDIA webpage for the GeForce GTX 660.

142. The **GeForce GTX 660 Features** accurately describes the NVIDIA GeForce GTX 660.

143. References to the document “**Tegra 2 Press Release**” refer to the document contemporaneously produced as SAMS-NVD-0001932 to 1934, entitled “New NVIDIA Tegra Processor Powers The Tablet Revolution,” NVIDIA Press Release (January 7, 2010).

144. The **Tegra 2 Press Release** is a true and correct copy of a document published by NVIDIA on or about January 7, 2010.

145. The **Tegra 2 Press Release** accurately describes SOCs in the Tegra 2 family.

146. References to the document “**NVIDIA Kepler Press Release**” refer to the document contemporaneously produced as SAMS-NVD-0001906 to 1910, entitled “NVIDIA Launches First GeForce GPUs Based on Next-Generation Kepler Architecture,” NVIDIA Press Release (March 22, 2012).

147. The **NVIDIA Kepler Press Release** is a true and correct copy of a document published by NVIDIA on or about March 22, 2012.

148. The **NVIDIA Kepler Press Release** accurately describes the NVIDIA Kepler architecture.

149. References to the document “**History on Side of Gate-Last High-k Approach**” refer to the document contemporaneously produced as SAMS-NVD-0001904 to 1905, entitled “TSMC’s Chiang Sees History on Side of Gate-Last High-k Approach,” David Lammers, Semiconductor International (February 10, 2010).

150. The **History on Side of Gate-Last High-k Approach** is a true and correct copy of a document published by Semiconductor International on or about February 10, 2012.

151. References to the document “**TSMC 65nm Process**” refer to the document contemporaneously produced as SAMS-NVD-0002090 to 2094, entitled “Under the Hood: Focus on Qualcomm and TI at 65 nm,” John Boyd, Semiconductor Insights, EE Times (May 14, 2007).

152. **TSMC 65nm Process** is a true and accurate copy of an article published by EE Times on or about May 14, 2007, available at http://www.eetimes.com/document.asp?doc_id=1281172&print=yes.

153. References to the document “**GeForce G100 GPU Page**” refer to the document contemporaneously produced as SAMS-NVD-0002107 to 2108, containing a Velocity webpage entitled “GeForce G100.”

154. **GeForce G100 GPU Page** is a true and accurate copy of a Velocity webpage.

155. References to the document “**Configure Your Z35 Archive**” refer to the document contemporaneously produced as SAMS-NVD-0002103 to 2104, containing a Velocity webpage entitled “Configure Your Z35.”

156. **Configure Your Z35 Archive** contains a true and accurate copy of a Velocity webpage.

157. References to the document “**About Velocity**” refer to the document contemporaneously produced as SAMS-NVD-0002099 to 2100, containing a Velocity webpage entitled “About Us.”

158. **About Velocity** is a true and accurate copy of a Velocity webpage.

159. References to the document “**GeForce G100 Purchase Page**” refer to the document contemporaneously produced as SAMS-NVD-0002109, containing a webpage entitled “NVIDIA Nvidia G100 DDR2 256MB DVI-i PCI Express Low Bracket Video Graphics.”

160. **GeForce G100 Purchase Page** is a true and accurate copy of a webpage offering for sale an NVIDIA GeForce G100.

161. References to the document “**Configure Your M17**” refer to the document contemporaneously produced as SAMS-NVD-0002101 to 2102, containing a Velocity webpage for the NoteMagix M17 laptop computer.

162. **Configure Your M17** is a true and accurate copy of a Velocity webpage.

163. References to the document “**Velocity Gaming and Enthusiast Laptops**” refer to the document contemporaneously produced as SAMS-NVD-0002134, containing a Velocity webpage entitled “Gaming and Enthusiast Laptops.”

164. **Velocity Gaming and Enthusiast Laptops** is a true and accurate copy of a Velocity webpage entitled “Gaming and Enthusiast Laptops.”

165. References to the document “**Velocity Grounding Instructions**” refer to the document contemporaneously produced as SAMS-NVD-0001618, entitled “What are the most common methods for grounding yourself?”

166. The **Velocity Grounding Instructions** is a true and accurate copy of a Velocity Micro PC support article.

167. The **Velocity Grounding Instructions** is a document prepared by Velocity and made available to Velocity’s customers.

168. The **Velocity Grounding Instructions** accurately describes common methods for grounding.

169. References to the document “**Vector Z25 Chassis Image**” refer to the document contemporaneously produced as SAMS-NVD-0001617.

170. The **Vector Z25 Chassis Image** document is a true and accurate copy of an original image available on Velocity's website.

171. The **Vector Z25 Chassis Image** is a document prepared by Velocity and made available to Velocity's customers.

172. The **Vector Z25 Chassis Image** accurately depicts the Velocity Vector Z25 chassis.

173. References to the document "**Configure Your Z25**" refer to the document contemporaneously produced as SAMS-NVD-0001335 to 1337, entitled "Configure Your Velocity Vector Z25"

174. **Configure Your Z25** is a true and accurate copy of the original document available on Velocity's Website.

175. **Configure Your Z25** is a document prepared by Velocity and made available to Velocity's customers.

176. **Configure Your Z25** is designed by Velocity to assist customers with configuring a Velocity personal computer for purchase.

177. References to the document "**DisplayPort Standard**" refer to the document contemporaneously produced as SAMS-NVD-0001375 to 1612, entitled "VESA DisplayPort Standard," VESA (Version 1, Revision 1a, 2008).

178. The **DisplayPort Standard** is a true and correct copy of a document published by the Video Electronics Standards Association ("VESA") in or around 2008.

179. The **DisplayPort Standard** accurately describes the technical requirements for products that support DisplayPort.

180. References to the document “**GeForce GTS 250M Specifications**” refer to the document contemporaneously produced as SAMS-NVD-0001613 to 1614, containing an NVIDIA webpage for the GeForce GTS 250M.

181. The **GeForce GTS 250M Specifications** is a true and correct copy of an NVIDIA webpage for the GeForce GTS 250M.

182. The **GeForce GTS 250M Specifications** accurately describes the NVIDIA GeForce GTS 250M.

183. References to the document “**GeForce GTX 770M Specifications**” refer to the document contemporaneously produced as SAMS-NVD-0001615 to 1616, containing an NVIDIA webpage for the GeForce GTX 770M.

184. The **GeForce GTX 770M Specifications** is a true and correct copy of an NVIDIA webpage for the GeForce GTX 770M.

185. The **GeForce GTX 770M Specifications** accurately describes the NVIDIA GeForce GTX 770M.

186. References to the document “**DCB 4.0**” refer to the document contemporaneously produced as SAMS-NVD-0001338 to 1374, entitled “Device Control Block 4.0 Specification,” NVIDIA.

187. **DCB 4.0** is a true and correct copy of a document available through NVIDIA online.

188. **DCB 4.0** accurately describes certain technical features of the **’724 Accused Mobile GPUs**.

189. References to the document “**Archived NoteMagix M15 HHD Webpage**” refer to the webpage printout contemporaneously produced as SAMS-NVD-0001744 to 1745, entitled “Velocity Micro: NoteMagix™ M15.”

190. The **Archived NoteMagix M15 HHD Webpage** document is a true and accurate copy of the original document available at <https://web.archive.org/web/20140701152730/http://www.velocitymicro.com/wizard.php?iid=241>.

191. The **Archived NoteMagix M15 HHD Webpage** is a document prepared by Velocity and made available to Velocity’s customers.

192. The **Archived NoteMagix M15 HHD Webpage** accurately describes the NoteMagix M15 laptop sold by Velocity.

193. References to the document “**GeForce G100 Specifications**” refer to the document contemporaneously produced as SAMS-NVD-0002271 to 2272, entitled “GeForce G100 Specifications.”

194. The **GeForce G100 Specifications** is a true and correct copy of a document published by NVIDIA and available at <http://www.geforce.com/hardware/desktop-gpus/geforce-g100/specifications>.

195. The **GeForce G100 Specifications** accurately describes the specifications of the GeForce G100 GPU.

196. References to the document “**Tegra 3 Report**” refer to the document contemporaneously produced as SAMS-NVD-0004050 to 4174, entitled “Partial Memory Interface on the NVIDIA T30 Tegra 3 Processor,” TechInsights (July 30, 2013).

197. The **Tegra 3 TRM** accurately describes SOCs in the Tegra 3 family.

198. References to the document “**Tegra 2 TRM**” refer to the document contemporaneously produced as SAMS-NVD-0002651 to 4049, entitled “NVIDIA® Tegra® 2 Series Mobile Processors Technical Reference Manual,” NVIDIA (2011).

199. The **Tegra 2 TRM** is a true and correct copy of a document published by NVIDIA in or around 2011.

200. The **Tegra 2 TRM** accurately describes SOCs in the Tegra 2 family.

201. NVIDIA makes the **Tegra 2 TRM** available to its customers and intends that its customers use and rely on the **Tegra 2 TRM**.

202. References to the document “**Tegra 3 TRM**” refer to the document contemporaneously produced as SAMS-NVD-0004175 to 6147, entitled “NVIDIA® Tegra® 3 HD Mobile Processors: T30 Series, AP30 Series Technical Reference Manual,” NVIDIA (2011).

203. The **Tegra 3 TRM** is a true and correct copy of a document published by NVIDIA in or around 2011.

204. The **Tegra 3 TRM** accurately describes SOCs in the Tegra 3 family.

205. NVIDIA makes the **Tegra 3 TRM** available to its customers and intends that its customers use and rely on the **Tegra 3 TRM**.

206. References to the document “**Tegra 4 TRM**” refer to the document contemporaneously produced as SAMS-NVD-0006148 to 8810, entitled “NVIDIA® Tegra® 4 Processors Technical Reference Manual,” NVIDIA (2013).

207. The **Tegra 4 TRM** is a true and correct copy of a document published by NVIDIA in or around 2013.

208. The **Tegra 4 TRM** accurately describes SOCs in the Tegra 4 family.

209. NVIDIA makes the **Tegra 4 TRM** available to its customers and intends that its customers use and rely on the **Tegra 4 TRM**.

210. References to the document “**Tegra K1 TRM**” refer to the document contemporaneously produced as SAMS-NVD-0008811 to 11153, entitled “NVIDIA® Tegra® K1 Mobile Processor Technical Reference Manual,” NVIDIA (2014).

211. The **Tegra K1 TRM** is a true and correct copy of a document published by NVIDIA in or around 2014.

212. The **Tegra K1 TRM** accurately describes SOCs in the Tegra K1 family.

213. NVIDIA makes the **Tegra K1 TRM** available to its customers and intends that its customers use and rely on the **Tegra K1 TRM**.

214. References to the document “**JEDEC Standard No. 79-3F**” refer to the document contemporaneously produced as SAMS-NVD-0002401 to 2626, entitled “JEDEC DDR3 Standard JESD79-3F,” JEDEC (July 2012).

215. The **JEDEC Standard No. 79-3F** is a true and correct copy of a document published by JEDEC in or around July 2012.

216. The **JEDEC Standard No. 79-3F** accurately describes the JEDEC standard for DDR3 SDRAM.

217. NVIDIA references the **JEDEC Standard No. 79-3F** in the design of its products.

218. NVIDIA relies, or relied, on the **JEDEC Standard No. 79-3F** in the design of its products.

219. References to the document “**JEDEC Standard No. 79-2F**” refer to the document contemporaneously produced as SAMS-NVD-0002273 to 2400, entitled “JEDEC DDR2 Standard JESD79-2F,” JEDEC (November 2009).

220. The **JEDEC Standard No. 79-2F** is a true and correct copy of a document published by JEDEC in or around November 2009.

221. The **JEDEC Standard No. 79-2F** accurately describes the JEDEC standard for DDR3 SDRAM.

222. NVIDIA references the **JEDEC Standard No. 79-2F** in the design of its products.

223. NVIDIA relies, or relied, on the **JEDEC Standard No. 79-2F** in the design of its products.

224. References to the document “**Micron Technical Note TN-41-01**” refer to the document contemporaneously produced as SAMS-NVD-0002627 to 2650, entitled “Micron Technical Note TN-41-01: Calculating Memory System Power for DDR3.”

225. The **Micron Technical Note TN-41-01** is a true and correct copy of a document published by Micron.

226. The **Micron Technical Note TN-41-01** accurately describes the specifications of Micron DDR3 SDRAM.

227. References to the document “**DDR2 Datasheet**” refer to the document contemporaneously produced as SAMS-NVD-0002138 to 2270, entitled “Micron DDR2 SDRAM Datasheet.”

228. The **DDR2 Datasheet** is a true and correct copy of a document published by Micron on an unknown date.

229. The **DDR2 Datasheet** accurately describes the specifications of Micron DDR2 SDRAM.

230. References to the document “**NVIDIA 2014 Form 10-K**” refer to the document contemporaneously produced as SAMS-NVD-0001619 to 1734, entitled “Form 10-K,” NVIDIA (2014).

231. The **NVIDIA 2014 Form 10-K** is a true and correct copy of a document published by NVIDIA on or about March 13, 2014.

232. The **NVIDIA 2014 Form 10-K** accurately describes NVIDIA’s business.

233. References to the document “**NVIDIA PartnerForce Info**” refer to the document contemporaneously produced as SAMS-NVD-0001736 to 1737, entitled “NVIDIA PartnerForce Info.”

234. **NVIDIA PartnerForce Info** is a true and accurate copy of a document authored by NVIDIA and available on NVIDIA’s website at http://www.nvidia.com/object/pf_boardpartners.html.

235. **NVIDIA PartnerForce Info** accurately describes certain aspects of NVIDIA’s business and is intended to be read and relied upon by NVIDIA’s partners.

236. References to the document “**NVIDIA PartnerForce Program**” refer to the document contemporaneously produced as SAMS-NVD-0001738 to 1739, entitled “NVIDIA PartnerForce Program.”

237. **NVIDIA PartnerForce Program** is a true and accurate copy of a document authored by NVIDIA and available on NVIDIA’s website at <http://www.nvidia.com/page/channel.html>.

238. **NVIDIA PartnerForce Program** accurately describes certain aspects of NVIDIA's business and is intended to be read and relied upon by NVIDIA's partners.

239. References to the document "**Velocity Company History**" refer to the document contemporaneously produced as SAMS-NVD-0001743, entitled "From Humble Beginnings to Critical Acclaim."

240. **Velocity Company History** is a true and accurate copy of a document authored by Velocity and available on Velocity's website at <http://www.velocitymicro.com/velocity-story.php>.

241. **Velocity Company History** accurately describes certain aspects of the history of Velocity's business.

242. References to the document "**Traveling the Road to Silicon Motown**" refer to the document contemporaneously produced as SAMS-NVD-0001740 to 1742, entitled "Traveling the Road to Silicon Motown," NVIDIA (Feb. 2, 2012).

243. **Traveling The Road To Silicon Motown** is a true and accurate copy of a document authored by NVIDIA and available on NVIDIA's website at <http://blogs.nvidia.com/blog/2012/02/02/traveling-the-road-to-silicon-motown/>.

244. **Traveling The Road To Silicon Motown** accurately describes certain aspects of NVIDIA's business.

245. References to the document "**NVIDIA Multiple Display Instructions**" refer to the document contemporaneously produced as SAMS-NVD-0001735, containing an NVIDIA webpage entitled "How to setup multiple displays," NVIDIA (2014).

246. The **NVIDIA Multiple Display Instructions** is a true and correct copy of a webpage published by NVIDIA.

247. The **NVIDIA Multiple Display Instructions** provide instructions regarding how to use multiple displays with **'724 Accused Mobile GPUs**.

CLAIMS FOR RELIEF

248. The allegations in the following Claims For Relief have evidentiary support or will likely have evidentiary support after a reasonable opportunity for further investigation or discovery. Samsung does not yet have the benefit of any discovery from NVIDIA or Velocity or third parties.

249. The court has not construed the meaning of any claims or terms in the **Asserted Patents**. In providing these detailed allegations, Samsung does not intend to convey or imply any particular claim constructions or the precise scope of the claims. Samsung's claim construction contentions regarding the full meaning and scope of the claim terms will be provided in compliance with the case schedule and any applicable orders.

250. Samsung contends that each element of each asserted claim of the **Asserted Patents** is literally present in the **Accused Products**, processes used to make the **Accused Products**, use of the **Accused Products**, or systems containing the **Accused Products**. If the Court's constructions or other determinations indicate that an element of an asserted claim of the **Asserted Patents** is not literally present, Samsung contends that each such element is present under the doctrine of equivalents. If necessary, Samsung will provide more detailed doctrine of equivalents contentions after discovery from NVIDIA and Velocity and third parties or a claim construction order by the court.

FIRST CLAIM FOR RELIEF
INFRINGEMENT OF U.S. PATENT NO. 5,860,158
(AGAINST NVIDIA AND VELOCITY)

251. Each of the above listed paragraphs is incorporated herein by reference and adopted, as if fully set forth again.

252. The **'158 Patent** was filed on November 15, 1996, issued on January 12, 1999, and is entitled "Cache Control Unit With A Cache Request Transaction-Oriented Protocol." The **'158 Patent** is generally directed to a cache control unit and a method of controlling cache.

253. The **'158 Patent** was assigned to SEC, and SEC continues to hold all rights, title, and interest in the **'158 Patent**. A true and correct copy of the **'158 Patent** is attached hereto as Exhibit A.

254. All **Accused SOCs** infringe the **'158 Patent** in substantially the same way.

255. All **Accused SOCs** are based on designs provided by ARM.

256. For example, the **Cortex-A9 SOCs** are based on the ARM Cortex-A9 design.

257. The **Cortex-A15 SOCs** are based on the ARM Cortex-A15 design.

258. Each **Cortex-A9 SOC** includes one or more ARM PL310 cache controllers.

259. The PL310 cache controllers in the **Cortex-A9 SOCs** provide L2 cache control functionality in the **Cortex-A9 SOCs**.

260. The PL310 cache controllers in the various **Cortex-A9 SOCs** function in substantially the same manner.

261. Each **Cortex-A15 SOCs** includes one or more L2 cache controllers.

262. The L2 cache controllers in the **Cortex-A15 SOCs** are also based on designs provided by ARM.

263. The L2 cache controllers in the **Cortex-A15 SOCs** are integrated with the CPU complex.

264. The L2 cache controllers in the **Cortex-A15 SOCs** perform the same general function—L2 cache control—as the PL310 cache controllers in the **Cortex-A9 SOCs**.

265. The L2 cache controllers in the **Cortex-A15 SOCs** meet the claim elements of the **'158 Patent** in substantially the same manner as the PL310 cache controllers in the **Cortex-A9 SOCs**.

266. All **Accused SOCs** also include one or more L1 cache controllers.

267. These cache controllers are also based on designs provided by ARM.

268. The L1 cache controllers in the **Accused SOCs** are integrated with the CPU complex.

269. The L1 cache controllers in the **Accused SOCs** interoperate with the L2 cache controllers in those SOCs.

270. The L1 cache controllers meet the claim elements of the **'158 Patent** in substantially the same manner as the PL310 cache controllers.

271. All **Unified Cache GPUs** infringe the **'158 Patent** in substantially the same way.

272. NVIDIA GPUs using the Fermi architecture include a unified L2 cache.

273. For example, the NVIDIA GF100 GPU has a 768 KB unified L2 cache that services all load, store, and texture requests. **NVIDIA GF100 Whitepaper** at 20.

274. The GF100's unified L2 cache dynamically load balances between different requests, allowing full utilization of the cache. **NVIDIA GF100 Whitepaper** at 20.

275. The **NVIDIA GF100 Whitepaper** depicts the L2 cache in the following block diagram:



GF100 block diagram showing the Host Interface, the GigaThread Engine, four GPCs, six Memory Controllers, six ROP partitions, and a 768 KB L2 cache. Each GPC contains four PolyMorph engines. The ROP partitions are immediately adjacent to the L2 cache.

NVIDIA GF100 Whitepaper at 11.

276. NVIDIA GPUs using the Kepler architecture include a unified L2 cache.

277. The unified L2 caches in the NVIDIA GPUs using the Kepler architecture are substantially similar to the L2 caches in GPUs using the Fermi architecture. **NVIDIA Kepler**

GK110 Whitepaper at 13.

278. The **NVIDIA Kepler GK110 Whitepaper** depicts the L2 cache in the following block diagram:



Kepler GK110 Full chip block diagram

NVIDIA Kepler GK110 Whitepaper at 6.

279. NVIDIA GPUs using the Maxwell architecture include a unified L2 cache.

280. The unified L2 caches in GPUs using the Maxwell architecture are substantially similar to the L2 caches in GPUs using the Fermi architecture.

281. The **NVIDIA GeForce GTX 750 Ti Whitepaper** depicts the L2 cache in the following block diagram:

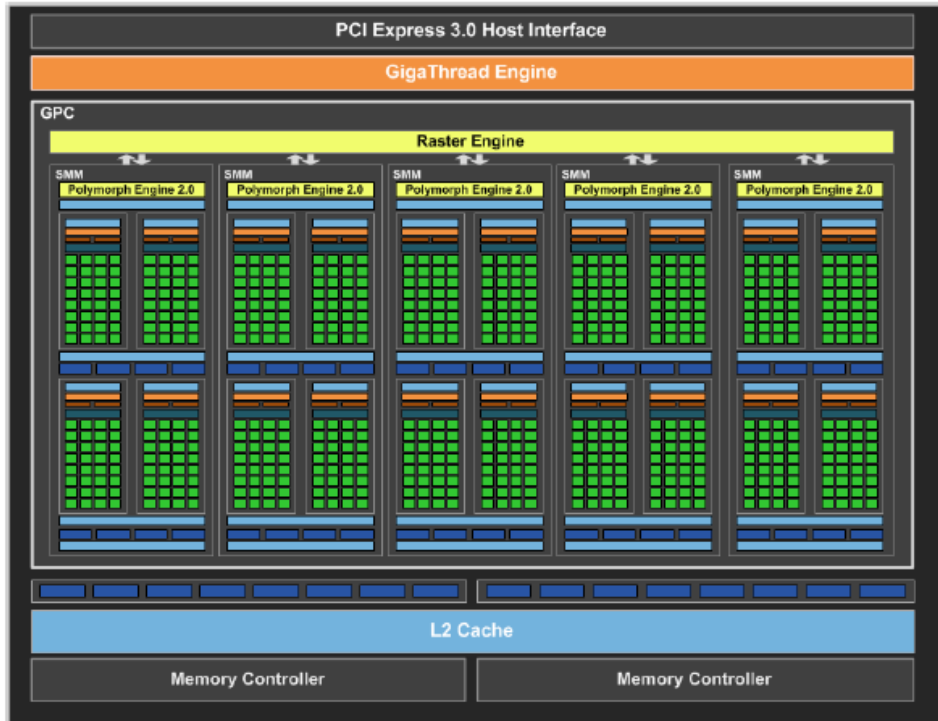


Figure 2: GM107 Full-Chip Block Diagram

NVIDIA GeForce GTX 750 Ti Whitepaper at 5.

282. “**Claim 1(a) of the ’158 Patent**” recites “[a] method for controlling a cache, the cache being coupled to a device.”

283. All **’158 Accused Products** practice a method for controlling a cache, the cache being coupled to a device.

284. For example, the **Cortex-A9 SOCs** include one or more PL310 cache controllers. One or more of these PL310 cache controllers are used for L2 cache control.

285. The PL310 cache controllers used for L2 cache control are coupled to cache memory.

286. The PL310 cache controllers used for L2 cache control are also coupled to one or more devices. For example, the PL310 cache controllers are coupled to the CPU cores, L1 caches, and snoop control units in the **Cortex-A9 SOCs**.

287. The cache memory is coupled to the devices via the PL310 cache controllers used for L2 cache control.

288. The **Cortex-A15 SOCs** include one or more L2 cache controllers. One or more of these L2 cache controllers are integrated or otherwise associated with the CPU complex.

289. The L2 cache controllers (referenced in the immediately preceding paragraph) are coupled to cache memory. These cache controllers are also coupled to one or more devices. For example, the L2 cache controllers are coupled to the CPU cores and L1 caches in the **Cortex-A15 SOCs**. The cache memory is coupled to the CPU cores via the L2 cache controllers.

290. All **Accused SOCs** include one or more L1 cache controllers. One or more of these L1 cache controllers are integrated with the CPU complex.

291. The L1 cache controllers (referenced in the immediately preceding paragraph) are coupled to cache memory. These cache controllers are each coupled to a CPU core. The cache memory is coupled to the CPU cores via the L1 cache controllers.

292. All **Unified Cache GPUs** include one or more L2 cache controllers. One or more of these L2 cache controllers are coupled to cache memory.

293. The L2 cache controllers (referenced in the immediately preceding paragraph) are also coupled to one or more upstream devices. The cache memory is coupled to the one or more upstream devices via the L2 cache controllers.

294. For example, the NVIDIA GF100 GPU includes 768 KB of L2 cache memory. **NVIDIA GF100 Whitepaper** at 11. This memory is coupled to the L2 cache controller.

295. The **NVIDIA GF100 Whitepaper** depicts the L2 cache in the following block diagram.



GF100 block diagram showing the Host Interface, the GigaThread Engine, four GPCs, six Memory Controllers, six ROP partitions, and a 768 KB L2 cache. Each GPC contains four PolyMorph engines. The ROP partitions are immediately adjacent to the L2 cache.

NVIDIA GF100 Whitepaper at 11.

296. As described in the preceding paragraphs, the **'158 Accused Products** practice a method for controlling a cache, the cache being coupled to a device.

297. **“Claim 1(b) of the '158 Patent”** recites “receiving by a cache controller a first cache request from the device.”

298. All **'158 Accused Products** receive by a cache controller a first cache request from the device.

299. For example, the **Accused SOCs** include cache controllers that receive cache requests from devices to which they are coupled.

300. The PL310 cache controllers in the **Cortex-A9 SOCs** receive cache requests from devices to which they are coupled, such as the CPU cores. These cache requests may include read requests and write requests.

301. The L2 cache controllers in the **Cortex-A15 SOCs** receive cache requests from devices to which they are coupled, such as the CPU cores. These cache requests may include read requests and write requests.

302. The L1 cache controllers in the **Accused SOCs** receive cache requests from the CPU cores to which they are coupled. These cache requests may include read requests and write requests.

303. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

304. As described in the preceding paragraphs, **Claim 1(b) of the '158 Patent** is met when an **Accused SOC** is used.

305. The **Unified Cache GPUs** practice **Claim 1(b) of the '158 Patent**.

306. For example, the L2 cache controllers in the **Unified Cache GPUs** receive cache requests from one or more of the upstream devices to which they are coupled.

307. The L2 cache controller in the NVIDIA GF100 GPU services all load, store, and texture requests. **NVIDIA GF100 Whitepaper** at 20.

308. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

309. As described in the preceding paragraphs, **Claim 1(b) of the '158 Patent** is met when a **Unified Cache GPU** is used.

310. As described in the preceding paragraphs, **Claim 1(b) of the '158 Patent** is met when a **'158 Accused Product** is used

311. “**Claim 1(c) of the ’158 Patent**” recites “providing by the cache controller a first request ID value corresponding to the first cache request to the device after receiving the first cache request.”

312. All **’158 Accused Products** provide by the cache controller a first request ID value corresponding to the first cache request to the device after receiving the first cache request.

313. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** provide request ID values corresponding to each cache request to the requesting device.

314. If a cache request is a read request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Read ID tag” or “RID.” **AXI Spec** at 35. This signal is the identification tag for the read data group of signals. *Id.*

315. If a cache request is a write request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Response ID tag” or “BID.” **AXI Spec** at 33. This signal is the ID tag of the of the write response. *Id.*

316. The request ID signals RID and BID are provided by the PL310 cache controller after receiving the corresponding request.

317. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 1(c) of the ’158 Patent** in a similar manner to that described above. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on designs provided by ARM. In addition, the L1 and L2 cache systems in the **Cortex-A9 SOCs** interoperate. The L1 cache controllers meet **Claim 1(c) of the ’158 Patent** in substantially the same manner as the PL310 cache controllers.

318. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 1(c) of the ’158 Patent** in a similar manner to that described above. The L1 and L2 cache controllers in

the **Cortex-A15 SOC**s are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOC**s perform a similar role as the L2 cache controllers in the **Cortex-A15 SOC**s. The L1 and L2 cache controllers in the **Cortex-A15 SOC**s meet **Claim 1(c) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

319. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

320. For the reasons described above, request ID values corresponding to each cache request will be provided.

321. As described in the preceding paragraphs, **Claim 1(c) of the '158 Patent** is met when an **Accused SOC** is used.

322. The **Unified Cache GPU**s practice **Claim 1(c) of the '158 Patent**.

323. For example, the L2 cache controllers in the **Unified Cache GPU**s provide request ID values corresponding to each cache request to the requesting device.

324. Such request ID values are provided after the L2 cache controller receives the cache request.

325. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

326. Request ID values corresponding to each cache request will be provided.

327. As described in the preceding paragraphs, **Claim 1(c) of the '158 Patent** is met when a **Unified Cache GPU** is used.

328. As described in the preceding paragraphs, **Claim 1(c) of the '158 Patent** is met when a **'158 Accused Product** is used.

329. “**Claim 1(d) of the ’158 Patent**” recites “initiating processing of the first cache request after receiving the first cache request.”

330. All **’158 Accused Products** initiate processing of the first cache request after receiving the first cache request.

331. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving the cache request.

332. If the first cache request is a read request, the PL310 cache controller will process that request after receiving it. If the first cache request is a write request, the PL310 cache controller will process that request after receiving it.

333. For example, in the timing diagram depicted below, the PL310 cache controller receives a read request for address A at time T1. The processing of that request occurs after time T1.

C.4 Outstanding read hit transaction

Figure C-4 shows the timing for an outstanding read hit transaction.

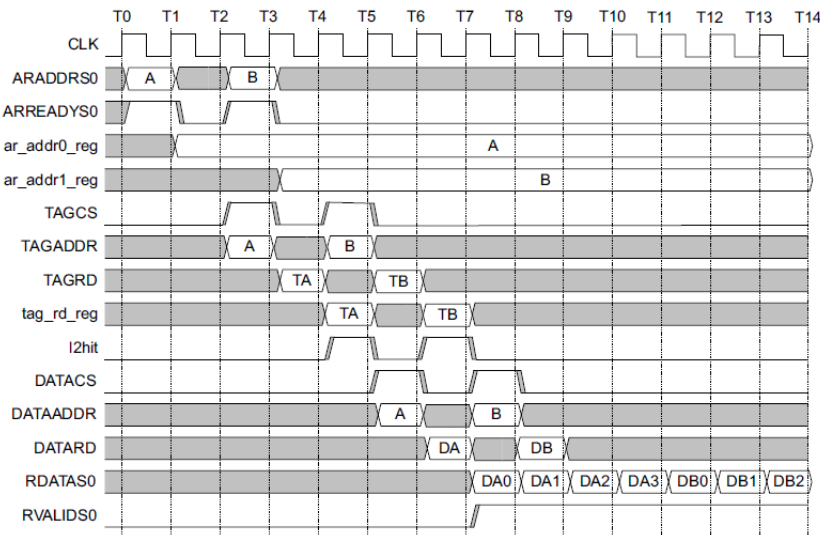


Figure C-4 Outstanding read hit transaction

PL310 TRM at Fig. C-4.

334. In the timing diagram depicted above, processing of the first cache request is initiated after the first cache request is received.

335. The L1 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving that cache request.

336. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** process each request after receiving that request.

337. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems. Each cache request will be processed by the cache controller after the cache controller receives the request.

338. As described in the preceding paragraphs, **Claim 1(d) of the '158 Patent** is met when an **Accused SOC** is used.

339. The **Unified Cache GPUs** practice **Claim 1(d) of the '158 Patent**.

340. For example, the L2 cache controllers in the **Unified Cache GPUs** process each cache request after receiving the cache request.

341. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

342. Each cache request will be processed by the cache controller after the cache controller receives the request.

343. As described in the preceding paragraphs, **Claim 1(d) of the '158 Patent** is met when a **Unified Cache GPU** is used.

344. As described in the preceding paragraphs, **Claim 1(d) of the '158 Patent** is met when a **'158 Accused Product** is used.

345. “**Claim 1(e) of the ’158 Patent**” recites “receiving by the cache controller a second cache request from the device after receiving the first cache request.”

346. All **’158 Accused Products** receive by the cache controller a second cache request from the device after receiving the first cache request.

347. For example, cache controllers in the **Cortex-A9 SOCs** can accept multiple outstanding read and write requests.

348. The PL310 cache controllers can accept up to 16 read requests and up to 8 write requests. **PL310 TRM** at 2-5.

349. The data cache in the L1 cache in the **Cortex-A9 SOCs** can accept up to four outstanding data cache read misses and up to four outstanding data cache write misses. **Cortex-A9 TRM** at 7-3.

350. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** can each also accept multiple outstanding read and write requests. **Cortex-A15 TRM** at 6-6 and 7-12.

351. Multiple cache requests likely occur in any reasonable usage scenario.

352. As described in the preceding paragraphs, **Claim 1(e) of the ’158 Patent** is met when an **Accused SOC** is used.

353. The **Unified Cache GPUs** practice **Claim 1(e) of the ’158 Patent**.

354. The **Unified Cache GPUs** can accept multiple outstanding cache requests.

355. For example, the NVIDIA GF100 GPU’s unified L2 cache dynamically load balances between different requests. **NVIDIA GF100 Whitepaper** at 20.

356. Multiple cache requests likely occur in any reasonable usage scenario.

357. As described in the preceding paragraphs, **Claim 1(e) of the ’158 Patent** is met when a **Unified Cache GPU** is used.

358. As described in the preceding paragraphs, **Claim 1(e) of the '158 Patent** is met when a **'158 Accused Product** is used.

359. “**Claim 1(f) of the '158 Patent**” recites “providing by the cache controller a second request ID value corresponding to the second cache request to the device after receiving the second cache request.”

360. All **'158 Accused Products** provide by the cache controller a second request ID value corresponding to the second cache request to the device after receiving the second cache request.

361. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** provide request ID values corresponding to each cache request to the requesting device.

362. If a cache request is a read request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Read ID tag” or “RID.” **AXI Spec** at 35. This signal is the identification tag for the read data group of signals. *Id.*

363. If a cache request is a write request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Response ID tag” or “BID.” **AXI Spec** at 33. This signal is the ID tag of the of the write response. *Id.*

364. The request ID signals RID and BID are provided by the PL310 cache controller after receiving the corresponding request.

365. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 1(f) of the '158 Patent** in a similar manner to that described above. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on designs provided by ARM. In addition, the L1 and L2 cache systems in

the **Cortex-A9 SOCs** interoperate. The L1 cache controllers meet **Claim 1(f) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

366. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 1(f) of the '158 Patent** in a similar manner to that described above. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 1(f) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

367. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

368. For the reasons described above, request ID values corresponding to each cache request will be provided.

369. Multiple cache requests likely occur in any reasonable usage scenario.

370. As described in the preceding paragraphs, **Claim 1(f) of the '158 Patent** is met when an **Accused SOC** is used.

371. The **Unified Cache GPUs** practice **Claim 1(f) of the '158 Patent**.

372. The L2 cache controllers in the **Unified Cache GPUs** provide request ID values corresponding to each cache request to the requesting device.

373. Such request ID values are provided after the L2 cache controller receives the cache request.

374. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

375. Request ID values corresponding to each cache request will be provided.

376. Multiple cache requests likely occur in any reasonable usage scenario.

377. As described in the preceding paragraphs, **Claim 1(f) of the '158 Patent** is met when a **Unified Cache GPU** is used.

378. As described in the preceding paragraphs, **Claim 1(f) of the '158 Patent** is met when a **'158 Accused Product** is used.

379. “**Claim 1(g) of the '158 Patent**” recites “initiating processing of the second cache request after receiving the second cache request.”

380. All **'158 Accused Products** initiate processing of the second cache request after receiving the second cache request.

381. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving the cache request.

382. If the second cache request is a read request, the PL310 cache controller will process that request after receiving it. If the second cache request is a write request, the PL310 cache controller will process that request after receiving it.

383. For example, in the timing diagram depicted below, the PL310 cache controller receives the second read request, for address B, at time T3. The processing of that request occurs after time T3.

C.4 Outstanding read hit transaction

Figure C-4 shows the timing for an outstanding read hit transaction.

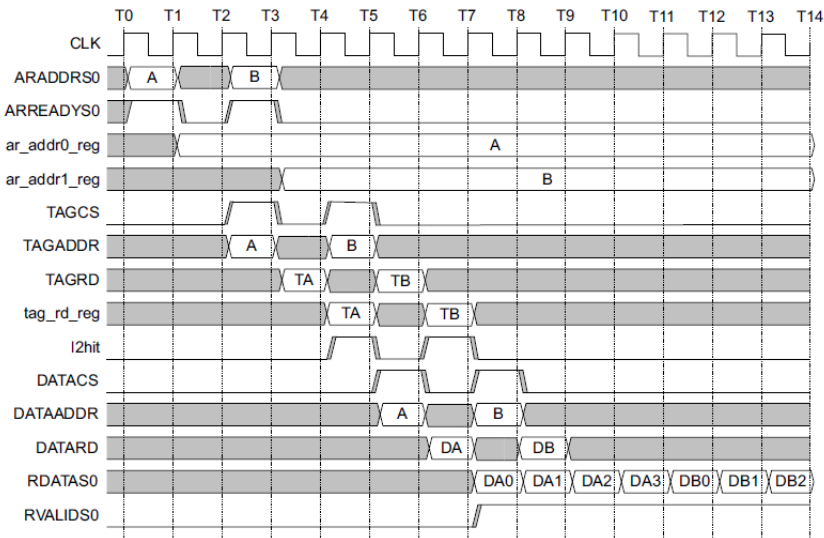


Figure C-4 Outstanding read hit transaction

PL310 TRM at Fig. C-4.

384. In the timing diagram depicted above, processing of the second cache request is initiated after receiving the second cache request.

385. The L1 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving that cache request.

386. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** process each request after receiving that request.

387. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems. Each cache request will be processed by the cache controller after the cache controller receives the request. Multiple cache requests occur in any reasonable usage scenario.

388. As described in the preceding paragraphs, **Claim 1(g) of the '158 Patent** is met when an **Accused SOC** is used.

389. The **Unified Cache GPUs** practice **Claim 1(g) of the '158 Patent**.

390. For example, the L2 cache controllers in the **Unified Cache GPUs** process each cache request after receiving the cache request.

391. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

392. Each cache request will be processed by the cache controller after the cache controller receives the request.

393. Multiple cache requests likely occur in any reasonable usage scenario.

394. As described in the preceding paragraphs, **Claim 1(g) of the '158 Patent** is met when a **Unified Cache GPU** is used.

395. As described in the preceding paragraphs, **Claim 1(g) of the '158 Patent** is met when a **'158 Accused Product** is used.

396. "**Claim 1(h) of the '158 Patent**" recites "completing the processing of the first and second cache requests after receiving the second cache request."

397. All **'158 Accused Products** complete processing of the first and second cache requests after receiving the second cache request.

398. Cache controllers in the **Cortex-A9 SOCs** can accept multiple outstanding read and write requests.

399. For example, the PL310 cache controllers can accept up to 16 read requests and up to 8 write requests. **PL310 TRM** at 2-5.

400. The data cache in the L1 cache in the **Cortex-A9 SOCs** can accept up to four outstanding data cache read misses and up to four outstanding data cache write misses. **Cortex-A9 TRM** at 7-3.

401. Cache controllers in the **Cortex-A9 SOCs** can accept two cache requests and complete processing of both cache requests after receiving the second cache request.

402. For example, the following timing diagram depicts the timing for an outstanding read hit transaction using a PL310 cache controller. The second cache request is received at time T3. The first cache request completes processing at time T11. The second cache request completes processing at time T14.

C.4 Outstanding read hit transaction

Figure C-4 shows the timing for an outstanding read hit transaction.

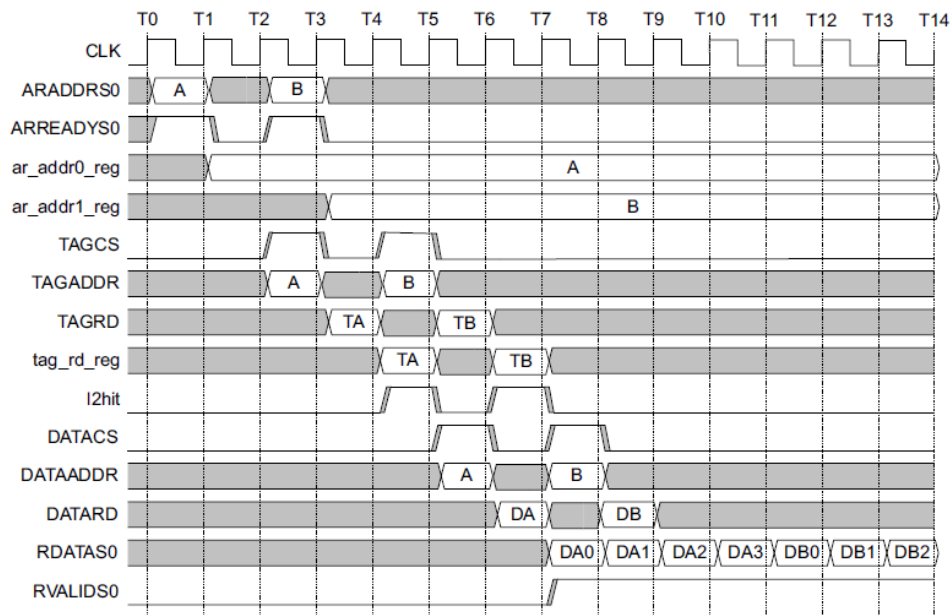


Figure C-4 Outstanding read hit transaction

PL310 TRM at Fig. C-4.

403. In the timing diagram depicted above, processing of the first and second cache requests is completed after receiving the second cache request.

404. The following timing diagram depicts the timing for a hit under miss read transaction using a PL310 cache controller. The second cache request is received at time T3.

The first cache request completes processing at time T16. The second cache request completes processing at time T13.

C.5 Hit under miss read transactions

Figure C-5 shows the timing for a hit under miss read transaction case.




Figure C-5 Hit under miss read transaction

PL310 TRM at Fig. C-5.

405. In the timing diagram depicted above, processing of the first and second cache requests is completed after receiving the second cache request.

406. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 1(h) of the '158 Patent** in a similar manner to that described above.

407. The data cache of the L1 cache controller is non-blocking, with up to four outstanding data cache read misses and up to four outstanding data cache write misses. **Cortex-A9 TRM** at 7-3.

408. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on designs provided by ARM. In addition, the L1 and L2 cache systems in the **Cortex-A9 SOCs** interoperate. The L1 cache controllers meet **Claim 1(h) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

409. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** can each also accept multiple outstanding read and write requests. **Cortex-A15 TRM** at 6-6 and 7-12.

410. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 1(h) of the '158 Patent** in a similar manner to that described above. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 1(h) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

411. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

412. Each cache request will be processed by the cache controller after the cache controller receives the request.

413. Multiple outstanding cache requests likely occur in any reasonable usage scenario.

414. As described in the preceding paragraphs, **Claim 1(h) of the '158 Patent** is met when an **Accused SOC** is used.

415. The **Unified Cache GPUs** practice **Claim 1(h) of the '158 Patent**.

416. The **Unified Cache GPUs** can accept multiple outstanding cache requests.

417. For example, the NVIDIA GF100 GPU's unified L2 cache dynamically load balances between different requests. **NVIDIA GF100 Whitepaper** at 20.

418. Cache controllers in the **Unified Cache GPUs** can accept two cache requests and complete processing of both cache requests after receiving the second cache request.

419. Cache requests are utilized during ordinary usage of GPUs that utilize L2 cache systems.

420. Each cache request will be processed by the cache controller after the cache controller receives the request.

421. Multiple outstanding cache requests occur in any reasonable usage scenario.

422. As described in the preceding paragraphs, **Claim 1(h) of the '158 Patent** is met when a **Unified Cache GPU** is used.

423. As described in the preceding paragraphs, **Claim 1(h) of the '158 Patent** is met when a **'158 Accused Product** is used.

424. "**Claim 15(a) of the '158 Patent**" recites "[a] method for controlling a cache, the cache being coupled to at least one device."

425. All **'158 Accused Products** practice a method for controlling a cache, the cache being coupled to at least one device.

426. For example, the **Cortex-A9 SOCs** include one or more PL310 cache controllers. One or more of these PL310 cache controllers are used for L2 cache control.

427. The PL310 cache controllers used for L2 cache control are coupled to cache memory.

428. The PL310 cache controllers used for L2 cache control are also coupled to one or more devices. For example, the PL310 cache controllers are coupled to the CPU cores, L1 caches, and snoop control units in the **Cortex-A9 SOCs**.

429. The cache memory is coupled to the devices via the PL310 cache controllers used for L2 cache control.

430. The **Cortex-A15 SOCs** include one or more L2 cache controllers. One or more of these L2 cache controllers are integrated with or otherwise associated with the CPU complex.

431. The L2 cache controllers (referenced in the immediately preceding paragraph) are coupled to cache memory. These cache controllers are also coupled to one or more devices. For example, the L2 cache controllers are coupled to the CPU cores and L1 caches in the **Cortex-A15 SOCs**. The cache memory is coupled to the devices via the L2 cache controllers.

432. All **Accused SOCs** include one or more L1 cache controllers. One or more of these L1 cache controllers are integrated with the CPU complex.

433. These L1 cache controllers (referenced in the immediately preceding paragraph) are coupled to cache memory. These cache controllers are each coupled to a CPU core. The cache memory is coupled to the CPU cores via the L1 cache controllers.

434. All **Unified Cache GPUs** include one or more L2 cache controllers. One or more of these L2 cache controllers are coupled to cache memory.

435. The L2 cache controllers (referenced in the immediately preceding paragraph) are also coupled to one or more upstream devices. The cache memory is coupled to the one or more upstream devices via the L2 cache controllers.

436. For example, the the NVIDIA GF100 GPU includes 768 KB of L2 cache memory. **NVIDIA GF100 Whitepaper** at 11. This memory is coupled to the L2 cache controller.

437. The **NVIDIA GF100 Whitepaper** depicts the L2 cache in the following block diagram.



GF100 block diagram showing the Host Interface, the GigaThread Engine, four GPCs, six Memory Controllers, six ROP partitions, and a 768 KB L2 cache. Each GPC contains four PolyMorph engines. The ROP partitions are immediately adjacent to the L2 cache.

NVIDIA GF100 Whitepaper at 11.

438. As described in the preceding paragraphs, the **'158 Accused Products** practice a method for controlling a cache, the cache being coupled to at least one device

439. **“Claim 15(b) of the '158 Patent”** recites “receiving by a cache controller a first cache request from a first device of the at least one device.”

440. All **'158 Accused Products** receive by a cache controller a first cache request from a first device of the at least one device.

441. For example, the **Accused SOCs** include cache controllers that receive cache requests from devices to which they are coupled.

442. The PL310 cache controllers in the **Cortex-A9 SOCs** receive cache requests from devices to which they are coupled, such as the CPU cores. These cache requests may include read requests and write requests.

443. The L2 cache controllers in the **Cortex-A15 SOCs** receive cache requests from devices to which they are coupled, such as the CPU cores. These cache requests may include read requests and write requests.

444. The L1 cache controllers in the **Accused SOCs** receive cache requests from the CPU cores to which they are coupled. These cache requests may include read requests and write requests.

445. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

446. As described in the preceding paragraphs, **Claim 15(b) of the '158 Patent** is met when an **Accused SOC** is used.

447. The **Unified Cache GPUs** practice **Claim 15(b) of the '158 Patent**.

448. The L2 cache controllers in the **Unified Cache GPUs** receive cache requests from one or more of the upstream devices to which they are coupled.

449. For example, the L2 cache controller in the NVIDIA GF100 GPU services all load, store, and texture requests. **NVIDIA GF100 Whitepaper** at 20.

450. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

451. As described in the preceding paragraphs, **Claim 15(b) of the '158 Patent** is met when a **Unified Cache GPU** is used.

452. As described in the preceding paragraphs, **Claim 15(b) of the '158 Patent** is met when a **'158 Accused Product** is used.

453. “**Claim 15(c) of the '158 Patent**” recites “providing by the cache controller a first request ID value corresponding to the first cache request to the first device after receiving the first cache request.”

454. All **'158 Accused Products** provide by the cache controller a first request ID value corresponding to the first cache request to the first device after receiving the first cache request.

455. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** provide request ID values corresponding to each cache request to the requesting device.

456. If a cache request is a read request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Read ID tag” or “RID.” **AXI Spec** at 35. This signal is the identification tag for the read data group of signals. *Id.*

457. If a cache request is a write request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Response ID tag” or “BID.” **AXI Spec** at 33. This signal is the ID tag of the of the write response. *Id.*

458. The request ID signals RID and BID are provided by the PL310 cache controller after receiving the corresponding request.

459. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 15(c) of the '158 Patent** in a similar manner to that described above. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the

Cortex-A9 SOCs are also based on designs provided by ARM. In addition, the L1 and L2 cache systems in the **Cortex-A9 SOC**s interoperate. The L1 cache controllers meet **Claim 15(c) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

460. The L1 and L2 cache controllers in the **Cortex-A15 SOC**s practice **Claim 15(c) of the '158 Patent** in a similar manner to that described above. The L1 and L2 cache controllers in the **Cortex-A15 SOC**s are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOC**s perform a similar role as the L2 cache controllers in the **Cortex-A15 SOC**s. The L1 and L2 cache controllers in the **Cortex-A15 SOC**s meet **Claim 15(c) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

461. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

462. For the reasons described above, request ID values corresponding to each cache request will be provided.

463. As described in the preceding paragraphs, **Claim 15(c) of the '158 Patent** is met when an **Accused SOC** is used.

464. The **Unified Cache GPU**s practice **Claim 15(c) of the '158 Patent**.

465. For example, the L2 cache controllers in the **Unified Cache GPU**s provide request ID values corresponding to each cache request to the requesting device. Such request ID values are provided after the L2 cache controller receives the cache request.

466. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

467. Request ID values corresponding to each cache request will be provided.

468. As described in the preceding paragraphs, **Claim 15(c) of the '158 Patent** is met when a **Unified Cache GPU** is used.

469. As described in the preceding paragraphs, **Claim 15(c) of the '158 Patent** is met when a **'158 Accused Product** is used.

470. “**Claim 15(d) of the '158 Patent**” recites “initiating processing of the first cache request after receiving the first cache request.”

471. All **'158 Accused Products** initiate processing of the first cache request after receiving the first cache request.

472. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving the cache request.

473. If the first cache request is a read request, the PL310 cache controller will process that request after receiving it. If the first cache request is a write request, the PL310 cache controller will process that request after receiving it.

474. For example, in the timing diagram depicted below, the PL310 cache controller receives a read request for address A at time T1. The processing of that request occurs after time T1.

C.4 Outstanding read hit transaction

Figure C-4 shows the timing for an outstanding read hit transaction.

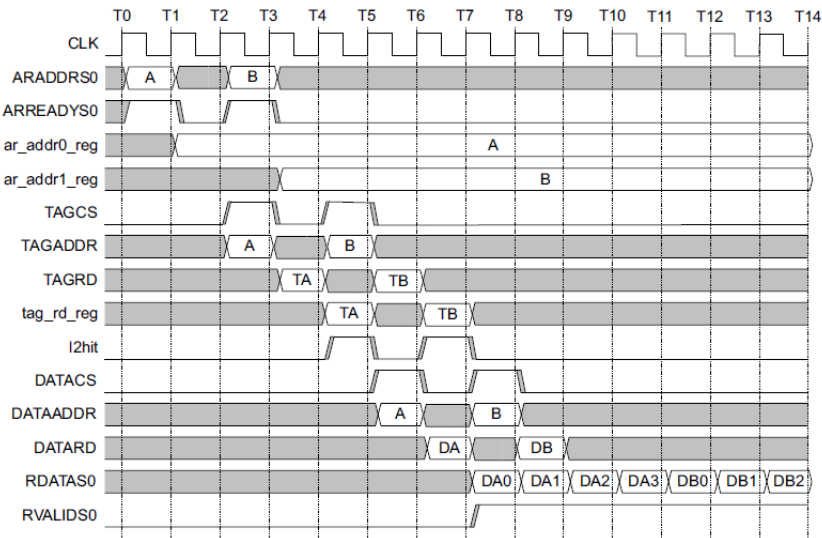


Figure C-4 Outstanding read hit transaction

PL310 TRM at Fig. C-4.

475. In the timing diagram depicted above, processing of the first cache request is initiated after receiving the first cache request.

476. The L1 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving that cache request.

477. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** process each request after receiving that request.

478. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

479. Each cache request will be processed by the cache controller after the cache controller receives the request.

480. As described in the preceding paragraphs, **Claim 15(d) of the '158 Patent** is met when an **Accused SOC** is used.

481. The **Unified Cache GPUs** practice **Claim 15(d) of the '158 Patent**.

482. For example, the L2 cache controllers in the **Unified Cache GPUs** process each cache request after receiving the cache request.

483. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

484. Each cache request will be processed by the cache controller after the cache controller receives the request.

485. As described in the preceding paragraphs, **Claim 15(d) of the '158 Patent** is met when a **Unified Cache GPU** is used.

486. As described in the preceding paragraphs, **Claim 15(d) of the '158 Patent** is met when a **'158 Accused Product** is used.

487. “**Claim 15(e) of the '158 Patent**” recites “receiving by the cache controller a second cache request from a second device of the at least one device after receiving the first cache request.”

488. All **'158 Accused Products** receive by the cache controller a second cache request from a second device of the at least one device after receiving the first cache request.

489. For example, cache controllers in the **Cortex-A9 SOCs** can accept multiple outstanding read and write requests.

490. The PL310 cache controllers can accept up to 16 read requests and up to 8 write requests. **PL310 TRM** at 2-5.

491. The data cache in the L1 cache in the **Cortex-A9 SOCs** can accept up to four outstanding data cache read misses and up to four outstanding data cache write misses. **Cortex-A9 TRM** at 7-3.

492. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** can each also accept multiple outstanding read and write requests. **Cortex-A15 TRM** at 6-6 and 7-12.

493. Multiple cache requests likely occur in any reasonable usage scenario.

494. Each of the **Accused SOCs** includes multiple CPU cores.

495. The second cache request can be received from the same device from which the first cache request was received, or from a different device. For example, if the first cache request is received from one CPU core, the second cache request may be received from the same CPU core, or from a different CPU core.

496. As described in the preceding paragraphs, **Claim 15(e) of the '158 Patent** is met when an **Accused SOC** is used.

497. The **Unified Cache GPUs** practice **Claim 15(e) of the '158 Patent**.

498. The **Unified Cache GPUs** can accept multiple outstanding cache requests.

499. For example, the NVIDIA GF100 GPU's unified L2 cache dynamically load balances between different requests. **NVIDIA GF100 Whitepaper** at 20.

500. Multiple cache requests likely occur in any reasonable usage scenario.

501. The second cache request can be received from the same device from which the first cache request was received, or from a different device.

502. As described in the preceding paragraphs, **Claim 15(e) of the '158 Patent** is met when a **Unified Cache GPU** is used.

503. As described in the preceding paragraphs, **Claim 15(e) of the '158 Patent** is met when a **'158 Accused Product** is used.

504. “**Claim 15(f) of the ’158 Patent**” recites “providing by the cache controller a second request ID value corresponding to the second cache request to the second device after receiving the second cache request.”

505. All **’158 Accused Products** provide by the cache controller a second request ID value corresponding to the second cache request to the second device after receiving the second cache request.

506. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** provide request ID values corresponding to each cache request to the requesting device.

507. If a cache request is a read request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Read ID tag” or “RID.” **AXI Spec** at 35. This signal is the identification tag for the read data group of signals. *Id.*

508. If a cache request is a write request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Response ID tag” or “BID.” **AXI Spec** at 33. This signal is the ID tag of the of the write response. *Id.*

509. The request ID signals RID and BID are provided by the PL310 cache controller after receiving the corresponding request.

510. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 15(f) of the ’158 Patent** in a similar manner to that described above. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on designs provided by ARM. In addition, the L1 and L2 cache systems in the **Cortex-A9 SOCs** interoperate. The L1 cache controllers meet **Claim 15(f) of the ’158 Patent** in substantially the same manner as the PL310 cache controllers.

511. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 15(f) of the '158 Patent** in a similar manner to that described above. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 15(f) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

512. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

513. For the reasons described above, request ID values corresponding to each cache request will be provided.

514. Multiple cache requests likely occur in any reasonable usage scenario.

515. As described in the preceding paragraphs, **Claim 15(f) of the '158 Patent** is met when an **Accused SOC** is used.

516. The **Unified Cache GPUs** practice **Claim 15(f) of the '158 Patent**.

517. For example, the L2 cache controllers in the **Unified Cache GPUs** provide request ID values corresponding to each cache request to the requesting device.

518. Such request ID values are provided after the L2 cache controller receives the cache request.

519. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

520. Request ID values corresponding to each cache request will be provided.

521. Multiple cache requests likely occur in any reasonable usage scenario.

522. As described in the preceding paragraphs, **Claim 15(f) of the '158 Patent** is met when a **Unified Cache GPU** is used.

523. As described in the preceding paragraphs, **Claim 15(f) of the '158 Patent** is met when a **'158 Accused Product** is used.

524. “**Claim 15(g) of the '158 Patent**” recites “initiating processing of the second cache request after receiving the second cache request.”

525. All **'158 Accused Products** initiate processing of the second cache request after receiving the second cache request.

526. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving the cache request.

527. If the second cache request is a read request, the PL310 cache controller will process that request after receiving it. If the second cache request is a write request, the PL310 cache controller will process that request after receiving it.

528. For example, in the timing diagram depicted below, the PL310 cache controller receives a second read request, for address B, at time T3. The processing of that request occurs after time T3.

C.4 Outstanding read hit transaction

Figure C-4 shows the timing for an outstanding read hit transaction.




Figure C-4 Outstanding read hit transaction

PL310 TRM at Fig. C-4.

529. In the timing diagram depicted above, processing of the second cache request is initiated after receiving the second cache request.

530. The L1 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving that cache request.

531. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** process each request after receiving that request.

532. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

533. Each cache request will be processed by the cache controller after the cache controller receives the request.

534. Multiple cache requests likely occur in any reasonable usage scenario.

535. As described in the preceding paragraphs, **Claim 15(g) of the '158 Patent** is met when an **Accused SOC** is used.

536. The **Unified Cache GPUs** practice **Claim 15(g) of the '158 Patent**.

537. For example, the L2 cache controllers in the **Unified Cache GPUs** process each cache request after receiving the cache request.

538. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

539. Each cache request will be processed by the cache controller after the cache controller receives the request.

540. Multiple cache requests likely occur in any reasonable usage scenario.

541. As described in the preceding paragraphs, **Claim 15(g) of the '158 Patent** is met when a **Unified Cache GPU** is used.

542. As described in the preceding paragraphs, **Claim 15(g) of the '158 Patent** is met when a **'158 Accused Product** is used.

543. “**Claim 15(h) of the '158 Patent**” recites “completing the processing of the first cache request after receiving the second cache request.”

544. All **'158 Accused Products** complete the processing of the first cache request after receiving the second cache request.

545. Cache controllers in the **Cortex-A9 SOCs** can accept multiple outstanding read and write requests.

546. For example, the PL310 cache controllers can accept up to 16 read requests and up to 8 write requests. **PL310 TRM** at 2-5.

547. The data cache in the L1 cache can accept up to four outstanding data cache read misses and up to four outstanding data cache write misses. **Cortex-A9 TRM** at 7-3.

548. Cache controllers in the **Cortex-A9 SOCs** can accept two cache requests and complete processing of the first cache request after receiving the second cache request. For example, the following timing diagram depicts the timing for an outstanding read hit transaction using a PL310 cache controller. The second cache request, for address B, is received at time T3. The first cache request completes processing at time T11.

C.4 Outstanding read hit transaction

Figure C-4 shows the timing for an outstanding read hit transaction.

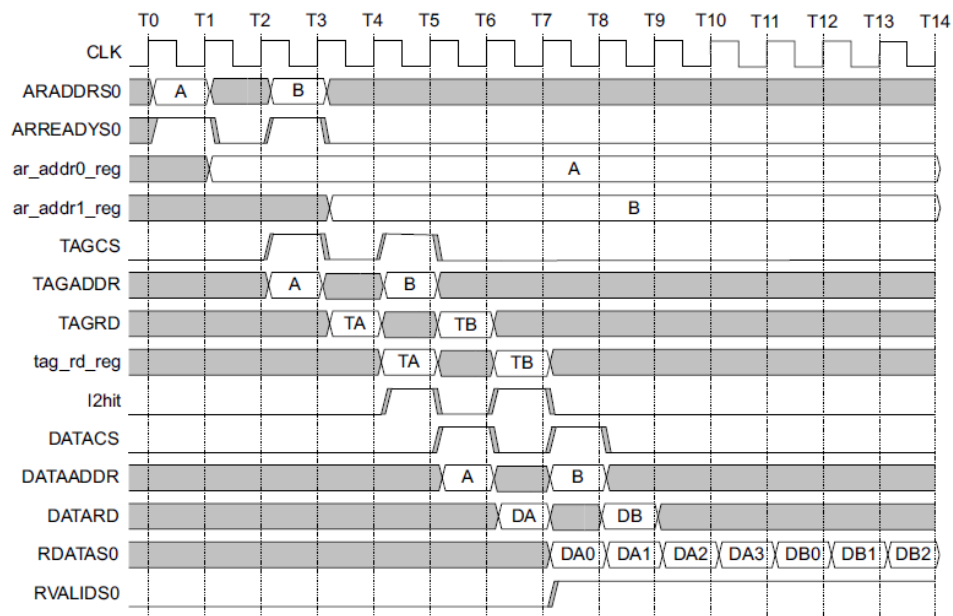


Figure C-4 Outstanding read hit transaction

PL310 TRM at Fig. C-4.

549. In the timing diagram depicted above, processing of the first cache request is completed after receiving the second cache request.

550. The following timing diagram depicts the timing for a hit under miss read transaction using a PL310 cache controller. The second cache request, for address B, is received at time T3. The first cache request completes processing at time T16.

C.5 Hit under miss read transactions

Figure C-5 shows the timing for a hit under miss read transaction case.




Figure C-5 Hit under miss read transaction

PL310 TRM at Fig. C-5.

551. In the timing diagram depicted above, processing of the first cache request is completed after receiving the second cache request.

552. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 15(h) of the '158 Patent** in a similar manner to that described above.

553. The data cache of the L1 cache controller is non-blocking, with up to four outstanding data cache read misses and up to four outstanding data cache write misses. **Cortex-A9 TRM** at 7-3.

554. In addition, the L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on

designs provided by ARM. In addition, the L1 and L2 cache systems in the **Cortex-A9 SOCs** interoperate. The L1 cache controllers meet **Claim 15(h) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

555. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** can each also accept multiple outstanding read and write requests. **Cortex-A15 TRM** at 6-6 and 7-12.

556. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 15(h) of the '158 Patent** in a similar manner to that described above. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 15(h) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

557. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

558. Each cache request will be processed by the cache controller after the cache controller receives the request.

559. Multiple outstanding cache requests likely occur in any reasonable usage scenario.

560. As described in the preceding paragraphs, **Claim 15(h) of the '158 Patent** is met when an **Accused SOC** is used.

561. The **Unified Cache GPUs** practice **Claim 15(h) of the '158 Patent**.

562. The **Unified Cache GPUs** can accept multiple outstanding cache requests.

563. For example, the NVIDIA GF100 GPU's unified L2 cache dynamically load balances between different requests. **NVIDIA GF100 Whitepaper** at 20.

564. Cache controllers in the **Unified Cache GPUs** can accept two cache requests and complete processing of both cache requests after receiving the second cache request.

565. Cache requests are utilized during ordinary usage of GPUs that utilize L2 cache systems.

566. Each cache request will be processed by the cache controller after the cache controller receives the request.

567. Multiple outstanding cache requests likely occur in any reasonable usage scenario.

568. As described in the preceding paragraphs, **Claim 15(h) of the '158 Patent** is met when a **Unified Cache GPU** is used.

569. As described in the preceding paragraphs, **Claim 15(h) of the '158 Patent** is met when a **'158 Accused Product** is used.

570. "**Claim 15(i) of the '158 Patent**" recites "completing the processing of the second cache request after receiving the second cache request."

571. All **'158 Accused Products** complete the processing of the second cache request after receiving the second cache request.

572. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving the cache request.

573. If the second cache request is a read request, the PL310 cache controller will process that request after receiving it. If the second cache request is a write request, the PL310 cache controller will process that request after receiving it.

574. For example, in the timing diagram depicted below, the PL310 cache controller receives the second read request, for address B, at time T3. The processing of that request occurs after time T3.

C.4 Outstanding read hit transaction

Figure C-4 shows the timing for an outstanding read hit transaction.

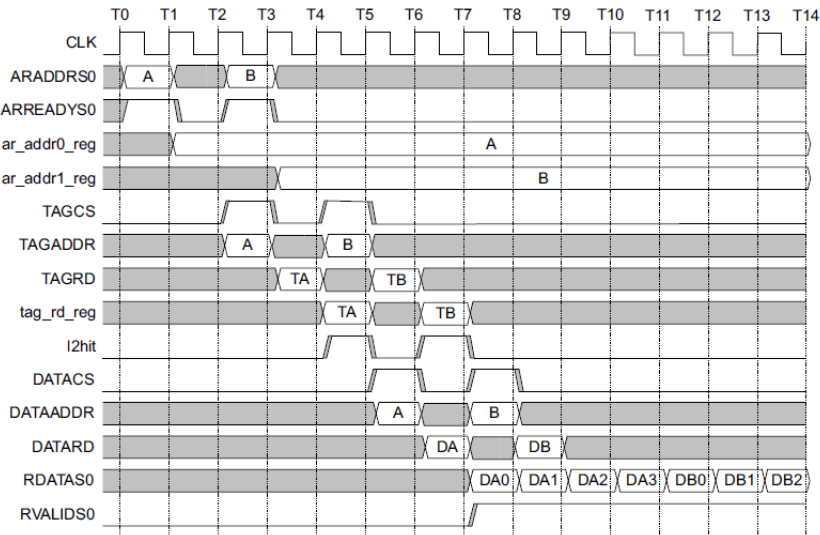


Figure C-4 Outstanding read hit transaction

PL310 TRM at Fig. C-4.

575. In the timing diagram depicted above, processing of the second cache request is completed after receiving the second cache request.

576. The L1 cache controllers in the **Cortex-A9 SOCs** process each cache request after receiving that cache request.

577. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** process each request after receiving that request.

578. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

579. Each cache request will be processed by the cache controller after the cache controller receives the request.

580. Multiple cache requests likely occur in any reasonable usage scenario.

581. As described in the preceding paragraphs, **Claim 15(i) of the '158 Patent** is met when an **Accused SOC** is used.

582. The **Unified Cache GPUs** practice **Claim 15(i) of the '158 Patent**.

583. For example, the L2 cache controllers in the **Unified Cache GPUs** process each cache request after receiving the cache request.

584. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

585. Each cache request will be processed by the cache controller after the cache controller receives the request.

586. Multiple cache requests likely occur in any reasonable usage scenario.

587. As described in the preceding paragraphs, **Claim 15(i) of the '158 Patent** is met when a **Unified Cache GPU** is used.

588. As described in the preceding paragraphs, **Claim 15(i) of the '158 Patent** is met when a **'158 Accused Product** is used.

589. "**Claim 17 of the '158 Patent**" recites "[t]he method of claim 15 wherein the first and second devices are the same device."

590. As described above, all **'158 Accused Products** practice claim 15.

591. In addition, all **'158 Accused Products** practice the method of claim 15 wherein the first and second devices are the same device.

592. As noted above, with respect to **Claim 15(e) of the '158 Patent**, the first device and the second device can be the same device.

593. “**Claim 18(a) of the '158 Patent**” recites “[a] method of controlling a cache, the cache coupled to a cache controller and a cache accessing device, the method comprising.”

594. All **'158 Accused Products** practice a method of controlling a cache, the cache coupled to a cache controller and a cache accessing device.

595. For example, the **Cortex-A9 SOCs** include one or more PL310 cache controllers. One or more of these PL310 cache controllers are used for L2 cache control.

596. The PL310 cache controllers used for L2 cache control are coupled to cache memory.

597. The PL310 cache controllers used for L2 cache control are also coupled to cache accessing devices. For example, the PL310 cache controllers are coupled to the CPU cores, L1 caches, and snoop control units in the **Cortex-A9 SOCs**.

598. The cache memory is coupled to the cache accessing devices via the PL310 cache controllers.

599. The **Cortex-A15 SOCs** use one or more L2 cache controllers. One or more of these L2 cache controllers are integrated or otherwise associated with the CPU complex.

600. The L2 cache controllers (referenced in the immediately preceding paragraph) are coupled to cache memory. These cache controllers are also coupled to one or more cache accessing devices. For example, the L2 cache controllers are coupled to the CPU cores and L1 caches in the **Cortex-A15 SOCs**. The cache memory is coupled to the cache accessing devices via the L2 cache controllers.

601. All **Accused SOCs** include one or more L1 cache controllers. One or more of these L1 cache controllers are integrated with the CPU complex.

602. The L1 cache controllers (referenced in the immediately preceding paragraph) are coupled to cache memory. These cache controllers are also to coupled one or more cache accessing devices. For example, the cache controllers are each coupled to a CPU core. The cache memories are coupled to the CPU cores via the L1 cache controllers.

603. All **Unified Cache GPUs** include one or more L2 cache controllers. One or more of these L2 cache controllers are coupled to cache memory.

604. The L2 cache controllers (referenced in the immediately preceding paragraph) are also coupled to one or more cache accessing upstream devices. The cache memory is coupled to the one or more upstream devices via the L2 cache controllers.

605. For example, the NVIDIA GF100 GPU includes 768 KB of L2 cache memory. **NVIDIA GF100 Whitepaper** at 11. This memory is coupled to the L2 cache controller.

606. The **NVIDIA GF100 Whitepaper** depicts the L2 cache in the following block diagram.



GF100 block diagram showing the Host Interface, the GigaThread Engine, four GPCs, six Memory Controllers, six ROP partitions, and a 768 KB L2 cache. Each GPC contains four PolyMorph engines. The ROP partitions are immediately adjacent to the L2 cache.

NVIDIA GF100 Whitepaper at 11.

607. As described in the preceding paragraphs, the **'158 Accused Products** practice a method of controlling a cache, the cache coupled to a cache controller and a cache accessing device.

608. **"Claim 18(b) of the '158 Patent"** recites "a first step of receiving a first cache request from the cache accessing device by the cache controller."

609. All **'158 Accused Products** practice a first step of receiving a first cache request from the cache accessing device by the cache controller.

610. For example, the cache controllers in **Accused SOCs** receive cache requests from devices to which they are coupled.

611. The PL310 cache controllers in the **Cortex-A9 SOCs** receive cache requests from the cache accessing devices to which they are coupled, such as the CPU cores. These cache requests may include write requests.

612. The L2 cache controllers in the **Cortex-A15 SOCs** receive cache requests from the cache accessing devices to which they are coupled, such as the CPU cores. These cache requests may include read requests and write requests.

613. The L1 cache controllers in the **Accused SOCs** receive cache requests from the CPU cores to which they are coupled. These cache requests may include read requests and write requests.

614. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

615. The **Unified Cache GPUs** practice **Claim 18(b) of the '158 Patent**.

616. For example, the L2 cache controllers in the **Unified Cache GPUs** receive cache requests from one or more of the upstream devices to which they are coupled.

617. The L2 cache controller in the NVIDIA GF100 GPU services all load, store, and texture requests. **NVIDIA GF100 Whitepaper** at 20.

618. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

619. As described in the preceding paragraphs, **Claim 18(b) of the '158 Patent** is met when a **'158 Accused Product** is used.

620. "**Claim 18(c) of the '158 Patent**" recites "a second step of providing a first request ID to the cache accessing device by the cache controller after the first step."

621. All **'158 Accused Products** practice a second step of providing a first request ID to the cache accessing device by the cache controller after the first step.

622. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** provide request ID values corresponding to each cache request to the requesting device.

623. If a cache request is a write request, the PL310 cache controller that receives that request will provide the request ID value in the form of a "Response ID tag" or "BID." **AXI Spec** at 33. This signal is the ID tag of the of the write response. *Id.*

624. The request ID is provided by the PL310 cache controller after receiving the corresponding request.

625. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 18(c) of the '158 Patent** in a similar manner to that described above. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on designs provided by ARM. In addition, the L1 and L2 cache systems in the **Cortex-A9 SOCs** interoperate. The L1 cache controllers meet **Claim 18(c) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

626. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 18(c) of the '158 Patent** in a similar manner to that described above. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 18(c) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

627. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

628. For the reasons described above, request ID values corresponding to each cache request will be provided.

629. The **Unified Cache GPUs** practice **Claim 18(c) of the '158 Patent**.

630. For example, the L2 cache controllers in the **Unified Cache GPUs** provide request ID values corresponding to each cache request to the requesting device.

631. Such request ID values are provided after the L2 cache controller receives the cache request.

632. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

633. Request ID values corresponding to each cache request will be provided.

634. As described in the preceding paragraphs, **Claim 18(c) of the '158 Patent** is met when a **'158 Accused Product** is used.

635. “**Claim 18(d) of the '158 Patent**” recites “a third step of storing the first request ID by the cache accessing device after the second step.”

636. All **'158 Accused Products** practice a third step of storing the first request ID by the cache accessing device after the second step.

637. For example, the requesting device in the **Cortex-A9 SOCs** stores the request IDs provided by the cache controllers.

638. If a cache request is a write request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Response ID tag” or “BID.” **AXI Spec** at 33. This signal is the ID tag of the of the write response. *Id.*

639. The request ID is stored by the cache accessing device.

640. For example, if the PL310 cache controller receives a write request from a requesting device, the PL310 cache controller will provide a request ID, which may be a BID signal. The requesting device will store the request ID.

641. The CPU cores in the **Cortex-A9 SOCs** practice **Claim 18(d) of the '158 Patent** in a similar manner by storing request IDs provided by the L1 cache controllers. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on designs provided by ARM. In addition, the L1 and L2 cache systems. The CPU cores in the **Cortex-A9 SOCs** store request IDs provided by the L1 cache controllers.

642. The cache accessing devices in the **Cortex-A15 SOCs** practice **Claim 18(d) of the '158 Patent** in a similar manner by storing request IDs provided by the L1 and/or L2 cache controllers. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The cache accessing devices in the **Cortex-A15 SOCs** store request IDs provided by the L1 and/or L2 cache controllers.

643. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

644. For the reasons described above, request ID values corresponding to each cache request will be provided.

645. The request ID values will be stored by the cache accessing device.

646. The **Unified Cache GPUs** practice **Claim 18(d) of the '158 Patent**.

647. For example, the requesting device in the **Unified Cache GPUs** stores the request ID values provided by the cache controllers.

648. Request ID values are provided after the L2 cache controller receives the cache request.

649. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

650. Request ID values corresponding to each cache request will be provided.

651. The request ID values will be stored by the cache accessing device.

652. As described in the preceding paragraphs, **Claim 18(d) of the '158 Patent** is met when a **'158 Accused Product** is used.

653. "**Claim 18(e) of the '158 Patent**" recites "a fourth step of receiving a second cache request from the cache accessing device by the cache controller after the first step."

654. All **'158 Accused Products** practice a fourth of receiving a second cache request from the cache accessing device by the cache controller after the first step.

655. Cache controllers in the **Cortex-A9 SOCs** can accept multiple outstanding read and write requests.

656. For example, the PL310 cache controllers can accept up to 16 read requests and up to 8 write requests. **PL310 TRM** at 2-5.

657. The data cache in the L1 cache can accept up to four outstanding data cache read misses and up to four outstanding data cache write misses. **Cortex-A9 TRM** at 7-3.

658. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** can each also accept multiple outstanding read and write requests. **Cortex-A15 TRM** at 6-6 and 7-12.

659. Multiple cache requests likely occur in any reasonable usage scenario.

660. The second cache request can be received from the same cache accessing device from which the first cache request was received, or from a different cache accessing device. For

example, if the first cache request is received from one CPU core, the second cache request may be received from the same CPU core, or from a different CPU core.

661. The **Unified Cache GPUs** practice **Claim 18(e) of the '158 Patent**.

662. The **Unified Cache GPUs** can accept multiple outstanding cache requests.

663. For example, the NVIDIA GF100 GPU's unified L2 cache dynamically load balances between different requests. **NVIDIA GF100 Whitepaper** at 20.

664. Multiple cache requests likely occur in any reasonable usage scenario.

665. As described in the preceding paragraphs, **Claim 18(e) of the '158 Patent** is met when a **'158 Accused Product** is used.

666. "**Claim 18(f) of the '158 Patent**" recites "a fifth step of providing a second request ID to the cache accessing device by the cache controller after the fourth step."

667. All **'158 Accused Products** practice a fifth step of providing a second request ID to the cache accessing device by the cache controller after the fourth step.

668. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** provide request ID values corresponding to each cache request to the requesting device.

669. If a cache request is a write request, the PL310 cache controller that receives that request will provide the request ID value in the form of a "Response ID tag" or "BID." **AXI Spec** at 33. This signal is the ID tag of the of the write response. *Id.*

670. The request ID is provided by the PL310 cache controller after receiving the corresponding request.

671. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 18(f) of the '158 Patent** in a similar manner to that described above. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the

Cortex-A9 SOCs are also based on designs provided by ARM. In addition, the L1 and L2 cache systems interoperate. The L1 cache controllers meet **Claim 18(f) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

672. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 18(f) of the '158 Patent** in a similar manner. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 18(f) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

673. Multiple cache requests likely occur in any reasonable usage scenario.

674. For the reasons described above, request ID values corresponding to each cache request will be provided.

675. The **Unified Cache GPUs** practice **Claim 18(f) of the '158 Patent**.

676. The L2 cache controllers in the **Unified Cache GPUs** provide request ID values corresponding to each cache request to the requesting device.

677. Such request ID values are provided after the L2 cache controller receives the cache request.

678. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

679. Request ID values corresponding to each cache request will be provided.

680. Multiple cache requests likely occur in any reasonable usage scenario.

681. As described in the preceding paragraphs, **Claim 18(f) of the '158 Patent** is met when a **'158 Accused Product** is used.

682. “**Claim 18(g) of the ’158 Patent**” recites “a sixth step of storing the second request ID by the cache accessing device after the fifth step.”

683. All **’158 Accused Products** practice a sixth step of storing the second request ID by the cache accessing device after the fifth step.

684. For example, the requesting device in the **Cortex-A9 SOCs** stores the request IDs provided by the cache controllers.

685. If a cache request is a write request, the PL310 cache controller that receives that request will provide the request ID value in the form of a “Response ID tag” or “BID.” **AXI Spec** at 33. This signal is the ID tag of the of the write response. *Id.*

686. The request ID is stored by the cache accessing device.

687. For example, if the PL310 cache controller receives a write request from a requesting device, the PL310 cache controller will provide a request ID, which may be a BID signal. The requesting device will store the request ID.

688. The CPU cores in the **Cortex-A9 SOCs** practice **Claim 18(g) of the ’158 Patent** in a similar manner by storing request IDs provided by the L1 cache controllers. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on designs provided by ARM. In addition, the L1 and L2 cache systems in the **Cortex-A9 SOCs** interoperate. The CPU cores in the **Cortex-A9 SOCs** store request IDs provided by the L1 cache controllers.

689. The cache accessing devices in the **Cortex-A15 SOCs** practice **Claim 18(g) of the ’158 Patent** in a similar manner by storing request IDs provided by the L1 and/or L2 cache controllers. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role

as the L2 cache controllers in the **Cortex-A15 SOCs**. The cache accessing devices in the **Cortex-A15 SOCs** store request IDs provided by the L1 and/or L2 cache controllers.

690. Multiple cache requests likely occur in any reasonable usage scenario.

691. For the reasons described above, request ID values corresponding to each cache request will be provided.

692. The request ID values will be stored by the cache accessing device.

693. The **Unified Cache GPUs** practice **Claim 18(g) of the '158 Patent**.

694. For example, the requesting device in the **Unified Cache GPUs** stores the request ID values provided by the cache controllers.

695. Request ID values are provided after the L2 cache controller receives the cache request.

696. Cache requests are utilized during ordinary usage of a GPU that utilizes an L2 cache.

697. Request ID values corresponding to each cache request will be provided.

698. The request ID values will be stored by the cache accessing device.

699. As described in the preceding paragraphs, **Claim 18(g) of the '158 Patent** is met when a **'158 Accused Product** is used.

700. "**Claim 18(h) of the '158 Patent**" recites "a seventh step of processing the first cache request after the second step."

701. All **'158 Accused Products** practice a seventh step of processing the first cache request after the second step.

702. Cache controllers in the **Cortex-A9 SOCs** process cache requests after providing a request ID for that cache request to the cache accessing device.

703. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** process write requests after providing the BID signal to the cache accessing device.

704. The PL310 cache controller can provide a BID signal in response to a write request once the memory system is able to guarantee that the effect of the write is visible to all processors in the system. **PL310 TRM** at Glossary-11.

705. Processing of the write request can occur after the BID signal is returned to the cache accessing device.

706. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 18(h) of the '158 Patent** in a similar manner to that described above. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on designs provided by ARM. In addition, the L1 and L2 cache systems interoperate. The L1 cache controllers meet **Claim 18(h) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

707. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 18(h) of the '158 Patent** in a similar manner. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 18(h) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

708. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L1 and L2 cache systems.

709. For the reasons described above, request ID values corresponding to each cache request will be provided.

710. Processing of the corresponding cache request will occur after the request ID value is provided.

711. As described in the preceding paragraphs, **Claim 18(h) of the '158 Patent** is met when a **'158 Accused Product** is used.

712. "**Claim 18(i) of the '158 Patent**" recites "an eighth step of processing the second cache request after the fifth step."

713. All **'158 Accused Products** practice an eighth step of processing the second cache request after the fifth step.

714. Cache controllers in the **Cortex-A9 SOCs** process cache requests after providing a request ID for that cache request to the cache accessing device.

715. For example, the PL310 cache controllers in the **Cortex-A9 SOCs** process write requests after providing the BID signal to the cache accessing device.

716. The PL310 cache controller can provide a BID signal in response to a write request once the memory system is able to guarantee that the effect of the write is visible to all processors in the system. **PL310 TRM** at Glossary-11.

717. Processing of the write request can occur after the BID signal is returned to the cache accessing device.

718. The L1 cache controllers in the **Cortex-A9 SOCs** practice **Claim 18(i) of the '158 Patent** in a similar manner. The L1 cache controllers in the **Cortex-A9 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** are also based on designs provided by ARM. In addition, the L1 and L2 cache systems in the **Cortex-A9 SOCs** interoperate. The L1 cache controllers meet **Claim 18(i) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

719. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 18(i) of the '158 Patent** in a similar manner. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L1 and L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 18(i) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

720. Multiple cache requests likely occur in any reasonable usage scenario.

721. For the reasons described above, request ID values corresponding to each cache request will be provided.

722. Processing of the corresponding cache request will occur after the request ID value is provided.

723. As described in the preceding paragraphs, **Claim 18(i) of the '158 Patent** is met when a **'158 Accused Product** is used.

724. “**Claim 21 of the '158 Patent**” recites “[t]he method of claim 18” wherein “the first cache request is received by the cache controller from a first cache accessing device” and “the second cache request is received by the cache controller from a second cache accessing device.”

725. As described above, all **'158 Accused Products** practice claim 18. In addition, all **'158 Accused Products** practice **Claim 21 of the '158 Patent**.

726. As noted above, with respect to **Claim 18(e) of the '158 Patent**, the second cache request can be received from the same cache accessing device from which the first cache request was received, or from a different cache accessing device.

727. For example, at least one PL310 cache controller in each **Cortex-A9 SOC** is coupled to multiple CPU cores. That PL310 cache controller can receive requests from each of those CPU cores.

728. Similarly, at least one L2 cache controller in each **Cortex-A15 SOC** is coupled to multiple CPU cores. That L2 cache controller can receive requests from each of those CPU cores.

729. It is likely that the L2 cache controllers coupled to multiple CPU cores will receive cache requests from more than one CPU core.

730. For example, the unified L2 cache controller in the **Unified Cache GPUs** is coupled to multiple upstream cache accessing devices.

731. It is likely that the unified L2 cache controller will receive requests from multiple upstream cache accessing devices.

732. As described in the preceding paragraphs, **Claim 21 of the '158 Patent** is met when a **'158 Accused Product** is used.

733. "**Claim 22(a) of the '158 Patent**" recites "[t]he method of claim 21."

734. As described above, all **'158 Accused Products** practice **Claim 21 of the '158 Patent**.

735. "**Claim 22(b) of the '158 Patent**" recites "wherein the first step comprises: receiving the first cache request by the cache controller [and] receiving a first device ID identifying the first cache accessing device to the cache controller."

736. All **'158 Accused Products** practice **Claim 22(b) of the '158 Patent**.

737. The PL310 cache controllers in the **Cortex-A9 SOCs** receive cache requests from devices to which they are coupled.

738. For example, the PL310 cache controllers can receive cache requests from the CPU cores. These cache requests may include write requests.

739. The write requests received by the PL310 cache controller include a signal called “write address ID” or “AWID.” **AXI Spec** at 31.

740. The AWID signal includes bits that identify the cache accessing device to the cache controller. **AXI Spec** at 82.

741. The L2 cache controllers in the **Cortex-A15 SOCs** receive cache requests from devices to which they are coupled.

742. For example, the L2 cache controllers can receive cache requests from the CPU cores. These cache requests may include read requests and write requests.

743. The L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 22(b) of the '158 Patent** in a similar manner to that described above. The L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 22(b) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

744. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L2 cache systems.

745. As described in the preceding paragraphs, **Claim 22(b) of the '158 Patent** is met when a **'158 Accused Product** is used.

746. “**Claim 22(c) of the '158 Patent**” recites “wherein the fourth step comprises: receiving the second cache request by the cache controller [and] receiving a second device ID identifying the second cache accessing device to the cache controller.”

747. All **'158 Accused Products** practice **Claim 22(c) of the '158 Patent**.

748. The PL310 cache controllers in the **Cortex-A9 SOCs** receive cache requests from devices to which they are coupled.

749. For example, the PL310 cache controllers can receive cache requests from the CPU cores. These cache requests may include write requests.

750. The write requests received by the PL310 cache controller include a signal called "write address ID" or "AWID." **AXI Spec** at 31.

751. The AWID signal includes bits that identify the cache accessing device to the cache controller. **AXI Spec** at 82.

752. The L2 cache controllers in the **Cortex-A15 SOCs** receive cache requests from devices to which they are coupled.

753. For example, the L2 cache controllers can receive cache requests from the CPU cores. These cache requests may include read requests and write requests.

754. The L2 cache controllers in the **Cortex-A15 SOCs** practice **Claim 22(c) of the '158 Patent** in a similar manner to that described above. The L2 cache controllers in the **Cortex-A15 SOCs** are based on designs provided by ARM. The PL310 cache controllers in the **Cortex-A9 SOCs** perform a similar role as the L2 cache controllers in the **Cortex-A15 SOCs**. The L2 cache controllers in the **Cortex-A15 SOCs** meet **Claim 22(c) of the '158 Patent** in substantially the same manner as the PL310 cache controllers.

755. Cache read and write requests are utilized during ordinary usage of computer systems that utilize L2 cache systems.

756. As described in the preceding paragraphs, **Claim 22(c) of the '158 Patent** is met when a **'158 Accused Product** is used.

757. NVIDIA and Velocity have directly infringed and continue to directly infringe the **'158 Patent** by using the **'158 Accused Products**.

758. NVIDIA uses the **'158 Accused Products**.

759. For example, NVIDIA uses its own products for demonstration, testing, and development purposes.

760. NVIDIA has used and continues to use its own products, including the **'158 Accused Products**, such as for demonstration, testing, and development purposes.

761. For example and without limitation, NVIDIA has used and continues to use the GM107, Tesla K10, NVIDIA Shield Tablet, and NVIDIA Shield Portable.

762. NVIDIA also has used and continues to use third-party products, including the **'158 Accused Products**, such as for demonstration, testing, and development purposes.

763. The third party graphics cards used by NVIDIA include one or more **'158 Accused Products**.

764. When NVIDIA uses the **'158 Accused Products**, such as for demonstration, testing, or development purposes, such use directly infringes the **'158 Patent**.

765. Velocity sells computers that incorporate a **Unified Cache GPU**.

766. For example, Velocity sells a computer called the Raptor Signature Edition.

767. The Raptor Signature Edition computer can be configured to include a variety of NVIDIA GPUs, including a variety of the **Unified Cache GPUs**.

768. For example, the Raptor Signature Edition computer can be configured to include an NVIDIA GeForce GTX Titan Black.

769. Velocity has used and continues to use its own products, including the **'158 Accused Products**, such as for testing and development purposes.

770. For example and without limitation, Velocity has used and continues to use the Cruz Tablet L510 incorporating the NVIDIA Tegra 250 SOC

771. For example and without limitation, Velocity has used and continues to use computers that incorporate a **Unified Cache GPU**.

772. When Velocity uses the **'158 Accused Products**, such as for demonstration, testing, or development purposes, such use directly infringes the **'158 Patent**.

773. In 2013, Samsung and NVIDIA were engaged in negotiations regarding rights to various patents.

774. On or around August 7, 2013, Samsung sent infringement claim charts and other documents to NVIDIA that specifically highlighted certain exemplary patents, claims, and products.

775. The claim charts Samsung sent to NVIDIA on or about August 7, 2013 included claim charts relating to the **'158 Patent**.

776. NVIDIA has had actual knowledge of the **'158 Patent** since at least as early as August 7, 2013.

777. Velocity has had actual knowledge of the **'158 Patent** at least as of November 11, 2014, the date the original complaint in this action was served.

778. NVIDIA and Velocity indirectly infringe the **'158 Patent** by inducing infringement by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(b), in this District and elsewhere in the United States.

779. Direct infringement is the result of activities performed by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, by, for example, using the **'158 Accused Products**.

780. NVIDIA and Velocity have induced and continue to induce infringement of the **'158 Patent** by intending that others infringe the **'158 Patent** by using the **'158 Accused Products**. NVIDIA and Velocity designed the **'158 Accused Products** such that they would each infringe one or more claims of the **'158 Patent** when used.

781. NVIDIA and Velocity provide the **'158 Accused Products** to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users.

782. By providing **'158 Accused Products** to others, NVIDIA and Velocity intend for **'158 Accused Products** to be used in the United States.

783. NVIDIA and Velocity also provide others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, instructions, user guides, and technical specifications.

784. When others follow such instructions, user guides, and/or other design documentation, they directly infringe one or more claims of the **'158 Patent**.

785. NVIDIA and Velocity know that by providing such instructions, user guides, and/or other design documentation, others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard

manufacturers, add-in board manufacturers, resellers, customers, and end users, follow those instructions, user guides, and other design documentation, and directly infringe one or more claims of the **'158 Patent**. NVIDIA and Velocity thus know that their actions actively induce infringement.

786. NVIDIA sells the **'158 Accused Products** directly to customers and end users.

787. For example, NVIDIA sells the NVIDIA Shield Tablet and NVIDIA Shield Portable via its online store at <http://store.nvidia.com>.

788. Through such sales, NVIDIA specifically intends that others, such as customers and end users, will infringe one or more claims of the **'158 Patent**.

789. NVIDIA provides the **'158 Accused Products**, and reference designs for the **'158 Accused Products**, to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, and resellers.

790. Through such activity, NVIDIA specifically intends that others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, and resellers, customers, and end users, infringe the **'158 Patent** by using in the United States the **'158 Accused Products**.

791. Through its manufacture (either directly, or through contract manufacturing facilities) and/or sale of the infringing products, NVIDIA specifically intends that others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, will infringe one or more claims of the **'158 Patent**.

792. NVIDIA specifically intends for others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'158 Patent** in the United States.

793. For example, NVIDIA advertised its “NVIDIA Authorized Board Partner Program ensures an exceptional customer experience when purchasing graphics cards and motherboards manufactured by partners that make use of NVIDIA’s latest technologies.”

NVIDIA PartnerForce Info.

794. NVIDIA claims that its “NVIDIA Authorized Board Partners offer the latest technologies from NVIDIA” and lists six Authorized Board Partners in the United States, six in Canada, and none in any other countries. *See* **NVIDIA PartnerForce Info.** NVIDIA also lists ten distributors in the United States (including “pre and post sales technical support”) and seven in Canada but none in any other country. *See* **NVIDIA PartnerForce Info.**

795. As an additional example, as of October 31, 2014, NVIDIA also advertised the “NVIDIA PartnerForce Program” which is “a sales and marketing program for value-added resellers, system builders, etailers and retailers who sell NVIDIA based components or systems.”

NVIDIA PartnerForce Program.

796. NVIDIA advertises in the United States and in this judicial district that “[b]y joining the PartnerForce Program, you may leverage the world-class brands and technology platforms from NVIDIA, sales and technical support from our experienced team and hundreds of turnkey sales and marketing tools.” **NVIDIA PartnerForce Program.**

797. NVIDIA also provides marketing materials and templates, including retail display items, web banners, copy blocks, logos, product shots, email and web page templates, to

members of the “NVIDIA PartnerForce Program” which includes value-added resellers, system builders, etailers and retailers inducing them to offer to sell and the **'158 Accused Products**. *See NVIDIA PartnerForce Program.*

798. At the Ford Event and Conference Center, in Dearborn, Michigan, NVIDIA showcased its latest processor technologies, which power everything from the CAD software that designers use to style cars to the infotainment systems that drivers use to map their trips and listen to music, in an effort to encourage domestic car manufacturers to include infringing technology in vehicles manufactured and sold in the United States. *See Traveling the Road to Silicon Motown.*

799. As described in the preceding paragraphs, NVIDIA specifically targets the United States market for **'158 Accused Products** actively induces others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'158 Patent**.

800. Velocity sells computers that incorporate a **Unified Cache GPU**.

801. For example, Velocity sells a computer called the Raptor Signature Edition.

802. The Raptor Signature Edition computer can be configured to include a variety of NVIDIA GPUs, including a variety of the **Unified Cache GPUs**.

803. For example, the Raptor Signature Edition computer can be configured to include an NVIDIA GeForce GTX Titan Black.

804. Velocity provides the **'158 Accused Products** to others, such as resellers, customers, and end users.

805. Through such activity, NVIDIA specifically intends that others, such as resellers, customers, and end users, infringe the **'158 Patent** by using in the United States the **'158 Accused Products**.

806. Velocity specifically targets the United States market for its products listed above and actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'158 Patent** in the United States.

807. For example, Velocity sells products via its website, <http://www.velocitymicro.com>, to customers in the United States. Such products include the **'158 Accused Products**, such as, for example, computers that are configured to include an NVIDIA GeForce GTX Titan Black.

808. The customers infringe the **'158 Patent** by using the **'158 Accused Products** purchased from Velocity.

809. Velocity also sells computers to resellers. Such computers include the **'158 Accused Products**, such as, for example, computers that are configured to include a **Unified Cache GPU**.

810. For example, Velocity products have been sold at nearly every major electronics retailer in North America including Best Buy, Circuit City, RadioShack, Costco, Sears, Target, and many others. *See Velocity Company History*.

811. Velocity introduces products and services that infringe the Asserted Patents intending that they would be used in this Judicial District and elsewhere in the United States.

812. The resellers infringe the **'158 Patent** by using the **'158 Accused Products**.

813. The resellers also sell the **'158 Accused Products** to customers and end users, who infringe the **'158 Patent** by using those products.

814. As described in the preceding paragraphs, Velocity specifically targets the United States market for **'158 Accused Products** actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'158 Patent**.

815. NVIDIA and Velocity indirectly infringe the **'158 Patent** by contributing to infringement by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(c), in this District and elsewhere in the United States.

816. Direct infringement is the result of activities performed by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users of the infringing products.

817. For example, the **'158 Accused Products** allow for the control of cache. When the infringing products are used as intended by others such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, the **'158 Accused Products** necessarily control the cache in an infringing manner. The infringing products cannot operate in an acceptable manner without controlling the cache as claimed in the **'158 Patent**.

818. From the facts set forth above, it is evident that NVIDIA and Velocity knew that the ability to control the cache in the infringing products is especially made or especially adapted to operate in the products of NVIDIA's and Velocity's distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard

manufacturers, add-in board manufacturers, resellers, customers, and end users, and is not a staple article or commodity of commerce and that its infringing use is required for operation of the infringing products. Any other use would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental.

819. NVIDIA's and Velocity's infringing products, with the ability to control the cache, are each a material part of the invention of the **'158 Patent** and are especially made or adapted to infringe one or more claims of the **'158 Patent**. Because the use of the **'158 Accused Products** necessarily infringes one or more claims of the **'158 Patent**, NVIDIA's and Velocity's sales of its infringing products have no substantial non-infringing uses.

820. Accordingly, NVIDIA and Velocity offer to sell, or sell a component, material, or apparatus for use in practicing one or more claims of the **'158 Patent** knowing the same to be especially made or especially adapted for use in an infringement of such patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

821. NVIDIA's infringement of the **'158 Patent** is willful and deliberate, entitling Samsung to increased damages under 35 U.S.C. § 284 and to attorneys' fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

822. Samsung has no adequate remedy at law for NVIDIA's and Velocity's infringement of the **'158 Patent** and is suffering irreparable harm, requiring permanent injunctive relief.

SECOND CLAIM FOR RELIEF
INFRINGEMENT OF U.S. PATENT NO. 6,262,938
(AGAINST NVIDIA AND VELOCITY)

823. Each of the above listed paragraphs is incorporated herein by reference and adopted, as if fully set forth again.

824. The **'938 Patent** was filed on March 3, 2000, issued on July 17, 2001, and is entitled "Synchronous DRAM Having Posted CAS Latency and Method for Controlling CAS Latency." The **'938 Patent** is generally directed to an SDRAM that supports an operation known as "posted CAS latency" and methods of controlling CAS latency.

825. The **'938 Patent** was assigned to SEC, and SEC continues to hold all rights, title, and interest in the **'938 Patent**. A true and correct copy of the **'938 Patent** is attached hereto as Exhibit B.

826. The **'938 Accused Products** infringe the **'938 Patent** in substantially the same way. All **Accused SOCs** and **'938 Accused GPUs** are designed to work in conjunction with memory, such as DDR3 or GDDR2, that is manufactured according to the standards set by JEDEC. All products designed to work with JEDEC standard memory must interface with the memory in the manner required by the standard.

827. As shown in the figure below, the GeForce G100 GPU, for example, is designed for use with GDDR2 memory.

The screenshot shows the NVIDIA GeForce website's hardware page for the GeForce G100. The 'Specifications' section is active, and the 'Memory Interface' is highlighted with a red box. The specifications table is as follows:

GPU Engine Specs	
CUDA Cores	8
Graphics Clock (MHz)	567MHz
Processor Clock (MHz)	1400MHz
Texture Fill Rate (billion/sec)	4.3
Memory Specs	
Memory Clock	500MHz
Standard Memory Config	512MB
Memory Interface	GDDR2
Memory Interface Width	64-bit
Memory Bandwidth (GB/sec)	8.0

GeForce G100 Specifications at 1 (red annotation added).

828. GDDR2 memory is JEDEC-standard memory. *See, e.g., JEDEC Standard No. 21-C.*

829. “**Claim 1(a) of the '938 Patent**” recites “[a] synchronous DRAM (SDRAM), operating in synchronization with a clock signal.”

830. The **'938 Accused Products** include an SDRAM that operates in synchronization with a clock signal.

831. The **'938 Accused GPUs** and **Accused SOCs** are designed to operate in conjunction with JEDEC standard SDRAM or SGRAM.

832. For example, the Tegra 2 is designed for use in conjunction with at least LP-DDR, LP-DDR2, or DDR2 memory. *See, e.g., Tegra 2 TRM* at 8, 492.

833. For example, the Tegra 3 is designed for use in conjunction with at least LP-DDR2 or DDR3 memory. *See, e.g., Tegra 3 TRM* at 10, 437.

834. For example, the Tegra K1 is designed for use in conjunction with at least LP-DDR3 and DDR3L memory. *See, e.g., Tegra K1 TRM* at 12, 643.

835. For example, the Tegra 4 is designed for use in conjunction with at least DDR3, low-voltage DDR, LP-DDR2, and LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 12, 1115.

836. DDR3 and DDR2 are synchronous DRAM. *See, e.g., JEDEC Standard No. 79-3F*.

837. Synchronous DRAM (SDRAM) operates in synchronization with a clock signal.

838. Any synchronous DRAM, including DDR3 and DDR2, provided by any manufacturer complies with the relevant JEDEC standard.

839. For example, DDR3 operates in synchronization with a clock signal. *See, e.g., JEDEC Standard No. 79-3F* at 56.

840. The DDR3 standard demonstrates that a read operation is synchronized with the rising and falling edges of a clock signal. *See, e.g., JEDEC Standard No. 79-3F* at 56.

841. For example, DDR2 operates in synchronization with a clock signal. *See, e.g., JEDEC Standard No. 79-2F* at 33.

842. The DDR2 standard demonstrates that read and write operations are synchronized with the rising and falling edges of a clock signal. *See, e.g., JEDEC Standard No. 79-2F* at 33.

843. As described in the preceding paragraphs, the **'938 Accused Products** include an SDRAM operating in synchronization with a clock signal.

844. "**Claim 1(b) of the '938 Patent**" recites "a memory bank having a plurality of memory cells arranged in rows and columns."

845. The **'938 Accused Products** include a memory bank having a plurality of memory cells arranged in rows and columns.

846. For example, DDR3 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-3F* at 15.

847. For example DDR2 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-2F* at 15.

848. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 1(b) of the '938 Patent.**"

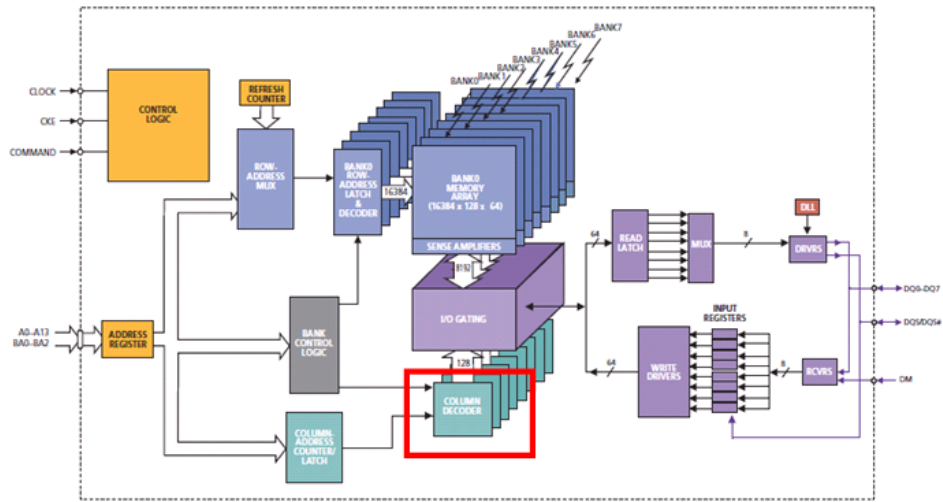
849. "**Claim 1(c) of the '938 Patent**" recites "a column decoder for selecting a column of the memory bank."

850. The **'938 Accused Products** include a column decoder for selecting a column of the memory bank.

851. For example, the Tegra 3 connects to DDR3 SDRAM. *See, e.g., Tegra 3 TRM* at 437.

852. For example, as shown in the figures below, Micron DDR3 SDRAM includes a column decoder that decodes a specific column upon a READ command.

Figure 1: 1Gb DDR3 SDRAM Functional Block Diagram



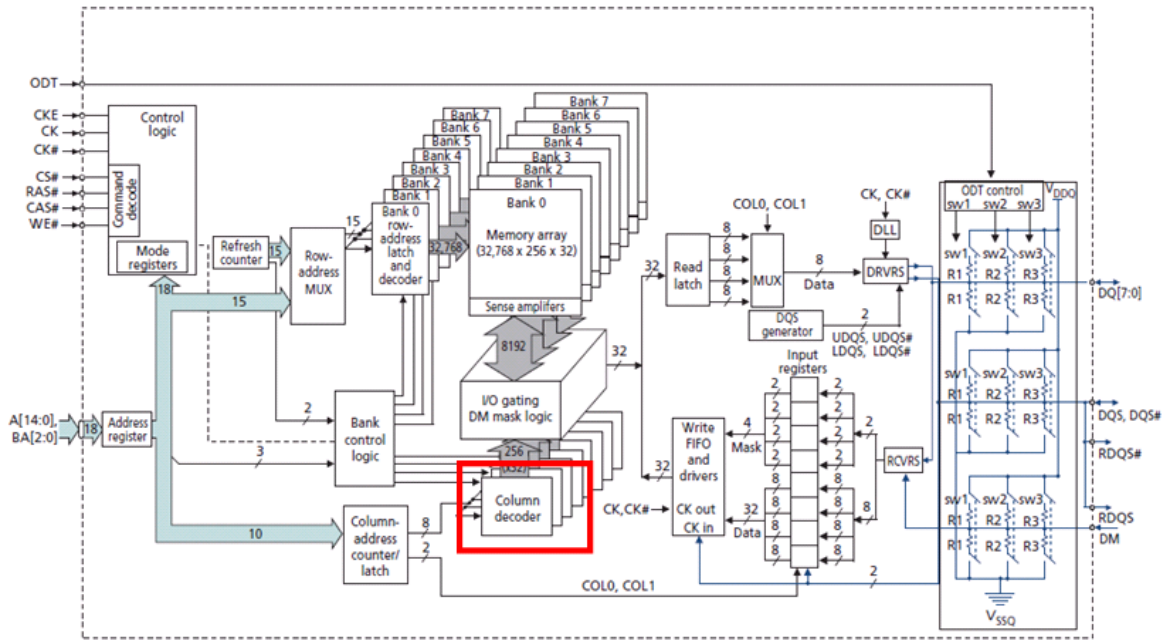
In the active state, the DDR3 device can perform READs and WRITEs. A READ command decodes a specific column address associated with the data that is stored in the sense amplifiers (shown in green in Figure 1). The data from this column is driven through the I/O, gating to the internal READ latch. From there, it is multiplexed onto the output drivers. The circuits used in this function are shown in purple in Figure 1.

Micron Technical Note TN-41-01 at 2 (annotations added).

853. For example, the Tegra 2 interfaces with DDR2 SDRAM. See, e.g., Tegra 2 TRM at 492.

854. For example, as shown in the figures below, Micron DDR2 SDRAM includes a column decoder that decodes a specific column upon a READ or WRITE command.

Figure 4: Functional Block Diagram – 256 Meg x 8



READ

The READ command is used to initiate a burst read access to an active row. The value on the bank address inputs determine the bank, and the address provided on address inputs A0–Ai (where Ai is the most significant column address bit for a given configuration) selects the starting column location. The value on input A10 determines whether or not auto precharge is used. If auto precharge is selected, the row being accessed will be precharged at the end of the read burst; if auto precharge is not selected, the row will remain open for subsequent accesses.

DDR2 SDRAM also supports the AL feature, which allows a READ or WRITE command to be issued prior to tRCD (MIN) by delaying the actual registration of the READ/WRITE command to the internal device by AL clock cycles.

WRITE

The WRITE command is used to initiate a burst write access to an active row. The value on the bank select inputs selects the bank, and the address provided on inputs A0–Ai (where Ai is the most significant column address bit for a given configuration) selects the starting column location. The value on input A10 determines whether or not auto precharge is used. If auto precharge is selected, the row being accessed will be precharged at the end of the WRITE burst; if auto precharge is not selected, the row will remain open for subsequent accesses.

DDR2 SDRAM also supports the AL feature, which allows a READ or WRITE command to be issued prior to tRCD (MIN) by delaying the actual registration of the READ/WRITE command to the internal device by AL clock cycles.

Input data appearing on the DQ is written to the memory array subject to the DM input logic level appearing coincident with the data. If a given DM signal is registered LOW, the corresponding data will be written to memory; if the DM signal is registered HIGH, the corresponding data inputs will be ignored, and a WRITE will not be executed to that byte/column location (see Figure 66 (page 112)).

DDR2 Datasheet at 13, 75 (annotations added).

855. As DDR2 and DDR3 SDRAMs are manufactured according to the JEDEC standard, all DDR2 and DDR3 SDRAM memory can and does operate in substantially the same way.

856. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 1(c) of the '938 Patent**.

857. "**Claim 1(d) of the '938 Patent**" recites "a column address input port for inputting a column address that selects the column of the memory bank."

858. The **'938 Accused Products** include a column address input port for inputting a column address that selects a column of the memory bank.

859. For example, DDR3 includes address inputs A0 - A15 to permit selection of the column address for Read/Write commands from a memory bank. *See, e.g., JEDEC Standard No. 79-3F* at 13.

860. For example, DDR2 includes address inputs A0 - A15 to permit selection of the column address for Read/Write commands from a memory bank. *See, e.g., JEDEC Standard No. 79-2F* at 13.

861. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 1(d) of the '938 Patent**.

862. "**Claim 1(e) of the '938 Patent**" recites "a first shift register for delaying the column address by a first number of delay clock cycles between the column address input port and the column decoder."

863. The **'938 Accused Products** include a shift register for delaying the column address by a number of delay clock cycles between the column address input port and the column decoder.

864. For example, DDR3 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 28.

865. For example, DDR2 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-2F* at 33.

866. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 1(e) of the '938 Patent**.

867. “**Claim 1(f) of the '938 Patent**” recites “a delay counter for sensing the number of clock cycles RCL of the clock signal from the application of the row access command to the application of the column access command with respect to the same bank, and for providing a first delay clock control signal to the first shift register.”

868. The **'938 Accused Products** include a delay counter for sensing the number of clock cycles RCL of the clock signal from the application of the row access command to the application of the column access command with respect to the same bank and for providing a delay clock control signal to the shift register.

869. For example, in DDR3 SDRAM, tRCD is the number of clock cycles between an active command row address command and an external CAS command. *See, e.g., JEDEC Standard No. 79-3F* at 157.

870. The DDR3 SDRAM delays the external CAS command by the Additive Latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 157.

871. For example, the receipt of the external CAS command triggers the Additive Latency delay mechanism. *See, e.g., JEDEC Standard No. 79-3F* at 157.

872. DDR3 holds a read or write command for the time of the additive latency. *See, e.g., JEDEC Standard No. 79-3F* at 28.

873. For example, in DDR2 SDRAM, tRCD is the difference between RL and CL, which is the number of clock cycles delay required between an active command row address strobe and a CAS. *See, e.g., JEDEC Standard No. 79-2F* at 33.

874. For example, the receipt of the external CAS command triggers the Additive Latency delay mechanism. *See, e.g., JEDEC Standard No. 79-2F* at 33.

875. DDR2 holds a read or write command for the time of the additive latency. *See, e.g., JEDEC Standard No. 79-2F* at 33.

876. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 1(f) of the '938 Patent**.

877. “**Claim 1(g) of the '938 Patent**” recites an SDRAM “wherein RLmin is the minimum number of clock cycles of the clock signal required from the application of a row access command to the output of the data of the memory and CLmin is the minimum number of clock cycles of the clock signal required from the application of a column access command to the output of the data of the memory cell.”

878. The **'938 Accused Products** include an SDRAM in which RLmin is the minimum number of clock cycles of the clock signal required from the application of a row access command to the output of the data of the memory.

879. The **'938 Accused Products** also include an SDRAM in which CLmin is the minimum number of clock cycles of the clock signal required from the application of a column access command to the output of the data of the memory cell.

880. For example, in DDR3 SDRAM, RLmin is equivalent to tAA plus tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

881. In DDR3 SDRAM, CL_{min} is equivalent to CL or tAA. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

882. For example, in DDR2 SDRAM, the difference between RL_{min} and CL_{min} is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

883. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 1(g) of the '938 Patent**.

884. “**Claim 1(h) of the '938 Patent**” recites an SDRAM “wherein the first delay clock control signal has information on the difference between RCL and (RL_{min}–CL_{min}), and the first number of delay clock cycles is determined in response to the difference between RCL and (RL_{min}–CL_{min}).”

885. The **'938 Accused Products** include an SDRAM in which the first delay clock control signal has information on the difference between RCL and (RL_{min}–CL_{min}).

886. The **'938 Accused Products** also include an SDRAM in which the first number of delay clock cycles is determined in response to the difference between RCL and (RL_{min}–CL_{min}).

887. For example, in DDR3 SDRAM, the difference between RL_{min} (equivalent to tAA plus tRCD) and CL_{min} (equivalent to CL or tAA) is tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

888. When the CAS command is received externally, the Additive Latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

889. The time between the externally supplied RAS command and the externally supplied CAS command (RCL) plus the Additive Latency must be equal to or greater than tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

890. Thus, the Additive Latency is determined based on the difference between RCL and (RLmin – CLmin). *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

891. For example, in DDR2 SDRAM, RLmin and CLmin are calculated and the difference between the two is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

892. When the CAS command is received externally, the additive latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

893. The time between the externally supplied RAS command and the externally supplied CAS command (RCL) plus the additive latency must be equal to or greater than tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

894. Thus, the additive latency is determined based on the difference between RCL and (RLmin – CLmin). *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

895. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 1(h) of the '938 Patent**.

896. **"Claim 2(a) of the '938 Patent"** recites "[a] synchronous DRAM, as recited in claim 1."

897. As described in the preceding paragraphs, the **'938 Accused Products** include an SDRAM as recited in claim 1.

898. **"Claim 2(b) of the '938 Patent"** recites "a plurality of registers serially coupled to each other for continuously transmitting the column address in response to the clock signal of every period."

899. The **'938 Accused Products** include a shift register that comprises a plurality of registers serially coupled to each other for continuously transmitting the column address in response to the clock signal of every period.

900. For example, DDR3 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 28.

901. For example, DDR2 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-2F* at 33.

902. Shift registers are typically implemented as a serial chain of registers, each of which transmits the data to the next register in each cycle.

903. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 2(b) of the '938 Patent**.

904. "**Claim 2(c) of the '938 Patent**" recites "a multiplexer for selectively providing one signal among the output signals of the plurality of registers to the column decoder."

905. The **'938 Accused Products** include a multiplexer that provides one signal among the output signals of the plurality of registers to the column decoder.

906. For example, DDR3 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 28.

907. For example, DDR2 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-2F* at 33.

908. The selection of a specific delay from a plurality of serial coupled registers requires a multiplexor and control logic to select the appropriate delay cycle.

909. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 2(c) of the '938 Patent**.

910. "**Claim 3 of the '938 Patent**" recites "[a] synchronous DRAM, as recited in claim 2, wherein the registers are D flip-flops."

911. As described in the preceding paragraphs, the **'938 Accused Products** include an SDRAM as recited in claim 2.

912. For example, a shift register is a group of flip-flops connected in a chain and driven by a common clock.

913. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 3 of the '938 Patent**.

914. "**Claim 8(a) of the '938 Patent**" recites "[a] synchronous DRAM (SDRAM) operating in synchronization with a clock signal."

915. The **'938 Accused Products** include an SDRAM that operates in synchronization with a clock signal.

916. The **'938 Accused GPUs** and **Accused SOCs** are designed to operate in conjunction with JEDEC standard SDRAM or SGRAM.

917. For example, the Tegra 3 is designed for use in conjunction with at least LP-DDR2 or DDR3 memory. *See, e.g., Tegra 3 TRM* at 10, 437.

918. For example, the Tegra 2 is designed for use in conjunction with at least LP-DDR, LP-DDR2, or DDR2 memory. *See, e.g., Tegra 2 TRM* at 8, 492.

919. For example, the Tegra 4 is designed for use in conjunction with at least DDR3, low-voltage DDR, LP-DDR2, and LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 12, 1115.

920. For example, the Tegra K1 is designed for use in conjunction with at least DDR3L and, LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 13, 643.

921. The SDRAM included in the **'938 Accused Products** complies with the relevant JEDEC standard.

922. JEDEC standard SDRAM operates synchronously with a clock signal.

923. As described in the preceding paragraphs, the **'938 Accused Products** include an SDRAM that operates in synchronization with a clock signal.

924. “**Claim 8(b) of the '938 Patent**” recites “a memory bank having a plurality of memory cells arranged in rows and columns.”

925. The **'938 Accused Products** include a memory bank having a plurality of memory cells arranged in rows and columns.

926. For example, DDR3 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-3F* at 15.

927. For example, DDR2 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-2F* at 15.

928. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 8(b) of the '938 Patent**.

929. “**Claim 8(c) of the '938 Patent**” recites “a column decoder for selecting a column of the memory bank.”

930. The **'938 Accused Products** include a column decoder for selecting a column of the memory bank.

931. For example, the Tegra 3 and Tegra 4 operate in conjunction with DDR3 SDRAM.

932. Micron DDR3 SDRAM includes a column decoder that decodes a specific column upon a READ command. *See, e.g., Micron Technical Note TN-41-01* at 2.

933. Micron DDR3 SDRAM complies with the JEDEC standard for DDR2 SDRAM.

934. For example, the Tegra 2 operates in conjunction with DDR2 SDRAM.

935. Micron DDR2 SDRAM includes a column decoder that decodes a specific column upon a READ or WRITE command. *See, e.g., DDR2 Datasheet* at 13, 75.

936. Micron DDR2 SDRAM complies with the JEDEC standard for DDR2 SDRAM.

937. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 8(c) of the '938 Patent**.

938. "**Claim 8(d) of the '938 Patent**" recites "a sense amplifier for amplifying data from the selected column."

939. The **'938 Accused Products** include a sense amplifier for amplifying data from the selected column.

940. For example, DDR3 includes internal sense amplifiers that amplify data from a column. *See, e.g., JEDEC Standard No. 79-3F* at 16.

941. For example, DDR2 includes internal sense amplifiers that amplify data from a column. *See, e.g., JEDEC Standard No. 79-2F* at 14.

942. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 8(d) of the '938 Patent**.

943. "**Claim 8(e) of the '938 Patent**" recites "a column address input port for inputting a column address for selecting the column of the memory bank."

944. The **'938 Accused Products** include a column address input port for inputting a column address for selecting the column of the memory bank.

945. For example, DDR3 includes address inputs A0 - A15 to permit selection of the column address for Read/Write commands from a memory bank. *See, e.g., JEDEC Standard No. 79-3F* at 13.

946. For example, DDR2 includes address inputs A0 - A15 to permit selection of the column address for Read/Write commands from a memory bank. *See, e.g., JEDEC Standard No. 79-2F* at 13.

947. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 8(e) of the '938 Patent**.

948. "**Claim 8(f) of the '938 Patent**" recites "a first shift register for delaying the column address by a first number of delay clock cycles between the column address input port and the column decoder."

949. The **'938 Accused Products** include a shift register for delaying the column address by a number of delay clock cycles between the column address input port and the column decoder.

950. For example, DDR3 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 28.

951. For example, DDR2 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-2F* at 33.

952. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 8(f) of the '938 Patent**.

953. "**Claim 8(g) of the '938 Patent**" recites "a delay counter for providing a first delay clock control signal having information on the difference between RCL and SAE to the first shift register."

954. The **'938 Accused Products** include a delay counter for providing a first delay clock control signal having information on the difference between RCL and SAE to the first shift register.

955. For example, DDR3 SDRAM delays the external CAS command by the Additive Latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 157.

956. The receipt of the external CAS command triggers the Additive Latency delay mechanism. *See, e.g., JEDEC Standard No. 79-3F* at 157.

957. DDR3 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 28.

958. For example, in DDR2, the number of clock cycles RCL of the clock signal from the application of the row access command to the application of the column access command with respect to the same bank is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

959. For example, DDR2 holds a read or write command for the time of the Additive Latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-2F* at 33.

960. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 8(g) of the '938 Patent**.

961. "**Claim 8(h) of the '938 Patent**" recites an SDRAM "wherein RCL is the number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the same bank."

962. The **'938 Accused Products** include an SDRAM wherein RCL is the number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the same bank.

963. For example, in DDR3 SDRAM, the difference between RLmin (equivalent to tAA plus tRCD) and CLmin (equivalent to CL or tAA) is tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

964. When the CAS command is received externally, the Additive Latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

965. The time between the externally supplied RAS command and the externally supplied CAS command is RCL, which when added to the Additive Latency must be equal to or greater than tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

966. For example, in DDR2 SDRAM, the difference between RLmin and CLmin is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

967. When the CAS command is received externally, the additive latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

968. The time between the externally supplied RAS command and the externally supplied CAS command (RCL) plus the Additive Latency must be equal to or greater than tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

969. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 8(h) of the '938 Patent**.

970. "**Claim 8(i) of the '938 Patent**" recites an SDRAM "wherein SAE is the number of clock cycles of the clock signal from the application of the row access command to the point of time at which the sense amplifier is enabled are determined."

971. The **'938 Accused Products** include an SDRAM wherein SAE is the number of clock cycles of the clock signal from application of the row access command to the point of time at which the sense amplifier is enabled are determined.

972. For example, in DDR2 SDRAM, SAE is equivalent to tRCD, which is the delay when converted to clock cycles from the application of the row access command to when the sense amplifier is enabled. *See, e.g., JEDEC Standard No. 79-2F* at 33.

973. For example, in DDR3 SDRAM, SAE is equivalent to tRCD, which is the delay when converted to clock cycles from the application of the row access command to when the sense amplifier is enabled. *See, e.g., JEDEC Standard No. 79-3F* at 28, 157.

974. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 8(i) of the '938 Patent**.

975. "**Claim 8(j) of the '938 Patent**" recites an SDRAM "wherein the first number of delay clock cycles is determined in response to the difference between RCL and SAE."

976. The **'938 Accused Products** include an SDRAM wherein the first number of delay clock cycles is determined in response to the difference between RCL and SAE.

977. For example, in DDR3 SDRAM, when the CAS command is received externally, the Additive Latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-3F* at 28, 157.

978. The time between the externally supplied RAS command and the externally supplied CAS command (RCL), plus the Additive Latency must be greater than tRCD (SAE). *See, e.g., JEDEC Standard No. 79-3F* at 28, 157.

979. Thus, the Additive Latency is determined based on the difference between RCL and SAE. *See, e.g., JEDEC Standard No. 79-3F* at 28, 157.

980. For example, in DDR2 SDRAM, the number of clock cycles RCL of the clock signal from the application of the row access command to the application of the column access command with respect to the same bank is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

981. When the CAS command is received externally, the additive latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

982. The time between the externally supplied RAS command and the externally supplied CAS command (RCL) plus the additive latency must be equal to or greater than tRCD (SAE). *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

983. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 8(j) of the '938 Patent**.

984. “**Claim 9(a) of the '938 Patent**” recites “[a] synchronous DRAM, as recited in claim 8.”

985. As described in the preceding paragraphs, the **'938 Accused Products** include an SDRAM as recited in claim 8 and thus meet the limitations of **Claim 9(a) of the '938 Patent**.

986. “**Claim 9(b) of the '938 Patent**” recites “a plurality of registers serially coupled to each other, for continuously transmitting the column address every cycle of the clock signal.”

987. The **'938 Accused Products** include a plurality of registers serially coupled to each other, for continuously transmitting the column address every cycle of the clock signal.

988. For example, DDR3 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 28.

989. For example, DDR2 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-2F* at 33.

990. Shift registers are typically implemented as a serial chain of registers, each of which transmits the data to the next register in each cycle.

991. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 9(b) of the '938 Patent**.”

992. “**Claim 9(c) of the '938 Patent**” recites “a multiplexer for selectively providing one signal among the output signals of the registers to the column decoder in response to the difference between RCL and SAE.”

993. The **'938 Accused Products** include a multiplexer for selectively providing one signal among the output signals of the registers to the column decoder in response to the difference between RCL and SAE.

994. For example, DDR3 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 28.

995. For example, DDR2 holds a read or write command for the time of the additive latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-2F* at 33.

996. The selection of a specific delay from a plurality of serial coupled registers requires a multiplexor and control logic to select the appropriate delay cycle.

997. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 9(c) of the '938 Patent.**”

998. “**Claim 10 of the '938 Patent**” recites “[a] synchronous DRAM, as recited in claim 9, wherein the registers are D flip-flops.”

999. As described in preceding paragraphs, the **'938 Accused Products** include an SDRAM as recited in claim 9.

1000. The **'938 Accused Products** include an SDRAM where the registers are D flip flops.

1001. For example, a shift register is a group of flip-flops connected in a chain and driven by a common clock.

1002. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 10 of the '938 Patent**.

1003. “**Claim 17 of the '938 Patent**” recites “[a] synchronous DRAM, as recited in claim 8, wherein the first delay clock signal is provided from outside of the SDRAM.”

1004. As described in preceding paragraphs, the **'938 Accused Products** include an SDRAM as recited in claim 8.

1005. The **'938 Accused Products** include an SDRAM, as recited in claim 8, wherein the first delay clock signal is provided from outside of the SDRAM.

1006. For example, DDR3 SDRAM requires Additive Latency to be set per the values programmed in the MR0 register. *See, e.g., JEDEC Standard No. 79-3F* at 28.

1007. For example, DDR2 SDRAM requires Additive Latency to be set per the values programmed in the Extended Mode Register. *See, e.g., JEDEC Standard No. 79-2F* at 33-34.

1008. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 17 of the '938 Patent**.

1009. “**Claim 18(a) of the '938 Patent**” recites “[a] synchronous DRAM (SDRAM) synchronized with a clock signal after predetermined column access strobe (CAS) latency has lapsed from a column access command.”

1010. For example, the Tegra 3 is designed for use in conjunction with at least LP-DDR2 or DDR3 memory. *See, e.g., Tegra 3 TRM* at 10, 437.

1011. For example, the Tegra 2 is designed for use in conjunction with at least LP-DDR, LP-DDR2, or DDR2 memory. *See, e.g., Tegra 2 TRM* at 8, 492.

1012. For example, the Tegra 4 is designed for use in conjunction with at least DDR3, low-voltage DDR, LP-DDR2, and LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 12, 1115.

1013. For example, the Tegra K1 is designed for use in conjunction with at least DDR3L and, LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 13, 643.

1014. As described in the preceding paragraphs, the **'938 Accused Products** include an SDRAM synchronized with a clock signal after predetermined column access strobe (CAS) latency has lapsed from a column access command.

1015. “**Claim 18(b) of the '938 Patent**” recites “a memory bank having a plurality of memory cells arranged in rows and columns.”

1016. The **'938 Accused Products** include a memory bank having a plurality of memory cells arranged in rows and columns.

1017. For example, DDR3 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-3F* at 15.

1018. For example, DDR2 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-2F* at 15.

1019. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 18(b) of the '938 Patent**.

1020. “**Claim 18(c) of the '938 Patent**” recites “a decoder for selecting one of the memory cells based on a column address and a row address.”

1021. The **'938 Accused Products** include a decoder for selecting one of the memory cells based on a column address and a row address.

1022. For example, the Tegra 3 and Tegra 4 operate in conjunction with DDR3 SDRAM.

1023. Micron DDR3 SDRAM includes a column decoder that decodes a specific column upon a READ command. *See, e.g., Micron Technical Note TN-41-01* at 2.

1024. For example, the Tegra 2 operates in conjunction with DDR2 SDRAM.

1025. Micron DDR2 SDRAM includes a column decoder that decodes a specific column upon a READ or WRITE command. *See, e.g., DDR2 Datasheet* at 13, 75.

1026. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 18(c) of the '938 Patent**.

1027. "**Claim 18(d) of the '938 Patent**" recites an SDRAM "wherein the CAS latency is determined by the number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the memory bank."

1028. The **'938 Accused Products** include an SDRAM in which the CAS latency is determined by the number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the memory bank.

1029. For example, CAS latency in DDR3 SDRAM is referred to as Read Latency (RL), which is the number of clock cycles of the clock signal from the application of an external column access command to availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1030. The JEDEC standard defines Read Latency as Additive Latency plus "CAS Latency." where "CAS Latency" refers to the number of clock cycles between the internal read command and the availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1031. Additive Latency is determined based on the number of clock cycles between the application of a row address command and the application of an external CAS command. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1032. For example, in DDR2, the JEDEC standard defines Read Latency as Additive Latency plus “CAS Latency.” *See, e.g., JEDEC Standard No. 79-2F* at 33.

1033. Additive Latency is determined based on the number of clock cycles between the application of a row address command and the application of an external CAS command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1034. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 18(d) of the '938 Patent**.

1035. “**Claim 19(a) of the '938 Patent**” recites “[a] synchronous DRAM (SDRAM).”

1036. The **'938 Accused Products** include an SDRAM.

1037. For example, the Tegra 3 is designed for use in conjunction with at least LP-DDR2 or DDR3 memory. *See, e.g., Tegra 3 TRM* at 10, 437.

1038. For example, the Tegra 2 is designed for use in conjunction with at least LP-DDR, LP-DDR2, or DDR2 memory. *See, e.g., Tegra 2 TRM* at 8, 492.

1039. For example, the Tegra 4 is designed for use in conjunction with at least DDR3, low-voltage DDR, LP-DDR2, and LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 12, 1115.

1040. For example, the Tegra K1 is designed for use in conjunction with at least DDR3L and, LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 13, 643.

1041. As described in the preceding paragraphs, the **'938 Accused Products include an SDRAM**.

1042. “**Claim 19(b) of the '938 Patent**” recites “a memory bank having a plurality of memory cells arranged in rows and columns.”

1043. The **'938 Accused Products** include a memory bank having a plurality of memory cells arranged in rows and columns.

1044. For example, DDR3 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-3F* at 15.

1045. For example, DDR2 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-2F* at 15.

1046. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 19(b) of the '938 Patent**.

1047. “**Claim 19(c) of the '938 Patent**” recites “a decoder for selecting one of the memory cells based on a column address and a row address.”

1048. The **'938 Accused Products** include a decoder for selecting one of the memory cells based on a column address and a row address.

1049. For example, the Tegra 3 and Tegra 4 operate in conjunction with DDR3 SDRAM.

1050. For example, Micron DDR3 SDRAM includes a column decoder that decodes a specific column upon a READ command. *See, e.g., Micron Technical Note TN-41-01* at 2.

1051. For example, the Tegra 2 operates in conjunction with DDR2 SDRAM.

1052. For example, Micron DDR2 SDRAM includes a column decoder that decodes a specific column upon a READ or WRITE command. *See, e.g., DDR2 Datasheet* at 13, 75.

1053. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 19(c) of the '938 Patent**.

1054. “**Claim 19(d) of the '938 Patent**” recites an SDRAM “wherein RLmin is the minimum number of clock cycles of the clock signal required from the application of a row access command to the output of the data of the selected memory cell.”

1055. The **'938 Accused Products** include an SDRAM in which RLmin is the minimum number of clock cycles of the clock signal required from the application of a row access command to the output of the data of the selected memory cell.

1056. For example, in DDR3 SDRAM, RLmin is equivalent to tAA plus tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1057. For example, in DDR2, the difference between the RLmin and CLmin is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1058. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 19(d) of the '938 Patent**.

1059. “**Claim 19(e) of the '938 Patent**” recites an SDRAM “wherein CLmin is the minimum number of clock cycles of the clock signal required from the application of a column access command to the output of the data of the selected memory cell.”

1060. The **'938 Accused Products** include an SDRAM in which CLmin is the minimum number of clock cycles of the clock signal required from the application of a column access command to the output of the data of the selected memory cell.

1061. For example, in DDR3 SDRAM, CLmin is equivalent to CL or tAA. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1062. For example, in DDR2, the difference between the RLmin and CLmin is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1063. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 19(e) of the '938 Patent**.

1064. “**Claim 19(f) of the '938 Patent**” recites an SDRAM “wherein RCL is the number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the memory bank.”

1065. The **'938 Accused Products** include an SDRAM wherein RCL is the number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the memory bank.

1066. For example, in DDR3 SDRAM, the difference between RLmin (equivalent to tAA plus tRCD) and CLmin (equivalent to CL or tAA) is tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1067. When the CAS command is received externally, the Additive Latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1068. For example, in DDR2 SDRAM, the number of clock cycles RCL of the clock signal from the application of the row access command to the application of the column access command with respect to the same bank is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1069. When the CAS command is received externally, the additive latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

1070. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 19(f) of the '938 Patent**.

1071. “**Claim 19(g) of the '938 Patent**” recites an SDRAM “wherein a CAS latency, which is the number of clock cycles of the clock signal required from the application of the column access command to the output of data, is determined to be (RLmin–RCL) when RCL is

less than $(RL_{min}-CL_{min})$, and is determined to be CL_{min} when RCL is not less than $(RL_{min}-CL_{min})$.”

1072. The **'938 Accused Products** include an SDRAM in which a CAS latency, which is the number of clock cycles of the clock signal required from the application of the column access command to the output of data, is determined to be $(RL_{min}-RCL)$ when RCL is less than $(RL_{min}-CL_{min})$, and is determined to be CL_{min} when RCL is not less than $(RL_{min}-CL_{min})$.

1073. For example, CAS latency in DDR3 SDRAM is referred to as Read Latency (RL), which is the number of clock cycles of the clock signal from the application of an external column access command to the availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1074. The JEDEC standard defines Read Latency as Additive Latency plus “CAS Latency,” where “CAS Latency” refers to the number of clock cycles between the internal Read command and the availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1075. When RCL is less than $(RL_{min} - CL_{min})$, Additive Latency is greater than zero, in which case the Read Latency is defined to be AL plus CL , which is equivalent to AL plus tAA , which is equivalent to $(tAA + tRCD) - (tRCD - AL)$, or $(RL_{min} - RCL)$. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1076. When RCL is not less than $(RL_{min} - CL_{min})$, Additive Latency is equal to zero, in which case Read Latency is defined to be CL , which is equivalent to tAA , or CL_{min} . *See, e.g., JEDEC Standard No. 79-3F* at 25.

1077. For example, in DDR2, the JEDEC standard defines Read Latency as Additive Latency plus “CAS Latency.” *See, e.g., JEDEC Standard No. 79-2F* at 33.

1078. When RCL is less than $(RL_{min} - Cl_{min})$, Additive Latency is greater than zero, in which case the Read Latency is defined to be AL plus CL. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1079. When RCL is not less than $(RL_{min} - Cl_{min})$, Additive Latency is equal to zero, in which case Read Latency is defined to be CL. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1080. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 19(g) of the '938 Patent**.

1081. "**Claim 20 of the '938 Patent**" recites "[a] synchronous DRAM, as recited in claim 19, wherein $(RL_{min} - CL_{min})$ is input from the outside of the SDRAM."

1082. As described in preceding paragraphs, the **'938 Accused Products** include an SDRAM as recited in claim 19.

1083. The **'938 Accused Products** include an SDRAM, as recited in claim 19, wherein $(RL_{min} - CL_{min})$ is input from the outside of the SDRAM.

1084. For example, DDR3 SDRAM requires Additive Latency to be set per the values programmed in the MR0 register. *See, e.g., JEDEC Standard No. 79-3F* at 28.

1085. For example, DDR2 SDRAM requires Additive Latency to be set per the values programmed in the Extended Mode Register. *See, e.g., JEDEC Standard No. 79-2F* at 33-34.

1086. For example, tRCD is input from outside the SDRAM by, for example, reading the Boot Configuration Tables from flash memory onto on-storage RAM. *See, e.g., Tegra 3 TRM* at 444-45.

1087. For example, tRCD is input from outside the SDRAM by, for example, reading the Boot Configuration Tables onto on-storage RAM. *See, e.g., Tegra 2 TRM* at 519.

1088. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 20 of the '938 Patent**.

1089. "**Claim 21(a) of the '938 Patent**" recites "[a] synchronous DRAM (SDRAM) operating in synchronization with a clock signal."

1090. The **'938 Accused Products** include an SDRAM operating in synchronization with a clock signal.

1091. For example, the Tegra 3 is designed for use in conjunction with at least LP-DDR2 or DDR3 memory. *See, e.g., Tegra 3 TRM* at 10, 437.

1092. For example, the Tegra 2 is designed for use in conjunction with at least LP-DDR, LP-DDR2, or DDR2 memory. *See, e.g., Tegra 2 TRM* at 8, 492.

1093. For example, the Tegra 4 is designed for use in conjunction with at least DDR3, low-voltage DDR, LP-DDR2, and LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 12, 1115.

1094. For example, the Tegra K1 is designed for use in conjunction with at least DDR3L and, LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 13, 643.

1095. As described in the preceding paragraphs, the **'938 Accused Products** include an SDRAM operating in synchronization with a clock signal.

1096. "**Claim 21(b) of the '938 Patent**" recites "a memory bank having a plurality of memory cells arranged in rows and columns."

1097. The **'938 Accused Products** include a memory bank having a plurality of memory cells arranged in rows and columns.

1098. For example, DDR3 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-3F* at 15.

1099. For example, DDR2 includes memory banks that include rows and columns. *See, e.g., JEDEC Standard No. 79-2F* at 15.

1100. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 21(b) of the '938 Patent**.

1101. "**Claim 21(c) of the '938 Patent**" recites "a column decoder for selecting the column of the memory bank."

1102. The **'938 Accused Products** include a column decoder for selecting the column of the memory bank.

1103. For example, the Tegra 3 and Tegra 4 operate in conjunction with DDR3 SDRAM.

1104. For example, Micron DDR3 SDRAM includes a column decoder that decodes a specific column upon a READ command. *See, e.g., Micron Technical Note TN-41-01* at 2.

1105. For example, the Tegra 2 operates in conjunction with DDR2 SDRAM.

1106. For example, Micron DDR2 SDRAM includes a column decoder that decodes a specific column upon a READ or WRITE command. *See, e.g., DDR2 Datasheet* at 13, 75.

1107. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 21(c) of the '938 Patent**.

1108. "**Claim 21(d) of the '938 Patent**" recites "a sense amplifier for amplifying data from the selected column."

1109. The **'938 Accused Products** include a sense amplifier for amplifying data from the selected column.

1110. For example, DDR3 includes internal sense amplifiers that amplify data from a column. *See, e.g., JEDEC Standard No. 79-3F* at 16.

1111. For example, DDR2 includes internal sense amplifiers that amplify data from a column. *See, e.g., JEDEC Standard No. 79-2F* at 14.

1112. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 21(d) of the '938 Patent**.

1113. "**Claim 21(e) of the '938 Patent**" recites "wherein RLmin is the minimum number of clock cycles of the clock signal required from the application of a row access command to the output of the data of the selected memory cell."

1114. The **'938 Accused Products** include an SDRAM where RLmin is the minimum number of clock cycles of the clock signal required from the application of a row access command to the output of the data of the selected memory cell.

1115. For example, in DDR3 SDRAM, RLmin is equivalent to tAA plus tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1116. For example, in DDR2, RCL is the difference between RLmin and CLmin, which is the number of clock cycles delay required between an active command row address strobe and a CAS. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1117. This is known as tRCD, which are the cycles required between the memory controller asserting a row address, and then asserting a column address during the subsequent read or write command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1118. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 21(e) of the '938 Patent**.

1119. "**Claim 21(f) of the '938 Patent**" recites an SDRAM "wherein CLmin is the minimum number of clock cycles of the clock signal required from the application of a column access command to the output of the data of the selected memory cell."

1120. The **'938 Accused Products** includes an SDRAM in which CLmin is the minimum number of clock cycles of the clock signal required from the application of a column access command to the output of the data of the selected memory cell.

1121. For example, in DDR3 SDRAM, CLmin is equivalent to CL or tAA. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1122. For example, in DDR2, RCL is the difference between RLmin and CLmin, which is the number of clock cycles delay required between an active command row address strobe and a CAS. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1123. This is known as tRCD, which are the cycles required between the memory controller asserting a row address, and then asserting a column address during the subsequent read or write command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1124. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 21(f) of the '938 Patent**.

1125. "**Claim 21(g) of the '938 Patent**" recites an SDRAM wherein "RCL is the number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the memory bank."

1126. The **'938 Accused Products** include an SDRAM in which RCL is the number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the memory bank.

1127. For example, in DDR3 SDRAM, the difference between RLmin (equivalent to tAA plus tRCD) and CLmin (equivalent to CL or tAA) is tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1128. When the CAS command is received externally, the Additive Latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1129. The time between the externally supplied RAS command and the externally supplied CAS command is RCL, which when added to the Additive Latency must be equal to or greater than tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1130. For example, in DDR2 SDRAM, the number of clock cycles RCL of the clock signal from the application of the row access command to the application of the column access command with respect to the same bank is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1131. When the CAS command is received externally, the additive latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

1132. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 21(g) of the '938 Patent**.

1133. "**Claim 21(h) of the '938 Patent**" recites an SDRAM wherein "SAE is the number of clock cycles of the clock signal from the application of the row access command to the point of time at which the sense amplifier is enabled."

1134. The **'938 Accused Products** include an SDRAM in which SAE is the number of clock cycles of the clock signal from the application of the row access command to the point of time at which the sense amplifier is enabled.

1135. For example, in DDR2 SDRAM, SAE is equivalent to tRCD, which is the delay when converted to clock cycles from the application of the row access command to when the sense amplifier is enabled. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1136. For example, in DDR3 SDRAM, SAE is equivalent to tRCD, which is the delay when converted to clock cycles from the application of the row access command to when the sense amplifier is enabled. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1137. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 21(h) of the '938 Patent**.

1138. "**Claim 21(i) of the '938 Patent**" recites an SDRAM "wherein a CAS latency, which is the number of clock cycles of the clock signal required from the application of the column access command to the output of data, is determined by the difference between RCL and SAE."

1139. The **'938 Accused Products** include an SDRAM in which a CAS latency, which is the number of clock cycles of the clock signal required from the application of the column access command to the output of data, is determined by the difference between RCL and SAE.

1140. For example, CAS latency in DDR3 SDRAM is referred to as Read Latency (RL), which is the number of clock cycles of the clock signal from the application of an external column access command to the availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1141. The JEDEC standard defines Read Latency as Additive Latency plus "CAS Latency," where "CAS Latency" refers to the number of clock cycles between the internal Read command and the availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1142. When the CAS command is received externally, the Additive Latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1143. The time between the externally supplied RAS command and the externally supplied CAS command (RCL), plus the Additive Latency must be greater than tRCD (SAE). *See, e.g., JEDEC Standard No. 79-3F* at 25.

1144. Thus, the Additive Latency is determined based on the difference between RCL and SAE. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1145. For example, in DDR2, the JEDEC standard defines Read Latency as Additive Latency plus “CAS Latency.” *See, e.g., JEDEC Standard No. 79-2F* at 33.

1146. When the CAS command is received externally, the Additive Latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1147. The time between the externally supplied RAS command and the externally supplied CAS command (RCL), plus the Additive Latency must be greater than tRCD (SAE). *See, e.g., JEDEC Standard No. 79-3F* at 33.

1148. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 21(i) of the '938 Patent**.

1149. “**Claim 23(a) of the '938 Patent**” recites “[a] method of controlling CAS latency of an SDRAM, synchronized with a clock signal, that includes a memory bank having a plurality of memory cells arranged in rows and columns and outputs the data of a selected memory cell.”

1150. The **'938 Accused Products** use a method of controlling CAS latency of an SDRAM, synchronized with a clock signal, that includes a memory bank having a plurality of memory cells arranged in rows and columns and outputs the data of a selected memory cell.

1151. For example, the Tegra 3 is designed for use in conjunction with at least LP-DDR2 or DDR3 memory. *See, e.g., Tegra 3 TRM* at 10, 437.

1152. For example, the Tegra 2 is designed for use in conjunction with at least LP-DDR, LP-DDR2, or DDR2 memory. *See, e.g., Tegra 2 TRM* at 8, 492.

1153. For example, the Tegra 4 is designed for use in conjunction with at least DDR3, low-voltage DDR, LP-DDR2, and LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 12, 1115.

1154. For example, the Tegra K1 is designed for use in conjunction with at least DDR3L and, LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 13, 643.

1155. For example, DDR3 includes memory banks that include rows and columns and that output the data of a memory cell. *See, e.g., JEDEC Standard No. 79-3F* at 15.

1156. For example, DDR2 includes memory banks that include rows and columns and that output the data of a memory cell. *See, e.g., JEDEC Standard No. 79-2F* at 15.

1157. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 23(a) of the '938 Patent**.

1158. “**Claim 23(b) of the '938 Patent**” recites “inputting a quantity (RLmin–CLmin) from the outside of the SDRAM, where RLmin is the minimum number of clock cycles of the clock signal required from the application of a row access command to the output of the data of the selected memory cell, and CLmin is the minimum number of clock cycles of the clock signal required from the application of a column access command to the output of the data of the selected memory cell.”

1159. The **'938 Accused Products** input a quantity (RLmin–CLmin) from the outside of the SDRAM, where RLmin is the minimum number of clock cycles of the clock signal required from the application of a row access command to the output of the data of the selected memory cell, and CLmin is the minimum number of clock cycles of the clock signal required

from the application of a column access command to the output of the data of the selected memory cell.

1160. For example, in DDR3 SDRAM, the difference between RLmin (equivalent to tAA plus tRCD) and CLmin (equivalent to CL or tAA) is tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1161. For example, in DDR2, the difference between the RLmin and CLmin is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1162. For example, tRCD is input from outside the SDRAM by, for example, reading the Boot Configuration Tables from flash memory onto on-storage RAM. *See, e.g., Tegra 3 TRM* at 444-45.

1163. For example, tRCD is input from outside the SDRAM by, for example, reading the Boot Configuration Tables onto on-storage RAM. *See, e.g., Tegra 2 TRM* at 519.

1164. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 23(b) of the '938 Patent**.

1165. "**Claim 23(c) of the '938 Patent**" recites "comparing RCL with (RLmin–CLmin), where RCL is a number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the memory bank."

1166. The **'938 Accused Products** compare RCL with (RLmin–CLmin), where RCL is a number of clock cycles of the clock signal from the application of a row access command to the application of a column access command with respect to the memory bank.

1167. For example, in DDR3 SDRAM, the difference between RLmin (equivalent to tAA plus tRCD) and CLmin (equivalent to CL or tAA) is tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1168. RCL is the number of clock cycles between the application of a row address command and an external CAS command. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1169. The DDR3 SDRAM delays the external CAS command by the Additive Latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1170. The receipt of the external CAS command triggers the Additive Latency delay mechanism. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1171. For example, in DDR2, the difference between the RLmin and CLmin is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1172. When the CAS command is received externally, the additive latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

1173. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 23(c) of the '938 Patent**.

1174. “**Claim 23(d) of the '938 Patent**” recites “determining CAS latency, which is the number of clock cycles of the clock signal required from the application of the column access command to the output of the data, to be (RLmin–RCL) when RCL is less than (RLmin–CLmin).”

1175. The **'938 Accused Products** determine CAS latency, which is the number of clock cycles of the clock signal required from the application of the column access command to the output of the data, to be (RLmin–RCL) when RCL is less than (RLmin–CLmin).

1176. For example, CAS latency in DDR3 SDRAM is referred to as Read Latency (RL), which is the number of clock cycles of the clock signal from the application of an external column access command to the availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1177. The JEDEC standard defines Read Latency as Additive Latency plus “CAS Latency,” where “CAS Latency” refers to the number of clock cycles between the internal Read command and the availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1178. When RCL is less than $(RL_{min} - Cl_{min})$, Additive Latency is greater than zero, in which case the Read Latency is defined to be AL plus CL, which is equivalent to AL plus t_{AA} , which is equivalent to $(t_{AA} + t_{RCD}) - (t_{RCD} - AL)$, or $(RL_{min} - RCL)$. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1179. For example, in DDR2, the difference between the RL_{min} and CL_{min} is t_{RCD} . *See, e.g., JEDEC Standard No. 79-2F* at 33.

1180. For example, in DDR2, the JEDEC standard defines Read Latency as Additive Latency plus “CAS Latency.” *See, e.g., JEDEC Standard No. 79-2F* at 33.

1181. When RCL is less than $(RL_{min} - Cl_{min})$, Additive Latency is greater than zero, in which case the Read Latency is defined to be AL plus CL. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1182. When RCL is not less than $(RL_{min} - Cl_{min})$, Additive Latency is equal to zero, in which case Read Latency is defined to be CL. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1183. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 23(d) of the '938 Patent**.

1184. “**Claim 23(e) of the '938 Patent**” recites “determining the CAS latency to be CL_{min} when RCL is no less than (RL_{min}–CL_{min}).”

1185. The **'938 Accused Products** determine the CAS latency to be CL_{min} when RCL is no less than (RL_{min}–CL_{min}).

1186. For example, when RCL is not less than (RL_{min} – Cl_{min}), Additive Latency is equal to zero, in which case Read Latency is defined to be CL, which is equivalent to tAA, or CL_{min}. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1187. For example, in DDR2, the difference between the RL_{min} and CL_{min} is tRCD. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1188. For example, when RCL is not less than (RL_{min} – Cl_{min}), Additive Latency is equal to zero, in which case Read Latency is defined to be CL. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1189. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 23(e) of the '938 Patent**.

1190. “**Claim 24(a) of the '938 Patent**” recites “[a] method of controlling CAS latency of an SDRAM which includes a bank having a plurality of memory cells arranged in rows and columns that outputs the data of a selected memory cell in synchronization with the clock signal.”

1191. The **'938 Accused Products** use a method controlling CAS latency of an SDRAM which includes a bank having a plurality of memory cells arranged in rows and columns that outputs the data of a selected memory cell in synchronization with the clock signal.

1192. For example, the Tegra 3 is designed for use in conjunction with at least LP-DDR2 or DDR3 memory. *See, e.g., Tegra 3 TRM* at 10, 437.

1193. For example, the Tegra 2 is designed for use in conjunction with at least LP-DDR, LP-DDR2, or DDR2 memory. *See, e.g., Tegra 2 TRM* at 8, 492.

1194. For example, the Tegra 4 is designed for use in conjunction with at least DDR3, low-voltage DDR, LP-DDR2, and LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 12, 1115.

1195. For example, the Tegra K1 is designed for use in conjunction with at least DDR3L and, LP-DDR3 memory. *See, e.g., Tegra 4 TRM* at 13, 643.

1196. For example, DDR3 includes memory banks that include rows and columns and that output the data of a memory cell. *See, e.g., JEDEC Standard No. 79-3F* at 15.

1197. For example, DDR2 includes memory banks that include rows and columns and that output the data of a memory cell. *See, e.g., JEDEC Standard No. 79-2F* at 15.

1198. As described in the preceding paragraphs, the **'938 Accused Products** use a method of controlling CAS latency of an SDRAM which includes a bank having a plurality of memory cells arranged in rows and columns that outputs the data of a selected memory cell in synchronization with the clock signal.

1199. “**Claim 24(b) of the '938 Patent**” recites “sensing RCL, where RCL is the number of clock cycles of the clock signal from an application of a row access command to an application of a column access command.”

1200. The **'938 Accused Products** senses RCL, where RCL is the number of clock cycles of the clock signal from an application of a row access command to an application of a column access command.

1201. For example, DDR3 senses the number of clock cycles RCL of the clock signal from the application of the row access command to the external application of the column access command with respect to the same bank. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1202. The difference between RLmin (equivalent to tAA plus tRCD) and CLmin (equivalent to CL or tAA) is tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1203. When the CAS command is received externally, the Additive Latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1204. The time between the externally supplied RAS command and the externally supplied CAS command is RCL, which when added to the Additive Latency must be equal to or greater than tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1205. For example, in DDR2, RCL is the difference between RLmin and CLmin, which is the number of clock cycles delay required between an active command row address strobe and a CAS. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1206. This is known as tRCD, which are the cycles required between the memory controller asserting a row address, and then asserting a column address during the subsequent read or write command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1207. When the CAS command is received externally, the additive latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

1208. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 24(b) of the '938 Patent**.

1209. "**Claim 24(c) of the '938 Patent**" recites "sensing SAE, where SAE is the number of clock cycles of the clock signal from application of the row access command to a point of time at which a sense amplifier is enabled."

1210. The **'938 Accused Products** sense the number of clock cycles of the clock signal from application of the row access command to a point of time at which a sense amplifier is enabled.

1211. For example, in DDR3 SDRAM, SAE is equivalent to tRCD, which is the delay when converted to clock cycles from the application of the row access command to when the sense amplifier is enabled. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1212. For example, in DDR2, RCL is the difference between RLmin and CLmin, which is the number of clock cycles delay required between an active command row address strobe and a CAS.

1213. This is known as tRCD (SAE), which are the cycles required between the memory controller asserting a row address, and then asserting a column address during the subsequent read or write command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1214. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 24(c) of the '938 Patent**.

1215. "**Claim 24(d) of the '938 Patent**" recites "comparing RCL with SAE."

1216. The **'938 Accused Products** compare RCL with SAE.

1217. For example, in DDR3 SDRAM, the difference between RLmin (equivalent to tAA plus tRCD) and CLmin (equivalent to CL or tAA) is tRCD, which is equivalent to SAE (in clock cycles). *See, e.g., JEDEC Standard No. 79-3F* at 157.

1218. RCL is the number of clock cycles between the application of a row address command and an external CAS command. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1219. The DDR3 SDRAM delays the external CAS command by the Additive Latency before it is issued inside the device. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1220. The receipt of the external CAS command triggers the Additive Latency delay mechanism. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1221. For example, in DDR2, RCL is the difference between RL_{min} and CL_{min}, which is the number of clock cycles delay required between an active command row address strobe and a CAS. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1222. This is known as tRCD, which are the cycles required between the memory controller asserting a row address, and then asserting a column address during the subsequent read or write command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1223. When the CAS command is received externally, the additive latency is added before the command is issued internally. *See, e.g., JEDEC Standard No. 79-2F* at 33, 78.

1224. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 24(d) of the '938 Patent**.

1225. “**Claim 24(e) of the '938 Patent**” recites “determining CAS latency, which is the number of clock cycles of the clock signal required from the application of the column access command to the output of the data, to be (RL_{min}–RCL) when RCL is less than SAE and the difference between RCL and SAE is not less than a predetermined number of reference clock cycles.”

1226. The **'938 Accused Products** determine CAS latency, which is the number of clock cycles of the clock signal required from the application of the column access command to the output of the data, to be (RL_{min}–RCL) when RCL is less than SAE and the difference between RCL and SAE is not less than a predetermined number of reference clock cycle.

1227. For example, CAS latency in DDR3 SDRAM is referred to as Read Latency (RL), which is the number of clock cycles of the clock signal from the application of an external column access command to the availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1228. The JEDEC standard defines Read Latency as Additive Latency plus “CAS Latency,” where “CAS Latency” refers to the number of clock cycles between the internal Read command and the availability of the first bit of output data. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1229. When RCL is less than SAE (tRCD), Additive Latency is greater than zero, in which case the Read Latency is defined to be AL plus CL, which is equivalent to AL plus tAA, which is equivalent to $(tAA + tRCD) - (tRCD - AL)$, or $(RL_{min} - RCL)$. *See, e.g., JEDEC Standard No. 79-3F* at 25.

1230. For example, in DDR2, RCL is the difference between RL_{min} and CL_{min} , which is the number of clock cycles delay required between an active command row address strobe and a CAS. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1231. This is known as tRCD (SAE), which are the cycles required between the memory controller asserting a row address, and then asserting a column address during the subsequent read or write command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1232. When RCL is less than SAE (tRCD), Additive Latency is greater than zero, in which case the Read Latency is defined to be AL plus CL. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1233. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 24(e) of the '938 Patent**.

1234. “**Claim 24(f) of the '938 Patent**” recites “determining the CAS latency to be CL_{min} when RCL is not less than SAE or the difference between RCL and SAE is less than the predetermined number of reference clock cycles.”

1235. The **'938 Accused Products** determine the CAS latency to be CLmin when RCL is not less than SAE or the difference between RCL and SAE is less than the predetermined number of reference clock cycles.

1236. For example, in DDR3, when RCL is not less than SAE (tRCD), Additive Latency is equal to zero, in which case Read Latency is defined to be CL, which is equivalent to tAA, or CLmin. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1237. For example, in DDR2, RCL is the difference between RLmin and CLmin, which is the number of clock cycles delay required between an active command row address strobe and a CAS. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1238. This is known as tRCD (SAE), which are the cycles required between the memory controller asserting a row address, and then asserting a column address during the subsequent read or write command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1239. For example, when RCL is not less than SAE (tRCD), Additive Latency is equal to zero, in which case Read Latency is defined to be CL. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1240. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 24(f) of the '938 Patent**.

1241. "**Claim 24(g) of the '938 Patent**" recites "wherein RLmin is the minimum number of clock cycles of a clock signal required from the application of a row access command to the output of the data of the selected memory cell."

1242. In the **'938 Accused Products**, RLmin is the minimum number of clock cycles of a clock signal required from the application of a row access command to the output of the data of the selected memory cell.

1243. For example, in DDR3 SDRAM, RLmin is equivalent to tAA plus tRCD. *See, e.g., JEDEC Standard No. 79-3F* at 25, 157.

1244. For example, in DDR2, RCL is the difference between RLmin and CLmin, which is the number of clock cycles delay required between an active command row address strobe and a CAS. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1245. This is known as tRCD, which are the cycles required between the memory controller asserting a row address, and then asserting a column address during the subsequent read or write command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1246. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 24(g) of the '938 Patent**.

1247. "**Claim 24(h) of the '938 Patent**" recites "CLmin is the minimum number of clock cycles of the clock signal required from the application of a column access command to the output of the data of the selected memory cell."

1248. In the **'938 Accused Products**, CLmin is the minimum number of clock cycles of the clock signal required from the application of a column access command to the output of the data of the selected memory cell.

1249. For example, in DDR3 SDRAM, CLmin is equivalent to CL or tAA. *See, e.g., JEDEC Standard No. 79-3F* at 157.

1250. For example, in DDR2, RCL is the difference between RLmin and CLmin, which is the number of clock cycles delay required between an active command row address strobe and a CAS. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1251. This is known as tRCD, which are the cycles required between the memory controller asserting a row address, and then asserting a column address during the subsequent read or write command. *See, e.g., JEDEC Standard No. 79-2F* at 33.

1252. As described in the preceding paragraphs, the **'938 Accused Products** meet the limitations of **Claim 24(h) of the '938 Patent**.

1253. NVIDIA has directly infringed and continue to directly infringe the **'938 Patent** by making, using, offering to sell, selling, and or importing the **'938 Accused Products**.

1254. NVIDIA uses the **'938 Accused Products**.

1255. For example, NVIDIA uses the **'938 Accused Products** for demonstration, testing, development, and configuration purposes.

1256. NVIDIA has used and continues to use its own products, including the **'938 Accused Products**, such as for demonstration, testing, development, and configuration purposes.

1257. For example and without limitation, NVIDIA has used and continues to use the GeForce G100 GPU.

1258. NVIDIA also has used and continues to use third-party products, including the **'938 Accused Products**, such as for demonstration, testing, and development purposes.

1259. When NVIDIA uses the **'938 Accused Products**, such as for testing or development purposes, such use directly infringes the **'938 Patent**.

1260. NVIDIA has in the past and continues to import, offer to sell, and sell its own products, including the **'938 Accused Products**.

1261. For example, NVIDIA has in the past and continues to import, offer to sell, and sell the NVIDIA Shield Tablet and NVIDIA Shield Portable.

1262. NVIDIA sells the **'938 Accused Products** to end users and customers.

1263. For example, NVIDIA sells the NVIDIA Shield Tablet and NVIDIA Shield Portable via its online store at <http://store.nvidia.com>.

1264. NVIDIA sells the **'938 Accused Products** to intermediate customers.

1265. For example, NVIDIA sells its products to distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers. **NVIDIA 2014 Form 10-K** at 8.

1266. The products NVIDIA sells to intermediate customers, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, include one or more of the **'938 Accused Products**.

1267. For example, NVIDIA sells NVIDIA Shield Tablet to third parties, such as Amazon, which advertises it on its website.

1268. When NVIDIA imports into the United States or offers for sale or sells in the United States the **'938 Accused Products**, such activity directly infringes the **'938 Patent**.

1269. Velocity has directly infringed and continue to directly infringe the **'938 Patent** by making, using, offering to sell, selling, and/or importing the **'938 Accused Products**.

1270. Velocity sells computers that incorporate the **'938 Accused GPUs**.

1271. For example, Velocity sells the Vector Z35, which includes one of the **'938 Accused GPUs**.

1272. The Vector Z35 can incorporate the NVIDIA GeForce 9500 GT, which can be configured to operate with GDDR2 memory.

1273. Velocity has used and continues to use its own products, including the **'938 Accused Products**, such as for demonstration, testing, development, and configuration purposes.

1274. For example and without limitation, Velocity has used and continues to use the Vector Z35.

1275. When Velocity uses the **'938 Accused Products**, such as for testing or development purposes, such use directly infringes the **'938 Patent**.

1276. Velocity has in the past and continues to import, offer for sale, and sell its own products, including the **'938 Accused Products**.

1277. When Velocity imports, offers for sale, or sells the **'938 Accused Products**, such activity directly infringes the **'938 Patent**.

1278. In 2013, Samsung and NVIDIA were engaged in negotiations regarding rights to various patents.

1279. On or around August 7, 2013, Samsung sent infringement claim charts and other documents to NVIDIA that specifically highlighted certain exemplary patents, claims, and products.

1280. The claim charts Samsung sent to NVIDIA on or about August 7, 2013 included claim charts relating to the **'938 Patent**.

1281. NVIDIA has had actual knowledge of the **'938 Patent** since at least as early as August 7, 2013.

1282. Velocity has had actual knowledge of the **'938 Patent** at least as of November 11, 2014, the date the original complaint in this action was served.

1283. NVIDIA and Velocity indirectly infringe the **'938 Patent** by inducing infringement by others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers,

customers, and end users, in accordance with 35 U.S.C. § 271(b), in this District and elsewhere in the United States.

1284. Direct infringement is the result of activities performed by others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, by, for example, making, using, offering to sell, selling, and or importing the **'938 Accused Products**.

1285. NVIDIA and Velocity have induced and continue to induce infringement of the **'938 Patent** by intending that others infringe the **'938 Patent** by making, using, offering to sell, selling, and/or importing the **'938 Accused Products**.

1286. NVIDIA and Velocity designed the **'938 Accused Products** such that they would each infringe one or more claims of the **'938 Patent**.

1287. NVIDIA and Velocity provide the **'938 Accused Products** to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users.

1288. By providing **'938 Accused Products** to others, NVIDIA and Velocity intend for **'938 Accused Products** to be made, used, offered for sale, sold, and/or imported in the United States.

1289. NVIDIA and Velocity also provide others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, instructions, user guides, and technical specifications.

1290. When others follow such instructions, user guides, and/or other design documentation, they directly infringe one or more claims of the **'938 Patent**.

1291. NVIDIA and Velocity know that by providing such instructions, user guides, and/or other design documentation, others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, follow those instructions, user guides, and other design documentation, and directly infringe one or more claims of the **'938 Patent**.

1292. NVIDIA and Velocity thus know that their actions actively induce infringement.

1293. NVIDIA sells the **'938 Accused Products** directly to customers and end users.

1294. For example, NVIDIA advertises the GeForce G100 GPU. *See, e.g., GeForce G100 GPU Page*.

1295. Graphics boards that incorporate the GeForce G100 GPU are advertised for sale on third party websites, such as **GeForce G100 Purchase Page**.

1296. Through such sales, NVIDIA specifically intends that others, such as customers and end users, will infringe one or more claims of the **'938 Patent**.

1297. NVIDIA provides the **'938 Accused Products**, and technical specifications for the **'938 Accused Products**, to others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users.

1298. Through such activity, NVIDIA specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, infringe the

'938 Patent by making, using, offering to sell, and selling in the United States, and importing into the United States the **'938 Accused Products**.

1299. Through its manufacture (either directly, or through third-parties) and/or sale of the infringing products, NVIDIA specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, will infringe one or more claims of the **'938 Patent**.

1300. NVIDIA specifically intends for others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'938 Patent** in the United States.

1301. For example, NVIDIA advertised its “NVIDIA Authorized Board Partner Program ensures an exceptional customer experience when purchasing graphics cards and motherboards manufactured by partners that make use of NVIDIA’s latest technologies.”

NVIDIA PartnerForce Info.

1302. NVIDIA claims that its “NVIDIA Authorized Board Partners offer the latest technologies from NVIDIA” and lists six Authorized Board Partners in the United States, six in Canada, and none in any other countries. *See NVIDIA PartnerForce Info.* NVIDIA also lists ten distributors in the United States (including “pre and post sales technical support”) and seven in Canada but none in any other country. *See NVIDIA PartnerForce Info.*

1303. As an additional example, as of October 31, 2014, NVIDIA also advertised the “NVIDIA PartnerForce Program” which is “a sales and marketing program for value-added

resellers, system builders, etailers and retailers who sell NVIDIA based components or systems.”

NVIDIA PartnerForce Program.

1304. NVIDIA advertises in the United States and in this judicial district that “[b]y joining the PartnerForce Program, you may leverage the world-class brands and technology platforms from NVIDIA, sales and technical support from our experienced team and hundreds of turnkey sales and marketing tools.” **NVIDIA PartnerForce Program.**

1305. NVIDIA also provides marketing materials and templates, including retail display items, web banners, copy blocks, logos, product shots, email and web page templates, to members of the “NVIDIA PartnerForce Program” which includes value-added resellers, system builders, and retailers inducing them to offer to sell and the **’938 Accused Products**. *See* **NVIDIA PartnerForce Program.**

1306. At the Ford Event and Conference Center, in Dearborn, Michigan, NVIDIA showcased its latest processor technologies, which power everything from the CAD software that designers use to style cars to the infotainment systems that drivers use to map their trips and listen to music, in an effort to encourage domestic car manufacturers to include infringing technology in vehicles manufactured and sold in the United States. *See* **Traveling the Road to Silicon Motown.**

1307. As described in the preceding paragraphs, NVIDIA specifically targets the United States market for **’938 Accused Products** actively induces others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **’938 Patent**.

1308. Velocity sells the **’938 Accused Products** directly to customers and end users.

1309. For example, Velocity advertised the Vector Z35, which incorporates the NVIDIA GeForce 9500 GT Graphics Card. *See, e.g., Configure Your Z35 Archive.*

1310. Through such sales, Velocity specifically intends that others, such as customers and end users, will infringe one or more claims of the **'938 Patent**.

1311. Velocity provides the **'938 Accused Products**, and technical specifications for the **'938 Accused Products**, to others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users.

1312. Through such activity, Velocity specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, infringe the **'938 Patent** by making, using, offering to sell, and selling in the United States, and importing into the United States the **'938 Accused Products**.

1313. Through its manufacture (either directly, or through third-parties) and/or sale of the infringing products, Velocity specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, will infringe one or more claims of the **'938 Patent**.

1314. NVIDIA and Velocity indirectly infringe the **'938 Patent** by contributing to infringement by others, such as distributors, wholesalers, retailers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(c), in this District and elsewhere in the United States.

1315. Direct infringement is the result of activities performed by the distributors, wholesalers, retailers, resellers, customers, and end users of the infringing products.

1316. For example, the **'938 Accused Products** allow for a method of controlling CAS latency of an SDRAM.

1317. When the infringing products are used as intended by distributors, wholesalers, retailers, resellers, customers, and end users, the **'938 Accused Products** necessarily control the CAS latency of an SDRAM in an infringing manner. The infringing products cannot operate in an acceptable manner without controlling the CAS latency as claimed in the **'938 Patent**.

1318. From the facts set forth above, it is evident that NVIDIA and Velocity knew that the ability to control the CAS latency of an SDRAM in the infringing products is especially made or especially adapted to operate in the products of NVIDIA's distributors, wholesalers, retailers, resellers, customers, and end users, and is not a staple article or commodity of commerce and that its infringing use is required for operation of the infringing products. Any other use would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental.

1319. NVIDIA's infringing products, with the ability to control the CAS latency of an SDRAM, are each a material part of the invention of the **'938 Patent** and are especially made or adapted to infringe one or more claims of the **'938 Patent**.

1320. Because the use of the **'938 Accused Products** necessarily infringes one or more claims of the **'938 Patent**, NVIDIA's sales of its infringing products have no substantial non-infringing uses.

1321. Accordingly, NVIDIA offers to sell, or sells a component, material, or apparatus for use in practicing one or more claims of the **'938 Patent** knowing the same to be especially made or especially adapted for use in an infringement of such patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

1322. NVIDIA's infringement of the **'938 Patent** is willful and deliberate, entitling Samsung to increased damages under 35 U.S.C. § 284 and to attorneys' fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

THIRD CLAIM FOR RELIEF
INFRINGEMENT OF U.S. PATENT NO. 6,287,902
(AGAINST NVIDIA AND VELOCITY)

1323. Each of the above listed paragraphs is incorporated herein by reference and adopted, as if fully set forth again.

1324. The **'902 Patent** was filed on May 25, 1999, issued on September 11, 2001, and is entitled "Methods Of Forming Etch Inhibiting Structures On Field Isolation Regions." The **'902 Patent** is generally directed to methods of forming microelectronic structures.

1325. The **'902 Patent** was assigned to SEC, and SEC continues to hold all rights, title, and interest in the **'902 Patent**. A true and correct copy of the **'902 Patent** is attached hereto as Exhibit C.

1326. For the **'902 Accused Products**, all **Accused 28 nm GPUs** and **Accused 28 nm SOCs** infringe the **'902 Patent** in substantially the same way.

1327. The **'902 Accused Products** are manufactured using one of TSMC's 28 nanometer high-k, metal gate processes ("**HKMG Processes**").

1328. The TSMC 28 nanometer **HKMG Processes** used to manufacture NVIDIA's 28 nanometer products are:

- 28HP (high performance with HKMG),
- 28HPL (low power with HKMG),
- 28HPM (high performance for mobile computing), or
- 28HPC (high performance compact).

See **TSMC 28HPC Process.**

1329. TSMC states that “[t]he 28nm high performance (HP) process is the first option to use high-k metal gate process technology.” **TSMC 28nm Technology.**

1330. TSMC uses the 28HP (high performance with HKMG) process in fabricating certain **’902 Accused Products.**

1331. TSMC states that the 28HPL process “adopts the same gate stack as HP technology.” **TSMC 28nm Technology.**

1332. TSMC uses the 28HPL (low power with HKMG) process in fabricating certain **’902 Accused Products.**

1333. TSMC states that the 28HPC (high performance compact) process is a “compact version of 28HPM.” **TSMC 28HPC Process.**

1334. TSMC uses the 28HPM (high performance for mobile computing) process in fabricating certain **’902 Accused Products.**

1335. TSMC uses the 28HPC process in fabricating certain **’902 Accused Products.**

1336. TSMC states that its **HKMG Processes** use a “replacement metal gate” technology. *See* **History on Side of Gate-Last High-k Approach.**

1337. The **Accused 28 nm GPUs** and **Accused 28 nm SOCs** include transistors formed by depositing one or more polysilicon gates that are replaced by metal layers.

1338. The **Accused 28 nm GPUs** and **Accused 28 nm SOCs** include spacers surrounding the metal gates. *See, e.g.,* **GK107 Report** at 36 and 47, Figures 2.3.1.4 and 2.3.1.13.

1339. The **Accused 28 nm GPUs** and **Accused 28 nm SOCs** include shallow trench isolation (“STI”) regions. *See, e.g.,* **GK107 Report** at 30 and 34, Figures 2.2.3 and 2.3.1.2.

1340. The **Accused 28 nm GPUs** and **Accused 28 nm SOCs** include transistors formed on the STI regions. *See, e.g., GK107 Report* at 30, Figure 2.2.3.

1341. The transistors on the STI regions in the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** are electrically isolated from transistors on PMOS and NMOS regions on the substrate. *See, e.g., GK107 Report* at 30, 33, 34, and 40, Figures 2.2.3, 2.3.1.1, 2.3.1.2, and 2.3.1.7.

1342. The GK107 GPU is one of the **Accused 28 nm GPUs**.

1343. The GK107 GPU is made by TSMC’s 28HP fabrication process. *See GK107 Report* at ix; *see also NVIDIA GeForce GTX 660 Features* (listing the “Kepler GPU Architecture” as a feature).

1344. NVIDIA states that “Kepler is based on 28-nanometer (nm) process technology and succeeds the 40-nm NVIDIA Fermi architecture, which was first introduced into the market in March 2010.” **NVIDIA Kepler Press Release**.

1345. The **GK107 Report** accurately depicts the gate stack for TSMC’s 28 nanometer **HKMG Processes**:




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.




Figure 2.3.1.9: Logic PMOS transistor gate stack, TEM cross-section.

GK107 Report at 36 and 42, Figure 2.3.1.4 and 2.3.1.9.

1346. For the **'902 Accused Products**, all **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** infringe the **'902 Patent** in the same way.

1347. The **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** are manufactured by TSMC.

1348. The Tegra 250 SOC is one of the **Accused 40 nm and Other SOCs**.

1349. The NVIDIA Tegra 250 SOC is made by TSMC's 40 nm process. *See Tegra 250 Report* at xii; *see also Tegra 2 Press Release*.

1350. The **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** include polysilicon gates. *See Tegra 250 Report* at xii and 29, Figure 2.3.5; **90-nm CMOS Device Technology** at 1.

1351. The polysilicon gates of the **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** are deposited using a patterning process. *See Tegra 250 Report* at 29, Figure 2.3.5; **90-nm CMOS Device Technology** at 1.

1352. The **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** include spacers surrounding the polysilicon gates. *See Tegra 250 Report* at 39 and 49, Figure 2.3.1.3 and 2.4.3; **TSMC 65 nm Process**.

1353. The **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** include an insulating layer covering the polysilicon gates. *See Tegra 250 Report* at 39 and 49, Figures 2.3.1.3 and 2.4.3; **TSMC 65 nm Process**.

1354. The **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** include STI regions. *See Tegra 250 Report* at 27, Figure 2.3.2; **90-nm CMOS Device Technology** at 1.

1355. The **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** include transistors formed on the STI regions. See **Tegra 250 Report** at 27, Figure 2.3.2.

1356. The transistors on the STI regions in the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** are electrically isolated from transistors on PMOS and NMOS regions on the substrate. See, e.g., **Tegra 250 Report** at 27; **Process Images Report** at 8.

1357. The **Tegra 250 Report** depicts the gate stack for TSMC’s 40 nanometer process:



Figure 2.4.3: Lower portion of PMD and dense S/D contacts; SEM cross-section.


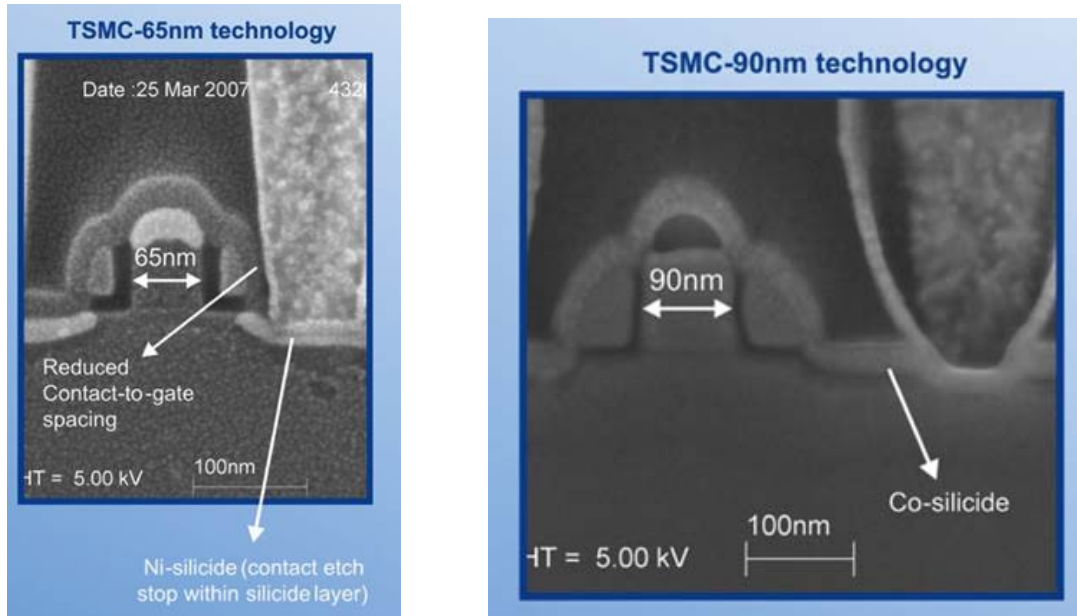


Figure 2.3.1.3: Location of a EELS line scan through the gate poly, offset spacer and sidewall spacer.

Tegra 250 Report at 49, 39, Figure 2.4.3, Figure 2.3.1.3.

1358. The **TSMC 65 nm Process** depicts the gate stack for TSMC’s 65 nanometer process and 90 nanometer process:



TSMC 65 nm Process.

1359. The **90-nm CMOS Device Technology** depicts the gate stack for TSMC’s 90 nanometer process:




Fig. 8. Representative TEM cross-section of a sub-nominal G-type transistor.

90-nm CMOS Device Technology at Figure 8.

1360. TSMC’s 65 nanometer and 90 technologies show a similar gate structure to TSMC’s 40 nanometer process.

1361. TSMC states that TSMC's "80-nm process is a lithographic shrink of TSMC's 90-nm process technology. As a consequence, this node supports most of the 90-nm libraries and intellectual property from TSMC and third-parties, requiring only simple re-characterization using 80-nm transistor models. Design rules are also a linear shrink from 90-nm." **Half-Node Process.**

1362. TSMC states that "TSMC's 55nm process is a 90% linear shrink process from the 65nm process." **TSMC 55nm Technology.**

1363. "**Claim 1(a) of the '902 Patent**" recites "[a] method for forming a contact hole for a microelectronic structure."

1364. The **'902 Accused Products** are made by a process that includes forming a contact hole for a microelectronic structure.

1365. For example, the NVIDIA GK107 GPU (a member of the **Accused 28 nm GPUs**) includes contact holes for a microelectronic structure on NMOS and PMOS regions of the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1366. The contact holes of the GK107 GPU are filled with one or more constituent materials.

1367. As shown below, the **GK107 Report** depicts contacts on NMOS and PMOS regions of the substrate.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 34, 40, Figures 2.3.1.2, 2.3.1.7. (red annotations added).

1368. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** are made by a process that includes forming a contact hole for a microelectronic structure.

1369. As another example, the Tegra 250 SOC (a member of the **Accused 40 nm and Other SOCs**) includes contact holes for a microelectronic structure on NMOS and PMOS regions of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1370. The contact holes are filled with one or more constituent materials. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1371. As shown below, the **Tegra 250 Report** depicts contacts on NMOS and PMOS regions of the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2. (red annotation added).

1372. “CA” in the **Tegra 250 Report** is a contact. **Tegra 250 Report** at 135.

1373. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** are made by a process that includes forming a contact hole for a microelectronic structure.

1374. As further described in the preceding paragraphs, the **'902 Accused Products** are made by a process that includes forming a contact hole for a microelectronic structure.

1375. “**Claim 1(b) of the '902 Patent**” recites “defining adjacent active and field regions on a substrate, and circuits thereon.”

1376. The **'902 Accused Products** are made by a process that includes defining adjacent active and field regions on a substrate.

1377. For example, the GK107 GPU is formed by a process that includes defining active NMOS and PMOS substrate regions adjacent to field regions containing STI regions. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1378. As shown below, the **GK107 Report** depicts NMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.

GK107 Report at 34, Figure 2.3.1.2. (red annotations added).

1379. As shown below, the **GK107 Report** depicts PMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 40, Figure 2.3.1.7. (red annotations added).

1380. Topographic images of the GK107 GPU also depict adjacent active and field regions on a substrate. See, e.g., **GK107 Report** at 33, Figure 2.3.1.1 depicting the field regions (in black) and the active PMOS and NMOS regions; see also **Process Images Report** at 21.

1381. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 1(b) of the '902 Patent**.

1382. As another example, the Tegra 250 SOC is formed by a process that includes defining active substrate regions adjacent to field regions containing STI regions.

1383. As shown below, the **Tegra 250 Report** depicts the adjacent active and field regions on the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2. (red annotations added).

1384. Topographical images of the Tegra 250 SOC also depict adjacent active and field regions on the substrate. *See, e.g., Tegra 250 Report* at 26, Figure 2.3.1 depicting the field regions (in black) and the active regions; *see also Process Images Report* at 8.

1385. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 1(b) of the '902 Patent**.

1386. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 1(b) of the '902 Patent**.

1387. “**Claim 1(c) of the '902 Patent**” recites “forming a field isolation layer on said field region.”

1388. The **'902 Accused Products** are made by a process that includes forming field isolation layers on the field region on the substrate.

1389. For example, the GK107 GPU includes STI regions filled with an oxide material. *See, e.g., GK107 Report* at x and 34, Figure 2.3.1.2.

1390. The STI regions are electrically insulated regions of the substrate.

1391. The **GK107 Report** depicts the field isolation layer on the field region on the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2, and 2.3.1.7.

1392. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 1(c) of the '902 Patent**.

1393. As another example, the Tegra 250 SOC includes STI regions filled with an oxide material. *See, e.g., Tegra 250 Report* at xiii and 27, Figure 2.3.2.

1394. The **Tegra 250 Report** depicts the field isolation layer on the field region on the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1395. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 1(c) of the '902 Patent**.

1396. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 1(c) of the '902 Patent**.

1397. "**Claim 1(d) of the '902 Patent**" recites "forming a first patterned conductive layer on said active region of said substrate spaced apart from said field region."

1398. The **'902 Accused Products** are made by a process that includes forming a first patterned conductive layer on the active region of the substrate spaced apart from the field region on the substrate.

1399. For example, the GK107 GPU includes NMOS and PMOS conductive transistors on active regions of the substrate. *See, e.g., GK107 Report* at 34, 35, and 40, Figures 2.3.1.2, 2.3.1.3, and 2.3.1.7.

1400. The NMOS and PMOS transistors on the active regions in the GK107 GPU are formed by a patterning process.

1401. As shown below, the **GK107 Report** depicts the NMOS and PMOS transistors on the active regions in the GK107 GPU.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.




Figure 2.3.1.3: Logic NMOS transistor, TEM cross-section.

GK107 Report at 34, 35, and 40, Figures 2.3.1.2, 2.3.1.3, and 2.3.1.7 (red annotations added).

1402. The **GK107 Report** describes the composition of the PMOS and NMOS transistors. *See GK107 Report* at x-xi.

1403. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 1(d) of the '902 Patent**.

1404. As another example, the Tegra 250 SOC includes transistors on active regions of the substrate. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1405. The transistors on the active regions in the Tegra 250 SOC include a patterned polysilicon gate. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1406. As shown below, the **Tegra 250 Report** depicts the transistors in the Tegra 250 SOC.




Figure 2.3.5: Logic transistor; TEM cross-section.

Tegra 250 Report at 29, Figure 2.3.5.

1407. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 1(d) of the '902 Patent**.

1408. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 1(d) of the '902 Patent**.

1409. “**Claim 1(e) of the '902 Patent**” recites “forming an etch inhibiting layer on said field isolation layer adjacent said active region of said substrate, the active region including the first patterned conductive layer.”

1410. The **'902 Accused Products** are made by a process that includes forming an etch inhibiting layer on the field isolation layer on the substrate. The etch inhibiting layer is adjacent to the active region of the substrate, which contains the first patterned conductive layer.

1411. For example, the GK107 GPU includes gate stacks on the STI region of the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1412. The gate stacks on the STI region of the GK107 GPU are adjacent to active NMOS and PMOS regions of the substrate containing NMOS and PMOS transistors. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1413. As shown below, the **GK107 Report** depicts gate stacks formed on the STI regions of the substrate adjacent to PMOS and NMOS gate stacks.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 34 and 40, Figures 2.3.1.2 and 2.3.1.7 (red annotations added).

1414. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 1(e) of the '902 Patent**.

1415. As another example, the Tegra 250 SOC includes gate stacks on the STI region of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1416. The gate stacks on the STI region of the Tegra 250 SOC are adjacent to transistors on the active regions of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1417. As shown below, the **Tegra 250 Report** depicts gate stacks on STI regions of the substrate adjacent to transistors on active regions of the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2 (red annotation added).

1418. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 1(e) of the '902 Patent**.

1419. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 1(e) of the '902 Patent**.

1420. “**Claim 1(f) of the '902 Patent**” recites “wherein said etch inhibiting layer comprises a second patterned conductive layer and an insulating spacer along a sidewall of the second patterned conductive layer.”

1421. The **'902 Accused Products** are made by a process wherein the etch inhibiting layer includes a second patterned conductive layer and an insulting spacer along the sidewalls of the second patterned conductive layer.

1422. For example, the gate stacks on the STI region of the GK107 GPU include a conductive gate surrounded by sidewall spacers. *See* **GK107 Report** at 30, Figure 2.2.3.

1423. The conductive gate of the gate stacks on the STI region of the GK107 GPU are formed by a patterning process.

1424. The sidewall spacers are on the left and right sides of the conductive gate of the gate stacks of the GK107 GPU. *See, e.g.,* **GK107 Report** at 30, Figure 2.2.3.

1425. The sidewall spacers of the GK107 GPU include silicon nitride.

1426. The sidewall spacers of the GK107 GPU are electrically insulating.

1427. As shown below, the **GK107 Report** depicts the conductive gate and insulating sidewall spacers formed on the STI regions of the substrate.




Figure 2.2.3: Shallow trench isolation, logic region, TEM cross-section.

GK107 Report at 30, Figure 2.2.3.

1428. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 1(f) of the '902 Patent**.

1429. As another example, the gate stacks on the STI region of the Tegra 250 SOC include a polysilicon gate and spacers. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

1430. The polysilicon gate of the Tegra 250 SOC is formed by a patterning process.

1431. The spacers are on the left and right sides of the polysilicon gate of the gate stacks of the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

1432. The **Tegra 250 Report** depicts the polysilicon gate and spacers formed on the STI regions of the substrate.




Figure 2.4.3: Lower portion of PMD and dense S/D contacts; SEM cross-section.

Tegra 250 Report at 49, Figure 2.4.3.

1433. The **Tegra 250 Report** describes the composition of the spacers. *See Tegra 250 Report* at xiii.

1434. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 1(f) of the '902 Patent**.

1435. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 1(f) of the '902 Patent**.

1436. “**Claim 1(g) of the '902 Patent**” recites “wherein the second patterned conductive layer does not extend over the active region of the substrate.”

1437. The **'902 Accused Products** are made by a process wherein the second patterned conductive layer on the field region on the substrate does not extend over the active region of the substrate.

1438. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active PMOS and NMOS regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

1439. The **GK107 Report** depicts that the conductive gates on the STI region do not extend over the active regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

1440. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 1(f) of the '902 Patent**.

1441. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

1442. The **Tegra 250 Report** depicts that the polysilicon gates on the STI region do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

1443. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 1(g) of the '902 Patent**.

1444. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 1(g) of the '902 Patent**.

1445. "**Claim 1(h) of the '902 Patent**" recites "wherein the second patterned conductive layer is a dummy pattern electrically isolated from the substrate and circuits thereon."

1446. The **'902 Accused Products** are made by a process wherein the second patterned conductive layer is electrically isolated from the substrate and circuits on the substrate and is a dummy pattern.

1447. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. *See* **GK107 Report** at 30, Figure 2.2.3.

1448. The conductive gates of the GK107 GPU on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

1449. The conductive gates of the GK107 GPU on the STI regions are dummy patterns. *See* **GK107 Report** at 33.

1450. The **GK107 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See* **GK107 Report** at 30 and 33, Figures 2.2.3 and 2.3.1; *see also* **Process Images Report** at 21.

1451. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 1(h) of the '902 Patent**.

1452. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. *See* **Tegra 250 Report** at 49, Figure 2.4.3.

1453. The polysilicon gates of the Tegra 250 SOC on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

1454. The polysilicon gates of the Tegra 250 SOC on the STI regions are dummy patterns. *See* **Tegra 250 Report** at 26.

1455. The **Tegra 250 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See* **Tegra 250 Report** at 26 and 27, Figures 2.3.1 and 2.3.2; **Process Images Report** at 8.

1456. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 1(h) of the '902 Patent**.

1457. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 1(h) of the '902 Patent**.

1458. “**Claim 1(i) of the '902 Patent**” recites “forming an insulating layer on said substrate, said field isolation layer, said first patterned conductive layer, and said etch inhibiting layer.”

1459. The **'902 Accused Products** are made by a process that includes forming an insulating layer on the etch inhibiting layer, the substrate, the field isolation layer, and the first patterned conductive layer.

1460. For example, the GK107 GPU is made by a process that includes depositing insulating pre-metal dielectric (PMD) oxide. *See* **GK107 Report** at 36, Figure 2.3.1.4.

1461. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 1(i) of the '902 Patent**.

1462. As another example, the Tegra 250 is made by a process that includes a liner. *See* **Tegra 250 Report** at 29, Figure 2.3.5.

1463. As shown below, the **Tegra 250 Report** depicts the liner.




Figure 2.3.5: Logic transistor; TEM cross-section.

Tegra 250 Report at 29, Figure 2.3.5.

1464. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 1(i) of the '902 Patent**.

1465. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 1(i) of the '902 Patent**.

1466. “**Claim 1(j) of the '902 Patent**” recites “forming a contact hole in said insulating layer exposing a portion of said active region between said etch inhibiting layer and said first patterned conductive layer.”

1467. The **'902 Accused Products** are made by a process that includes forming a contact hole in the insulating layer exposing a portion of the active region between the first patterned conductive layer and the etch inhibiting layer.

1468. For example, the GK107 GPU is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI

region and the PMOS and NMOS transistors on the active regions. *See, e.g., GK107 Report* at 62, Figure 2.4.5.

1469. As shown below, the **GK107 Report** depicts the contacts on the active region of the substrate between the dummy gate stacks on the STI regions and the NMOS and PMOS transistors on the active regions.




Figure 2.4.5: Contacts to NMOS S/D, TEM cross-section.

GK107 Report at 62, Figure 2.4.5. (red annotation added).

1470. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 1(j) of the '902 Patent**.

1471. As another example, the Tegra 250 SOC is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI region and the transistors on the active regions. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1472. As shown below, the **Tegra 250 Report** depicts the contacts on the active region of the substrate between the dummy gate stacks on the STI regions and the transistors on the active regions.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2. (red annotation added).

1473. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 1(j) of the '902 Patent**.

1474. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 1(j) of the '902 Patent**.

1475. “**Claim 2 of the '902 Patent**” recites “[a] method according to claim 1 wherein said insulating layer comprises nitride.”

1476. The **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** are made by a process wherein the insulating layer includes nitride.

1477. For example, the liner of the Tegra 250 SOC includes silicon nitride. *See Tegra 250 Report* at 40-41.

1478. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 2 of the '902 Patent**.

1479. “**Claim 3 of the '902 Patent**” recites “[a] method according to claim 1 wherein said first patterned conductive layer comprises a conductive portion and insulating spacers along sidewalls thereof.”

1480. The **'902 Accused Products** are made by a process wherein the first patterned conductive layer includes a conductive portion and insulating spacers along sidewalls.

1481. For example, the PMOS and NMOS transistors of the GK107 GPU on the active regions of the substrate include a conductive gate surrounded by sidewall spacers. *See* **GK107 Report** at 36, Figure 2.3.1.4

1482. The sidewall spacers are on the left and right sides of the conductive gate of the PMOS and NMOS transistors in the GK107 GPU. *See, e.g.,* **GK107 Report** at 36, Figure 2.3.1.4.

1483. As shown below, the **GK107 Report** depicts the conductive gate and insulating sidewall spacers for transistors in the GK107 GPU on the active region of the substrate.




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.

GK107 Report at 36, Figure 2.3.1.4.

1484. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 3 of the '902 Patent**.

1485. As another example, the transistors of the Tegra 250 SOC on the active regions of the substrate include a polysilicon gate and spacers. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1486. The spacers are on the left and right sides of the polysilicon gate of the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1487. **Tegra 250 Report** depicts the polysilicon gate and spacers formed on the active regions of the substrate.




Figure 2.3.5: Logic transistor; TEM cross-section.

Tegra 250 Report at 29, Figure 2.3.5.

1488. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 3 of the '902 Patent**.

1489. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 3 of the '902 Patent**.

1490. “**Claim 4 of the '902 Patent**” recites “[a] method according to claim 3 wherein said steps of forming said etch inhibiting layer and forming said patterned conductive layer are performed simultaneously.”

1491. The **'902 Accused Products** are made by a process that includes forming the etch inhibiting layer and the patterned conductive layers simultaneously.

1492. For example, in the GK107 GPU, the gate stacks on the STI regions of the substrate and transistors on the active regions of the substrate are formed simultaneously.

1493. As shown in the **GK107 Report**, SEM and TEM images depict nearly identical dummy gate stacks on the STI and transistors on active regions of the substrate. *See GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1494. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 4 of the '902 Patent**.

1495. As another example, in the Tegra 250 SOC, the gate stacks on the STI regions of the substrate and transistors on the active regions of the substrate are formed simultaneously.

1496. As shown in the **Tegra 250 Report**, SEM and TEM images depict nearly identical dummy gate stacks on the STI and transistors on active regions of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1497. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 4 of the '902 Patent**.

1498. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 4 of the '902 Patent**.

1499. "**Claim 5 of the '902 Patent**" recites "[a] method according to claim 1 wherein said field isolation layer comprises oxide."

1500. The **'902 Accused Products** are made by a process wherein the field isolation layer includes oxide.

1501. For example, the GK107 GPU is made by a process that includes a STI region filled with an oxide material. *See GK107 Report* at x.

1502. The **GK107 Report** depicts the STI layer filled with an oxide material. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1503. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 5 of the '902 Patent**.

1504. As another example, the Tegra 250 SOC is made by a process that includes a STI region filled with an oxide material. *See Tegra 250 Report* at xiii.

1505. The **Tegra 250 Report** depicts the STI layer filled with an oxide material. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1506. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 5 of the '902 Patent**.

1507. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 5 of the '902 Patent**.

1508. “**Claim 6 of the '902 Patent**” recites “[a] method according to claim 1 wherein said step of defining said field and active regions comprises the step of forming a trench in said substrate, and wherein said field isolation layer fills said trench.”

1509. The **'902 Accused Products** are made by a process that includes forming a trench in the substrate. The process includes filling the trench with the field isolation layer.

1510. For example, the GK107 GPU is made by a process that includes forming a shallow trench isolation region filled with an oxide material. *See GK107 Report* at x and 34, Figure 2.3.1.2.

1511. The **GK107 Report** depicts the STI layer filled with an oxide material. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1512. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 6 of the '902 Patent**.

1513. As another example, the Tegra 250 SOC is made by a process that includes a shallow trench isolation region filled with an oxide material. *See Tegra 250 Report* at xiii.

1514. The **Tegra 250 Report** depicts the STI layer filled with an oxide material. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1515. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 6 of the '902 Patent**.

1516. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 6 of the '902 Patent**.

1517. "**Claim 7 of the '902 Patent**" recites "[a] method according to claim 1 wherein the insulating spacer of the etch inhibiting layer extends to the active region of the substrate."

1518. The **'902 Accused Products** are made by a process wherein the insulating spacer of the etch inhibiting layer extends to the active regions of the substrate.

1519. For example, the GK107 GPU includes sidewall spacers on the STI region of the substrate that extend to the active regions of the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1520. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 7 of the '902 Patent**.

1521. As another example, the Tegra 250 includes spacers on the STI region of the substrate that extend to the active regions of the substrate. *See, e.g., Tegra 250 Report* at 51, Figure 2.4.6.

1522. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 7 of the '902 Patent**.

1523. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 7 of the '902 Patent**.

1524. "**Claim 8 of the '902 Patent**" recites "[a] method according to claim 1 wherein the insulating spacer and the insulating layer comprise different materials so that the insulating layer can be etched selectively with respect to the insulating spacer."

1525. The **'902 Accused Products** are made by a process that includes different materials for the insulating spacer and the insulating layer such that the insulating layer can be etched selectively with respect to the insulating spacer.

1526. For example, in the GK107 GPU, sidewall spacers and PMD oxide comprise different materials. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

1527. In the GK107 GPU, the sidewalls spacers include silicon nitride and PMD oxide includes oxide. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

1528. PMD oxide in the GK107 GPU can therefore be selectively etched with respect to the silicon nitride.

1529. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 8 of the '902 Patent**.

1530. As another example, in the Tegra 250 SOC, the spacers and the liner comprise different materials. *See, e.g., Tegra 250 Report* at 40-41.

1531. The liner in the Tegra 250 SOC can therefore be selectively etched with respect to the spacers. *See, e.g., Tegra 250 Report* at 55.

1532. The liner functions as an etch stop for the contact. *See, e.g., Tegra 250 Report* at 55.

1533. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 8 of the '902 Patent**.

1534. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 8 of the '902 Patent**.

1535. “**Claim 9 of the '902 Patent**” recites “[a] method according to claim 1 wherein contact hole is formed over the field isolation layer and exposes a portion of the etch inhibiting layer without exposing the field isolation layer to thereby increase the area over which the contact hole can be formed without damaging the field isolation layer.”

1536. The **'902 Accused Products** are made by a process that includes forming the contact hole over the field isolation layer, exposing a portion of the etch inhibiting layer without exposing the field isolation layer to thereby increase the area over which the contact hole can be formed without damaging the field isolation layer.

1537. For example, the GK107 GPU is made by a process includes forming a contact hole extending over the STI region of the substrate. *See* **GK107 Report** at 30 and 40, Figures 2.2.3 and 2.3.1.7.

1538. The **GK107 Report** depicts contact holes in the GK107 GPU exposing a portion of the dummy gate stacks without exposing the oxide of the STI region. *See* **GK107 Report** at 34, Figure 2.3.1.2.

1539. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 9 of the '902 Patent**.

1540. As another example, the Tegra 250 SOC is made by a process includes forming a contact hole extending over the STI region of the substrate. *See* **Tegra 250 Report** at 27, Figure 2.3.2.

1541. The **Tegra 250 Report** depicts contact holes in the Tegra 250 SOC exposing a portion of the dummy gate stacks without exposing the oxide of the STI region. *See Tegra 250 Report* at 27, Figure 2.3.2.

1542. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 9 of the '902 Patent**.

1543. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 9 of the '902 Patent**.

1544. “**Claim 10 of the '902 Patent**” recites “[a] method according to claim 1 wherein at least a portion of the contact hole is self-aligned with respect to the first patterned conductive layer and the etch inhibiting layer.”

1545. The **'902 Accused Products** are made by a process that includes forming at least portions of the contact holes as being self-aligned with respect to the first patterned conductive layer and the etch inhibiting layer.

1546. For example, the contact holes in the GK107 GPU are self-aligned features with respect to the PMOS and NMOS transistors on the active regions of the substrate. *See, e.g., GK107 Report* at 34, Figure 2.3.1.2.

1547. The contact holes in the GK107 GPU are self-aligned features with respect to the dummy gate stacks on the STI regions of the substrate. *See, e.g., GK107 Report* at 34, Figure 2.3.1.2.

1548. As shown below, the **GK107 Report** depicts self-aligned contacts with respect to transistors on the active regions of the substrate and dummy gate stacks on the STI regions of the substrate.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.

GK107 Report at 34, Figure 2.3.1.2.

1549. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 10 of the '902 Patent**.

1550. As another example, the contact holes in the Tegra 250 SOC are self-aligned features with respect to the transistors on the active regions of the substrate. *See Tegra 250 Report* at 27, Figure 2.3.2.

1551. The contact holes in the Tegra 250 SOC are self-aligned features with respect to the dummy gate stacks on the STI regions of the substrate. *See Tegra 250 Report* at 27, Figure 2.3.2.

1552. As shown below, the **Tegra 250 Report** depicts self-aligned contacts with respect to transistors on the active regions of the substrate and dummy gate stacks on the STI regions of the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2.

1553. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 10 of the '902 Patent**.

1554. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 10 of the '902 Patent**.

1555. “**Claim 11(a) of the '902 Patent**” recites “[a] method for forming a microelectronic structure.”

1556. The **'902 Accused Products** are made by a process that includes forming a microelectronic structure on a substrate.

1557. For example, the NVIDIA GK107 GPU includes NMOS and PMOS transistors and contacts on a substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1558. The **GK107 Report** depicts NMOS and PMOS transistors and contacts on a substrate. *See GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1559. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** are made by a process that includes forming a microelectronic structure.

1560. As another example, the Tegra 250 SOC includes transistors and contacts on a substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1561. The **Tegra 250 Report** depicts transistors and contacts on a substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1562. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** are made by a process that includes forming a microelectronic structure.

1563. As further described in the preceding paragraphs, the **'902 Accused Products** are made by a process that includes forming a microelectronic structure.

1564. “**Claim 11(b) of the '902 Patent**” recites “defining adjacent active and field regions on a substrate, and circuits thereon.”

1565. The **'902 Accused Products** are made by a process that includes defining adjacent active and field regions on a substrate.

1566. For example, the GK107 GPU is formed by a process that includes defining active NMOS and PMOS substrate regions adjacent to field regions containing STI regions. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1567. As shown below, the **GK107 Report** depicts NMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.

GK107 Report at 34, Figure 2.3.1.2. (red annotations added).

1568. As shown below, the GK107 Report depicts PMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 40, Figure 2.3.1.7. (red annotations added).

1569. Topographic images of the GK107 GPU also depict adjacent active and field regions on a substrate. *See, e.g.*, **GK107 Report** at 33, Figure 2.3.1.1 depicting the field regions (in black) and the active PMOS and NMOS regions; *see also* **Process Images Report** at 21.

1570. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(b) of the '902 Patent**.

1571. As another example, the Tegra 250 SOC is formed by a process that includes defining active substrate regions adjacent to field regions containing STI regions.

1572. As shown below, the **Tegra 250 Report** depicts the adjacent active and field regions on the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2. (red annotations added).

1573. Topographical images of the Tegra 250 SOC also depict adjacent active and field regions on the substrate. *See, e.g.*, **Tegra 250 Report** at 26, Figure 2.3.1 depicting the field regions (in black) and the active regions; *see also* **Process Images Report** at 8.

1574. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(b) of the '902 Patent**.

1575. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(b) of the '902 Patent**.

1576. "**Claim 11(c) of the '902 Patent**" recites "forming a field isolation layer which covers said field region."

1577. The **'902 Accused Products** are made by a process that includes forming a field isolation layer which covers the field region on the substrate.

1578. For example, the GK107 GPU includes STI regions filled with an oxide material. *See, e.g., GK107 Report* at x and 34, Figure 2.3.1.2.

1579. The STI regions are electrically insulated regions of the substrate.

1580. The **GK107 Report** depicts the field isolation layer which covers the field region on the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2, and 2.3.1.7.

1581. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(c) of the '902 Patent**.

1582. As another example, the Tegra 250 SOC includes STI regions filled with an oxide material. *See, e.g., Tegra 250 Report* at xiii and 27, Figure 2.3.2.

1583. The **Tegra 250 Report** depicts the field isolation layer which covers the field region on the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1584. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(c) of the '902 Patent**.

1585. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(c) of the '902 Patent**.

1586. “**Claim 11(d) of the ’902 Patent**” recites “forming a first patterned layer on said active region of said substrate spaced apart from said field region.”

1587. The **’902 Accused Products** are made by a process that includes forming a first patterned layer on the active region of the substrate spaced apart from the field region on the substrate.

1588. For example, the GK107 GPU includes PMOS conductive transistors on active regions of the substrate. *See, e.g., GK107 Report* at 40, Figure 2.3.1.7.

1589. The GK107 GPU includes NMOS conductive transistors on active regions of the substrate. *See, e.g., GK107 Report* at 34 and 35, Figures 2.3.1.2 and 2.3.1.3.

1590. The NMOS and PMOS transistors on the active regions in the GK107 GPU are formed by a patterning process.

1591. The **GK107 Report** depicts PMOS and NMOS transistors on the active regions in the GK107 GPU. *See, e.g., GK107 Report* at 34, 35, and 40, Figures 2.3.1.2, 2.3.1.3, and 2.3.1.7.

1592. The **GK107 Report** describes the composition of the PMOS and NMOS transistors. *See GK107 Report* at x-xi.

1593. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(d) of the ’902 Patent**.

1594. As another example, the Tegra 250 SOC includes transistors on active regions of the substrate. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1595. The transistors on the active regions in the Tegra 250 SOC include a patterned polysilicon gate. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1596. The **Tegra 250 Report** depicts the transistors in the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1597. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(d) of the '902 Patent**.

1598. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(d) of the '902 Patent**.

1599. "**Claim 11(e) of the '902 Patent**" recites "forming a second patterned layer on said field isolation layer adjacent said active region of said substrate, the active region including the first patterned layer."

1600. The **'902 Accused Products** are made by a process that includes forming a second patterned layer on the field isolation layer on the substrate. The second patterned layer is adjacent to the active region of the substrate, which contains the first patterned layer.

1601. For example, the GK107 GPU includes gate stacks on the STI region of the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1602. The gate stacks on the STI region of the GK107 GPU are adjacent to active NMOS and PMOS regions of the substrate containing NMOS and PMOS transistors. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1603. As shown below, the **GK107 Report** depicts gate stacks formed on the STI regions of the substrate adjacent to PMOS and NMOS gate stacks.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 34 and 40, Figures 2.3.1.2 and 2.3.1.7 (red annotations added).

1604. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(e) of the '902 Patent**.

1605. As another example, the Tegra 250 SOC includes gate stacks on the STI region of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1606. The gate stacks on the STI region of the Tegra 250 SOC are adjacent to transistors on the active regions of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1607. As shown below, the **Tegra 250 Report** depicts gate stacks on STI regions of the substrate adjacent to transistors on active regions of the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2 (red annotation added).

1608. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(e) of the '902 Patent**.

1609. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(e) of the '902 Patent**.

1610. “**Claim 11(f) of the '902 Patent**” recites “wherein said second patterned layer comprises a patterned conductive layer and an insulating spacer along a sidewall of the patterned conductive layer.”

1611. The '902 Accused Products are made by a process wherein the second patterned layer includes a patterned conductive layer and an insulating spacer along the sidewalls of the patterned conductive layer.

1612. For example, the gate stacks on the STI region of the GK107 GPU include a conductive gate surrounded by sidewall spacers. See **GK107 Report** at 30, Figure 2.2.3.

1613. The conductive gate of the gate stacks on the STI region of the GK107 GPU are formed by a patterning process.

1614. The sidewall spacers are on the left and right sides of the conductive gate of the gate stacks of the GK107 GPU. See, e.g., **GK107 Report** at 30, Figure 2.2.3.

1615. The sidewall spacers of the GK107 GPU include silicon nitride.

1616. The sidewall spacers of the GK107 GPU are electrically insulating.

1617. As shown below, the **GK107 Report** depicts the conductive gate and insulating sidewall spacers formed on the STI regions of the substrate.




Figure 2.2.3: Shallow trench isolation, logic region, TEM cross-section.

GK107 Report at 30, Figure 2.2.3.

1618. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(f) of the '902 Patent**.

1619. As another example, the gate stacks on the STI region of the Tegra 250 SOC include a polysilicon gate and spacers. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

1620. The polysilicon gate of the Tegra 250 SOC is formed by a patterning process.

1621. The spacers are on the left and right sides of the polysilicon gate of the gate stacks of the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

1622. **Tegra 250 Report** depicts the polysilicon gate and spacers formed on the STI regions of the substrate.




Figure 2.4.3: Lower portion of PMD and dense S/D contacts; SEM cross-section.

Tegra 250 Report at 49, Figure 2.4.3.

1623. The **Tegra 250 Report** describes the composition of the spacers. *See Tegra 250 Report* at xiii.

1624. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(f) of the '902 Patent**.

1625. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(f) of the '902 Patent**.

1626. "**Claim 11(g) of the '902 Patent**" recites "wherein the patterned conductive layer does not extend over the active region of the substrate."

1627. The **'902 Accused Products** are made by a process wherein the patterned conductive layer on the field region on the substrate does not extend over the active region of the substrate.

1628. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active PMOS and NMOS regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

1629. The **GK107 Report** depicts that the conductive gates on the STI region do not extend over the active regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

1630. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(g) of the '902 Patent**.

1631. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

1632. The **Tegra 250 Report** depicts that the polysilicon gates on the STI region do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

1633. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(g) of the '902 Patent**.

1634. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(g) of the '902 Patent**.

1635. “**Claim 11(h) of the '902 Patent**” recites “wherein the patterned conductive layer is a dummy pattern electrically isolated from the substrate and circuits thereon.”

1636. The **'902 Accused Products** are made by a process wherein the patterned conductive layer is electrically isolated from the substrate and circuits on the substrate and is a dummy pattern.

1637. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. *See* **GK107 Report** at 30, Figure 2.2.3.

1638. The conductive gates of the GK107 GPU on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

1639. The conductive gates of the GK107 GPU on the STI regions are dummy patterns. *See* **GK107 Report** at 33.

1640. The **GK107 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See* **GK107 Report** at 30 and 33, Figures 2.2.3 and 2.3.1; *see also* **Process Images Report** at 21.

1641. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(h) of the '902 Patent**.

1642. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. **Tegra 250 Report** at 49, Figure 2.4.3.

1643. The polysilicon gates of the Tegra 250 SOC on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

1644. The polysilicon gates of the Tegra 250 SOC on the STI regions are dummy patterns. *See Tegra 250 Report* at 26.

1645. The **Tegra 250 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See Tegra 250 Report* at 26 and 27, Figures 2.3.1 and 2.3.2; **Process Images Report** at 8.

1646. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(h) of the '902 Patent**.

1647. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(h) of the '902 Patent**.

1648. “**Claim 11(i) of the '902 Patent**” recites “forming an insulating layer covering said substrate, said field isolation layer, and said first and second patterned layers.”

1649. The **'902 Accused Products** are made by a process that includes forming an insulating layer covering the first patterned layer, the second patterned layer, the substrate, and the field isolation layer.

1650. For example, the GK107 GPU is made by a process that includes depositing insulating PMD oxide. *See GK107 Report* at 36, Figure 2.3.1.4.

1651. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(i) of the '902 Patent**.

1652. As another example, the Tegra 250 is made by a process that includes a liner. *See Tegra 250 Report* at 29, Figure 2.3.5.

1653. The **Tegra 250 Report** depicts the liner. *See Tegra 250 Report* at 29, Figure 2.3.5.

1654. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(i) of the '902 Patent**.

1655. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(i) of the '902 Patent**.

1656. "**Claim 11(j) of the '902 Patent**" recites "forming a contact hole in said insulating layer wherein said contact hole exposes a portion of said active region between said first and second patterned layers."

1657. The **'902 Accused Products** are made by a process that includes forming a contact hole in the insulating layer exposing a portion of the active region between the first patterned layer and the second patterned layer.

1658. For example, the GK107 GPU is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI region and the PMOS and NMOS transistors on the active regions. *See, e.g., GK107 Report* at 62, Figure 2.4.5.

1659. The **GK107 Report** depicts the contacts on the active region of the substrate between the dummy gate stacks on the STI regions and the NMOS and PMOS transistors on the active regions. *See GK107 Report* at 62, Figure 2.4.5.

1660. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(j) of the '902 Patent**.

1661. As another example, the Tegra 250 SOC is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI region and the transistors on the active regions. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1662. As shown below, the **Tegra 250 Report** depicts the contacts on the active region of the substrate between the dummy gate stacks on the STI regions and the transistors on the active regions. *See Tegra 250 Report* at 27, Figure 2.3.2.

1663. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(j) of the '902 Patent**.

1664. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(j) of the '902 Patent**.

1665. “**Claim 11(k) of the '902 Patent**” recites “wherein said steps of forming said first patterned layer and forming said second patterned layer are performed simultaneously.”

1666. The **'902 Accused Products** are made by a process that includes forming the first patterned layer and the second patterned layers simultaneously.

1667. For example, in the GK107 GPU, the gate stacks on the STI regions of the substrate and transistors on the active regions of the substrate are formed simultaneously.

1668. As shown in the **GK107 Report**, SEM and TEM images depict nearly identical dummy gate stacks on the STI and transistors on active regions of the substrate. *See GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1669. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 11(k) of the '902 Patent**.

1670. As another example, in the Tegra 250 SOC, the gate stacks on the STI regions of the substrate and transistors on the active regions of the substrate are formed simultaneously.

1671. As shown in the **Tegra 250 Report**, SEM and TEM images depict nearly identical dummy gate stacks on the STI and transistors on active regions of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1672. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 11(k) of the '902 Patent**.

1673. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 11(k) of the '902 Patent**.

1674. "**Claim 12(a) of the '902 Patent**" recites "[a] method for forming a microelectronic structure."

1675. The **'902 Accused Products** are made by a process that includes forming a microelectronic structure on a substrate.

1676. For example, the NVIDIA GK107 GPU includes NMOS and PMOS transistors and contacts on a substrate. *See, e.g., GK107 Report* at 34, 40, Figures 2.3.1.2, 2.3.1.7.

1677. The **GK107 Report** depicts NMOS and PMOS transistors and contacts on a substrate. *See GK107 Report* at 34, 40, Figures 2.3.1.2, 2.3.1.7.

1678. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** are made by a process that includes forming a microelectronic structure.

1679. As another example, the Tegra 250 SOC includes transistors and contacts on a substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1680. The **Tegra 250 Report** depicts transistors and contacts on a substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1681. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** are made by a process that includes forming a microelectronic structure.

1682. As described in the preceding paragraphs, the **'902 Accused Products** are made by a process that includes forming a microelectronic structure.

1683. “**Claim 12(b) of the '902 Patent**” recites “defining adjacent active and field regions on a substrate, and circuits thereon.”

1684. The **'902 Accused Products** are made by a process that includes defining adjacent active and field regions on a substrate.

1685. For example, the GK107 GPU is formed by a process that includes defining active NMOS and PMOS substrate regions adjacent to field regions containing STI regions. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1686. As shown below, the **GK107 Report** depicts NMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.

GK107 Report at 34, Figure 2.3.1.2. (red annotations added).

1687. As shown below, the **GK107 Report** depicts PMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 40, Figure 2.3.1.7. (red annotations added).

1688. Topographic images of the GK107 GPU also depict adjacent active and field regions on a substrate. *See, e.g., GK107 Report* at 33, Figure 2.3.1.1 depicting the field regions (in black) and the active PMOS and NMOS regions; *see also Process Images Report* at 21.

1689. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 12(b) of the '902 Patent**.

1690. As another example, the Tegra 250 SOC is formed by a process that includes defining active substrate regions adjacent to field regions containing STI regions.

1691. As shown below, the **Tegra 250 Report** depicts the adjacent active and field regions on the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2. (red annotations added).

1692. Topographical images of the Tegra 250 SOC also depict adjacent active and field regions on the substrate. *See, e.g., Tegra 250 Report* at 26, Figure 2.3.1 depicting the field regions (in black) and the active regions; *see also Process Images Report* at 8.

1693. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 12(b) of the '902 Patent**.

1694. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 12(b) of the '902 Patent**.

1695. “**Claim 12(c) of the '902 Patent**” recites “forming a field isolation layer which covers said field region.”

1696. The **'902 Accused Products** are made by a process that includes forming a field isolation layer which covers the field region on the substrate.

1697. For example, the GK107 GPU includes STI regions filled with an oxide material. *See, e.g., GK107 Report* at x and 34, Figure 2.3.1.2.

1698. The STI regions are electrically insulated regions of the substrate.

1699. The **GK107 Report** depicts the field isolation layer which covers the field region on the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2, and 2.3.1.7.

1700. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 12(c) of the '902 Patent**.

1701. As another example, the Tegra 250 SOC includes STI regions filled with an oxide material. *See, e.g., Tegra 250 Report* at xiii and 27, Figure 2.3.2.

1702. The **Tegra 250 Report** depicts the field isolation layer which covers the field region on the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1703. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 12(c) of the '902 Patent**.

1704. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 12(c) of the '902 Patent**.

1705. "**Claim 12(d) of the '902 Patent**" recites "forming a first patterned layer on said active region of said substrate spaced apart from said field region."

1706. The **'902 Accused Products** are made by a process that includes forming a first patterned layer on the active region of the substrate spaced apart from the field region on the substrate.

1707. For example, the GK107 GPU includes PMOS conductive transistors on active regions of the substrate. *See GK107 Report* at 40, Figure 2.3.1.7.

1708. The GK107 GPU includes NMOS conductive transistors on active regions of the substrate. *See GK107 Report* at 34 and 35, Figures 2.3.1.2 and 2.3.1.3.

1709. The NMOS and PMOS transistors on the active regions in the GK107 GPU are formed by a patterning process.

1710. The **GK107 Report** depicts PMOS and NMOS transistors on the active regions in the GK107 GPU. *See, e.g., GK107 Report* at 34, 35, and 40, Figures 2.3.1.2, 2.3.1.3, and 2.3.1.7

1711. The **GK107 Report** describes the composition of the PMOS and NMOS transistors. *See GK107 Report* at x-xi.

1712. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 12(d) of the '902 Patent**.

1713. As another example, the Tegra 250 SOC includes transistors on active regions of the substrate. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1714. The transistors on the active regions in the Tegra 250 SOC include a patterned polysilicon gate. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1715. The **Tegra 250 Report** depicts the transistors in the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1716. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 12(d) of the '902 Patent**.

1717. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 12(d) of the '902 Patent**.

1718. “**Claim 12(e) of the '902 Patent**” recites “forming a second patterned layer on said field isolation layer adjacent said active region of said substrate, the active region including the first patterned layer.”

1719. The **'902 Accused Products** are made by a process that includes forming a second patterned layer on the field isolation layer on the substrate. The second patterned layer is adjacent to the active region of the substrate, which contains the first patterned layer.

1720. For example, the GK107 GPU includes gate stacks on the STI region of the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1721. The gate stacks on the STI region of the GK107 GPU are adjacent to active NMOS and PMOS regions of the substrate containing NMOS and PMOS transistors. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1722. The **GK107 Report** depicts gate stacks formed on the STI regions of the substrate adjacent to PMOS and NMOS gate stacks. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1723. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 12(e) of the '902 Patent**.

1724. As another example, the Tegra 250 SOC includes gate stacks on the STI region of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1725. The gate stacks on the STI region of the Tegra 250 SOC are adjacent to transistors on the active regions of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1726. The **Tegra 250 Report** depicts gate stacks on STI regions of the substrate adjacent to transistors on active regions of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1727. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 12(e) of the '902 Patent**.

1728. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 12(e) of the '902 Patent**.

1729. "**Claim 12(f) of the '902 Patent**" recites "wherein said second patterned layer comprises a patterned conductive layer and an insulating spacer along a sidewall of the patterned conductive layer."

1730. The **'902 Accused Products** are made by a process wherein the second patterned layer includes a patterned conductive layer and an insulating spacer along the sidewalls of the patterned conductive layer.

1731. For example, the gate stacks on the STI region of the GK107 GPU include a conductive gate surrounded by sidewall spacers. *See* **GK107 Report** at 30, Figure 2.2.3.

1732. The conductive gate of the gate stacks on the STI region of the GK107 GPU are formed by a patterning process.

1733. The sidewall spacers are on the left and right sides of the conductive gate of the gate stacks of the GK107 GPU. *See, e.g.,* **GK107 Report** at 30, Figure 2.2.3.

1734. The sidewall spacers of the GK107 GPU include silicon nitride.

1735. The sidewall spacers of the GK107 GPU are electrically insulating.

1736. As shown below, the **GK107 Report** depicts the conductive gate and insulating sidewall spacers formed on the STI regions of the substrate.




Figure 2.2.3: Shallow trench isolation, logic region, TEM cross-section.

GK107 Report at 30, Figure 2.2.3.

1737. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 12(f) of the '902 Patent**.

1738. As another example, the gate stacks on the STI region of the Tegra 250 SOC include a polysilicon gate and spacers. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

1739. The polysilicon gate of the Tegra 250 SOC is formed by a patterning process.

1740. The spacers are on the left and right sides of the polysilicon gate of the gate stacks of the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

1741. **Tegra 250 Report** depicts the polysilicon gate and spacers formed on the STI regions of the substrate.




Figure 2.4.3: Lower portion of PMD and dense S/D contacts; SEM cross-section.

Tegra 250 Report at 49, Figure 2.4.3.

1742. The **Tegra 250 Report** describes the composition of the spacers. *See Tegra 250 Report* at xiii.

1743. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 12(f) of the '902 Patent**.

1744. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 12(f) of the '902 Patent**.

1745. “**Claim 12(g) of the '902 Patent**” recites “wherein the patterned conductive layer does not extend over the active region of the substrate.”

1746. The **'902 Accused Products** are made by a process wherein the patterned conductive layer on the field region on the substrate does not extend over the active region of the substrate.

1747. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active PMOS and NMOS regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

1748. The **GK107 Report** depicts that the conductive gates on the STI region do not extend over the active regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

1749. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 12(g) of the '902 Patent**.

1750. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

1751. The **Tegra 250 Report** depicts that the polysilicon gates on the STI region do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

1752. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 12(g) of the '902 Patent**.

1753. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 12(g) of the '902 Patent**.

1754. "**Claim 12(h) of the '902 Patent**" recites "wherein the patterned conductive layer is a dummy pattern electrically isolated from the substrate and circuits thereon."

1755. The **'902 Accused Products** are made by a process wherein the patterned conductive layer is electrically isolated from the substrate and circuits on the substrate and is a dummy pattern.

1756. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. *See* **GK107 Report** at 30, Figure 2.2.3.

1757. The conductive gates of the GK107 GPU on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

1758. The conductive gates of the GK107 GPU on the STI regions are dummy patterns. *See* **GK107 Report** at 33.

1759. The **GK107 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See* **GK107 Report** at 30 and 33, Figures 2.2.3 and 2.3.1; *see also* **Process Images Report** at 21.

1760. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 12(h) of the '902 Patent**.

1761. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. **Tegra 250 Report** at 49, Figure 2.4.3.

1762. The polysilicon gates of the Tegra 250 SOC on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

1763. The polysilicon gates of the Tegra 250 SOC on the STI regions are dummy patterns. *See* **Tegra 250 Report** at 26.

1764. The **Tegra 250 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See* **Tegra 250 Report** at 26 and 27, Figures 2.3.1 and 2.3.2; **Process Images Report** at 8.

1765. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 12(h) of the '902 Patent**.

1766. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 12(h) of the '902 Patent**.

1767. "**Claim 12(i) of the '902 Patent**" recites "forming an insulating layer covering said substrate, said field isolation layer, and said first and second patterned layers."

1768. The **'902 Accused Products** are made by a process that includes forming an insulating layer covering the first patterned layer, the second patterned layer, the substrate, and the field isolation layer.

1769. For example, the GK107 GPU is made by a process that includes depositing insulating PMD oxide. *See* **GK107 Report** at 36, Figure 2.3.1.4.

1770. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 12(i) of the '902 Patent**.

1771. As another example, the Tegra 250 is made by a process that includes a liner. *See* **Tegra 250 Report** at 29, Figure 2.3.5.

1772. As shown below, the **Tegra 250 Report** depicts the liner.




Figure 2.3.5: Logic transistor; TEM cross-section.

Tegra 250 Report at 29, Figure 2.3.5.

1773. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 12(i) of the '902 Patent**.

1774. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 12(i) of the '902 Patent**.

1775. “**Claim 12(j) of the '902 Patent**” recites “forming a contact hole in said insulating layer wherein said contact hole exposes a portion of said active region between said first and second patterned layers.”

1776. The **'902 Accused Products** are made by a process that includes forming a contact hole in the insulating layer exposing a portion of the active region between the first patterned layer and the second patterned layer.

1777. For example, the GK107 GPU is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI

region and the PMOS and NMOS transistors on the active regions. *See, e.g., GK107 Report* at 62, Figure 2.4.5.

1778. As shown below, the **GK107 Report** depicts the contacts on the active region of the substrate between the dummy gate stacks on the STI regions and the NMOS and PMOS transistors on the active regions.




Figure 2.4.5: Contacts to NMOS S/D, TEM cross-section.

GK107 Report at 62, Figure 2.4.5. (red annotation added).

1779. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 12(j) of the '902 Patent**.

1780. As another example, the Tegra 250 SOC is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI region and the transistors on the active regions. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1786. For example, the GK107 GPU is made by a process that includes forming a shallow trench isolation region in the substrate filled with an oxide material. *See* **GK107 Report** at x and 34, Figure 2.3.1.2.

1787. The **GK107 Report** depicts the STI layer filled with an oxide material. *See, e.g.,* **GK107 Report** at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1788. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 13 of the '902 Patent**.

1789. As another example, the Tegra 250 SOC is made by a process that includes a shallow trench isolation region filled with an oxide material. *See* **Tegra 250 Report** at xiii.

1790. The **Tegra 250 Report** depicts the STI layer filled with an oxide material. *See, e.g.,* **Tegra 250 Report** at 27, Figure 2.3.2.

1791. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 13 of the '902 Patent**.

1792. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 13 of the '902 Patent**.

1793. “**Claim 14 of the '902 Patent**” recites “[a] method according to claim 12 wherein the insulating layer and the insulating spacer comprise different materials so that the insulating layer can be etched selectively with respect to the insulating spacer.”

1794. The **'902 Accused Products** are made by a process that includes different materials for the insulating spacer and the insulating layer such that the insulating layer can be etched selectively with respect to the insulating spacer.

1795. For example, in the GK107 GPU, sidewall spacers and PMD oxide comprise different materials. *See, e.g.,* **GK107 Report** at 36, Figure 2.3.1.4.

1796. In the GK107 GPU, the sidewalls spacers include silicon nitride and PMD oxide includes oxide. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

1797. PMD oxide in the GK107 GPU can therefore be selectively etched with respect to the silicon nitride.

1798. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 14 of the '902 Patent**.

1799. As another example, in the Tegra 250 SOC, the spacers and the liner comprise different materials. *See, e.g., Tegra 250 Report* at 40-41.

1800. The liner in the Tegra 250 SOC can therefore be selectively etched with respect to the spacers. *See, e.g., Tegra 250 Report* at 55.

1801. The liner functions as an etch stop for the contact. *See, e.g., Tegra 250 Report* at 55.

1802. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 14 of the '902 Patent**.

1803. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 14 of the '902 Patent**.

1804. “**Claim 15(a) of the '902 Patent**” recites “[a] method for forming a microelectronic structure.”

1805. The **'902 Accused Products** are made by a process that includes forming a microelectronic structure on a substrate.

1806. For example, the NVIDIA GK107 GPU includes NMOS and PMOS transistors and contacts on a substrate. *See, e.g., GK107 Report* at 34, 40, Figures 2.3.1.2, 2.3.1.7.

1807. The **GK107 Report** depicts NMOS and PMOS transistors and contacts on a substrate. *See GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1808. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** are made by a process that includes forming a microelectronic structure.

1809. As another example, the Tegra 250 SOC includes transistors and contacts on a substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1810. The **Tegra 250 Report** depicts transistors and contacts on a substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1811. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** are made by a process that includes forming a microelectronic structure.

1812. As further described in the preceding paragraphs, the **'902 Accused Products** are made by a process that includes forming a microelectronic structure.

1813. “**Claim 15(b) of the '902 Patent**” recites “defining adjacent active and field regions on a substrate, and circuits thereon.”

1814. The **'902 Accused Products** are made by a process that includes defining adjacent active and field regions on a substrate.

1815. For example, the GK107 GPU is formed by a process that includes defining active NMOS and PMOS substrate regions adjacent to field regions containing STI regions. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1816. As shown below, the **GK107 Report** depicts NMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.

GK107 Report at 34, Figure 2.3.1.2. (red annotations added).

1817. As shown below, the GK107 Report depicts PMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 40, Figure 2.3.1.7. (red annotations added).

1818. Topographic images of the GK107 GPU also depict adjacent active and field regions on a substrate. *See, e.g., GK107 Report* at 33, Figure 2.3.1.1 depicting the field regions (in black) and the active PMOS and NMOS regions; *see also Process Images Report* at 21.

1819. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(b) of the '902 Patent**.

1820. As another example, the Tegra 250 SOC is formed by a process that includes defining active substrate regions adjacent to field regions containing STI regions.

1821. As shown below, the **Tegra 250 Report** depicts the adjacent active and field regions on the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

See, e.g., Tegra 250 Report at 27, Figure 2.3.2. (red annotations added).

1822. Topographical images of the Tegra 250 SOC also depict adjacent active and field regions on the substrate. *See, e.g., Tegra 250 Report* at 26, Figure 2.3.1 depicting the field regions (in black) and the active regions; *see also Process Images Report* at 8.

1823. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(b) of the '902 Patent**.

1824. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(b) of the '902 Patent**.

1825. "**Claim 15(c) of the '902 Patent**" recites "forming a field isolation layer which covers said field region."

1826. The **'902 Accused Products** are made by a process that includes forming a field isolation layer which covers the field region on the substrate.

1827. For example, the GK107 GPU includes STI regions filled with an oxide material. *See, e.g., GK107 Report* at x and 34, Figure 2.3.1.2.

1828. The STI regions are electrically insulated regions of the substrate.

1829. The **GK107 Report** depicts the field isolation layer which covers the field region on the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2, and 2.3.1.7.

1830. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(c) of the '902 Patent**.

1831. As another example, the Tegra 250 SOC includes STI regions filled with an oxide material. *See, e.g., Tegra 250 Report* at xiii and 27, Figure 2.3.2.

1832. The **Tegra 250 Report** depicts the field isolation layer which covers the field region on the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1833. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(c) of the '902 Patent**.

1834. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(c) of the '902 Patent**.

1835. “**Claim 15(d) of the ’902 Patent**” recites “forming a first patterned layer on said active region of said substrate spaced apart from said field region.”

1836. The **’902 Accused Products** are made by a process that includes forming a first patterned layer on the active region of the substrate spaced apart from the field region on the substrate.

1837. For example, the GK107 GPU includes PMOS conductive transistors on active regions of the substrate. *See* **GK107 Report** at 40, Figure 2.3.1.7.

1838. The GK107 GPU includes NMOS conductive transistors on active regions of the substrate. *See* **GK107 Report** at 34 and 35, Figures 2.3.1.2 and 2.3.1.3.

1839. The NMOS and PMOS transistors on the active regions in the GK107 GPU are formed by a patterning process.

1840. The **GK107 Report** depicts PMOS and NMOS transistors on the active regions in the GK107 GPU. *See, e.g.,* **GK107 Report** at 34, 35, and 40, Figures 2.3.1.2, 2.3.1.3, and 2.3.1.7.

1841. The **GK107 Report** describes the composition of the PMOS and NMOS transistors. *See* **GK107 Report** at x-xi.

1842. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(d) of the ’902 Patent**.

1843. As another example, the Tegra 250 SOC includes transistors on active regions of the substrate. *See, e.g.,* **Tegra 250 Report** at 29, Figure 2.3.5.

1844. The transistors on the active regions in the Tegra 250 SOC include a patterned polysilicon gate. *See, e.g.,* **Tegra 250 Report** at 29, Figure 2.3.5.

1845. The **Tegra 250 Report** depicts the transistors in the Tegra 250 SOC. *See Tegra 250 Report* at 29, Figure 2.3.5.

1846. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(d) of the '902 Patent**.

1847. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(d) of the '902 Patent**.

1848. “**Claim 15(e) of the '902 Patent**” recites “forming a second patterned layer on said field isolation layer adjacent said active region of said substrate, the active region including the first patterned layer.”

1849. The **'902 Accused Products** are made by a process that includes forming a second patterned layer on the field isolation layer on the substrate. The second patterned layer is adjacent to the active region of the substrate, which contains the first patterned layer.

1850. For example, the GK107 GPU includes gate stacks on the STI region of the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1851. The gate stacks on the STI region of the GK107 GPU are adjacent to active NMOS and PMOS regions of the substrate containing NMOS and PMOS transistors. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1852. As shown below, the **GK107 Report** depicts gate stacks formed on the STI regions of the substrate adjacent to PMOS and NMOS gate stacks.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 34 and 40, Figures 2.3.1.2 and 2.3.1.7 (red annotations added).

1853. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(e) of the '902 Patent**.

1854. As another example, the Tegra 250 SOC includes gate stacks on the STI region of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1855. The gate stacks on the STI region of the Tegra 250 SOC are adjacent to transistors on the active regions of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1856. As shown below, the **Tegra 250 Report** depicts gate stacks on STI regions of the substrate adjacent to transistors on active regions of the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

See, e.g., Tegra 250 Report at 27, Figure 2.3.2 (red annotation added).

1857. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(e) of the '902 Patent**.

1858. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(e) of the '902 Patent**.

1859. “**Claim 15(f) of the '902 Patent**” recites “wherein said second patterned layer comprises a patterned conductive layer and an insulating spacer along a sidewall of the patterned conductive layer.”

1860. The '902 Accused Products are made by a process wherein the second patterned layer includes a patterned conductive layer and an insulating spacer along the sidewalls of the patterned conductive layer.

1861. For example, the gate stacks on the STI region of the GK107 GPU include a conductive gate surrounded by sidewall spacers. See **GK107 Report** at 30, Figure 2.2.3.

1862. The conductive gate of the gate stacks on the STI region of the GK107 GPU are formed by a patterning process.

1863. The sidewall spacers are on the left and right sides of the conductive gate of the gate stacks of the GK107 GPU. See, e.g., **GK107 Report** at 30, Figure 2.2.3.

1864. The sidewall spacers of the GK107 GPU include silicon nitride.

1865. The sidewall spacers of the GK107 GPU are electrically insulating.

1866. As shown below, the **GK107 Report** depicts the conductive gate and insulating sidewall spacers formed on the STI regions of the substrate.




Figure 2.2.3: Shallow trench isolation, logic region, TEM cross-section.

GK107 Report at 30, Figure 2.2.3.

1867. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(f) of the '902 Patent**.

1868. As another example, the gate stacks on the STI region of the Tegra 250 SOC include a polysilicon gate and spacers. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

1869. The polysilicon gate of the Tegra 250 SOC is formed by a patterning process.

1870. The spacers are on the left and right sides of the polysilicon gate of the gate stacks of the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

1871. As shown below, the **Tegra 250 Report** depicts the polysilicon gate and spacers formed on the STI regions of the substrate.




Figure 2.4.3: Lower portion of PMD and dense S/D contacts; SEM cross-section.

Tegra 250 Report at 49, Figure 2.4.3.

1872. The **Tegra 250 Report** describes the composition of the spacers. *See Tegra 250 Report* at xiii.

1873. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(f) of the '902 Patent**.

1874. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(f) of the '902 Patent**.

1875. “**Claim 15(g) of the '902 Patent**” recites “wherein the patterned conductive layer does not extend over the active region of the substrate.”

1876. The **'902 Accused Products** are made by a process wherein the patterned conductive layer on the field region on the substrate does not extend over the active region of the substrate.

1877. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active PMOS and NMOS regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

1878. The **GK107 Report** depicts that the conductive gates on the STI region do not extend over the active regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

1879. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(g) of the '902 Patent**.

1880. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

1881. The **Tegra 250 Report** depicts that the polysilicon gates on the STI region do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

1882. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(g) of the '902 Patent**.

1883. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(g) of the '902 Patent**.

1884. “**Claim 15(h) of the ’902 Patent**” recites “wherein the patterned conductive layer is a dummy pattern electrically isolated from the substrate and circuits thereon.”

1885. The **’902 Accused Products** are made by a process wherein the patterned conductive layer is electrically isolated from the substrate and circuits on the substrate and is a dummy pattern.

1886. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. *See* **GK107 Report** at 30, Figure 2.2.3.

1887. The conductive gates of the GK107 GPU on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

1888. The conductive gates of the GK107 GPU on the STI regions are dummy patterns. *See* **GK107 Report** at 33.

1889. The **GK107 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See* **GK107 Report** at 30 and 33, Figures 2.2.3 and 2.3.1; *see also* **Process Images Report** at 21.

1890. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(h) of the ’902 Patent**.

1891. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. **Tegra 250 Report** at 49, Figure 2.4.3.

1892. The polysilicon gates of the Tegra 250 SOC on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

1893. The polysilicon gates of the Tegra 250 SOC on the STI regions are dummy patterns. *See Tegra 250 Report* at 26.

1894. The **Tegra 250 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See Tegra 250 Report* at 26 and 27, Figures 2.3.1 and 2.3.2; **Process Images Report** at 8.

1895. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(h) of the '902 Patent**.

1896. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(h) of the '902 Patent**.

1897. "**Claim 15(i) of the '902 Patent**" recites "forming an insulating layer covering said substrate, said field isolation layer, and said first and second patterned layers."

1898. The **'902 Accused Products** are made by a process that includes forming an insulating layer covering the first patterned layer, the second patterned layer, the substrate, and the field isolation layer.

1899. For example, the GK107 GPU is made by a process that includes depositing insulating PMD oxide. *See GK107 Report* at 36, Figure 2.3.1.4.

1900. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(i) of the '902 Patent**.

1901. As another example, the Tegra 250 is made by a process that includes a liner. *See Tegra 250 Report* at 29, Figure 2.3.5.

1902. The **Tegra 250 Report** depicts the liner. *See Tegra 250 Report* at 29, Figure 2.3.5.

1903. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(i) of the '902 Patent**.

1904. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(i) of the '902 Patent**.

1905. “**Claim 15(j) of the '902 Patent**” recites “forming a contact hole in said insulating layer wherein said contact hole exposes a portion of said active region between said first and second patterned layers.”

1906. The **'902 Accused Products** are made by a process that includes forming a contact hole in the insulating layer exposing a portion of the active region between the first patterned layer and the second patterned layer.

1907. For example, the GK107 GPU is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI region and the PMOS and NMOS transistors on the active regions. *See, e.g., GK107 Report* at 62, Figure 2.4.5.

1908. The **GK107 Report** depicts the contacts on the active region of the substrate between the dummy gate stacks on the STI regions and the NMOS and PMOS transistors on the active regions. *See, e.g., GK107 Report* at 62, Figure 2.4.5.

1909. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(j) of the '902 Patent**.

1910. As another example, the Tegra 250 SOC is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI region and the transistors on the active regions. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1911. The **Tegra 250 Report** depicts the contacts on the active region of the substrate between the dummy gate stacks on the STI regions and the transistors on the active regions. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1912. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(j) of the '902 Patent**.

1913. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(j) of the '902 Patent**.

1914. "**Claim 15(k) of the '902 Patent**" recites "wherein each of said first and second patterned layers comprises insulating spacers along sidewalls thereof."

1915. The **'902 Accused Products** are made by a process wherein each of the first and second patterned layers include spacers along sidewalls.

1916. For example, the GK107 GPU includes sidewall spacers on the left and right sides of the conductive gates on the STI region. *See, e.g., GK107 Report* at 30, Figure 2.2.3.

1917. The GK107 GPU additionally includes sidewall spacers on the left and right sides of the conductive gates of the transistors on the active region of the substrate. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

1918. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 15(k) of the '902 Patent**.

1919. As another example, the transistors of the Tegra 250 SOC on the active regions of the substrate include spacers on the left and right sides of the polysilicon gates. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1920. The Tegra 250 SOC additionally includes spacers on the left and right sides of the polysilicon gates on the STI region. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

1921. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 15(k) of the '902 Patent**.

1922. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 15(k) of the '902 Patent**.

1923. "**Claim 16 of the '902 Patent**" recites "[a] method according to claim 15 wherein said insulating layer can be selectively etched with respect to said insulating spacers."

1924. The **'902 Accused Products** are made by a process wherein the insulating layer can be etched selectively with respect to the insulating spacer.

1925. For example, in the GK107 GPU, sidewall spacers and PMD oxide comprise different materials. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

1926. In the GK107 GPU, the sidewalls spacers include silicon nitride and PMD oxide includes oxide. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

1927. PMD oxide in the GK107 GPU can therefore be selectively etched with respect to the silicon nitride.

1928. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 16 of the '902 Patent**.

1929. As another example, in the Tegra 250 SOC, the spacers and the liner comprise different materials. *See, e.g., Tegra 250 Report* at 40-41.

1930. The liner in the Tegra 250 SOC can therefore be selectively etched with respect to the spacers. *See, e.g., Tegra 250 Report* at 55.

1931. The liner functions as an etch stop for the contact. *See, e.g., Tegra 250 Report* at 55.

1932. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 16 of the '902 Patent**.

1933. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 16 of the '902 Patent**.

1934. "**Claim 17 of the '902 Patent**" recites "[a] method according to claim 16 wherein said insulating layer comprises nitride."

1935. The **Accused 40 nm and Other GPUs** and **Accused 40 nm and Other SOCs** are made by a process wherein the insulating layer includes nitride.

1936. For example, the liner of the Tegra 250 SOC includes silicon nitride. *See Tegra 250 Report* at 40-41.

1937. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 17 of the '902 Patent**.

1938. "**Claim 18(a) of the '902 Patent**" recites "[a] method for forming a microelectronic structure."

1939. The **'902 Accused Products** are made by a process that includes forming a microelectronic structure on a substrate.

1940. For example, the NVIDIA GK107 GPU includes NMOS and PMOS transistors and contacts on a substrate. *See, e.g., GK107 Report* at 34, 40, Figures 2.3.1.2, 2.3.1.7.

1941. The **GK107 Report** depicts NMOS and PMOS transistors and contacts on a substrate. *See GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1942. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** are made by a process that includes forming a microelectronic structure.

1943. As another example, the Tegra 250 SOC includes transistors and contacts on a substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1944. The **Tegra 250 Report** depicts transistors and contacts on a substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1945. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** are made by a process that includes forming a microelectronic structure.

1946. As further described in the preceding paragraphs, the **'902 Accused Products** are made by a process that includes forming a microelectronic structure.

1947. “**Claim 18(b) of the '902 Patent**” recites “defining adjacent active and field regions on a substrate, and circuits thereon.”

1948. The **'902 Accused Products** are made by a process that includes defining adjacent active and field regions on a substrate.

1949. For example, the GK107 GPU is formed by a process that includes defining active NMOS and PMOS substrate regions adjacent to field regions containing STI regions. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1950. As shown below, the **GK107 Report** depicts NMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.

GK107 Report at 34, Figure 2.3.1.2. (red annotations added).

1951. As shown below, the GK107 Report depicts PMOS active regions adjacent to field regions on the substrate.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 40, Figure 2.3.1.7. (red annotations added).

1952. Topographic images of the GK107 GPU also depict adjacent active and field regions on a substrate. *See, e.g., GK107 Report* at 33, Figure 2.3.1.1 depicting the field regions (in black) and the active PMOS and NMOS regions; *see also Process Images Report* at 21.

1953. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 18(b) of the '902 Patent**.

1954. As another example, the Tegra 250 SOC is formed by a process that includes defining active substrate regions adjacent to field regions containing STI regions.

1955. As shown below, the **Tegra 250 Report** depicts the adjacent active and field regions on the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2. (red annotations added).

1956. Topographical images of the Tegra 250 SOC also depict adjacent active and field regions on the substrate. *See, e.g., Tegra 250 Report* at 26, Figure 2.3.1 depicting the field regions (in black) and the active regions; *see also Process Images Report* at 8.

1957. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(b) of the '902 Patent**.

1958. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(b) of the '902 Patent**.

1959. "**Claim 18(c) of the '902 Patent**" recites "forming a field isolation layer which covers said field region."

1960. The **'902 Accused Products** are made by a process that includes forming a field isolation layer which covers the field region on the substrate.

1961. For example, the GK107 GPU includes STI regions filled with an oxide material. *See, e.g., GK107 Report* at x and 34, Figure 2.3.1.2.

1962. The STI regions are electrically insulated regions of the substrate.

1963. The **GK107 Report** depicts the field isolation layer which covers the field region on the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2, and 2.3.1.7.

1964. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 18(c) of the '902 Patent**.

1965. As another example, the Tegra 250 SOC includes STI regions filled with an oxide material. *See, e.g., Tegra 250 Report* at xiii and 27, Figure 2.3.2.

1966. The **Tegra 250 Report** depicts the field isolation layer which covers the field region on the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1967. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(c) of the '902 Patent**.

1968. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(c) of the '902 Patent**.

1969. “**Claim 18(d) of the ’902 Patent**” recites “forming a first patterned layer on said active region of said substrate spaced apart from said field region.”

1970. The **’902 Accused Products** are made by a process that includes forming a first patterned layer on the active region of the substrate spaced apart from the field region on the substrate.

1971. For example, the GK107 GPU includes PMOS conductive transistors on active regions of the substrate. *See* **GK107 Report** at 40, Figure 2.3.1.7.

1972. The GK107 GPU includes NMOS conductive transistors on active regions of the substrate. *See* **GK107 Report** at 34 and 35, Figures 2.3.1.2 and 2.3.1.3.

1973. The NMOS and PMOS transistors on the active regions in the GK107 GPU are formed by a patterning process.

1974. The **GK107 Report** depicts PMOS and NMOS transistors on the active regions in the GK107 GPU. *See, e.g.,* **GK107 Report** at 34, 35, and 40, Figures 2.3.1.2, 2.3.1.3, and 2.3.1.7.

1975. The **GK107 Report** describes the composition of the PMOS and NMOS transistors. *See* **GK107 Report** at x-xi.

1976. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 18(d) of the ’902 Patent**.

1977. As another example, the Tegra 250 SOC includes transistors on active regions of the substrate. *See, e.g.,* **Tegra 250 Report** at 29, Figure 2.3.5.

1978. The transistors on the active regions in the Tegra 250 SOC include a patterned polysilicon gate. *See, e.g.,* **Tegra 250 Report** at 29, Figure 2.3.5.

1979. The **Tegra 250 Report** depicts the transistors in the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

1980. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(d) of the '902 Patent**.

1981. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(d) of the '902 Patent**.

1982. “**Claim 18(e) of the '902 Patent**” recites “forming a second patterned layer on said field isolation layer adjacent said active region of said substrate, the active region including the first patterned layer.”

1983. The **'902 Accused Products** are made by a process that includes forming a second patterned layer on the field isolation layer on the substrate. The second patterned layer is adjacent to the active region of the substrate, which contains the first patterned layer.

1984. For example, the GK107 GPU includes gate stacks on the STI region of the substrate. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1985. The gate stacks on the STI region of the GK107 GPU are adjacent to active NMOS and PMOS regions of the substrate containing NMOS and PMOS transistors. *See, e.g., GK107 Report* at 34 and 40, Figures 2.3.1.2 and 2.3.1.7.

1986. As shown below, the **GK107 Report** depicts gate stacks formed on the STI regions of the substrate adjacent to PMOS and NMOS gate stacks.




Figure 2.3.1.2: Logic NMOS transistors, SEM cross-section.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 34 and 40, Figures 2.3.1.2 and 2.3.1.7 (red annotations added).

1987. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 18(e) of the '902 Patent**.

1988. As another example, the Tegra 250 SOC includes gate stacks on the STI region of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1989. The gate stacks on the STI region of the Tegra 250 SOC are adjacent to transistors on the active regions of the substrate. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

1990. As shown below, the **Tegra 250 Report** depicts gate stacks on STI regions of the substrate adjacent to transistors on active regions of the substrate.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2 (red annotation added).

1991. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(e) of the '902 Patent**.

1992. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(e) of the '902 Patent**.

1993. “**Claim 18(f) of the '902 Patent**” recites “wherein said second patterned layer comprises a patterned conductive layer and an insulating spacer along a sidewall of the patterned conductive layer.”

1994. The '902 **Accused Products** are made by a process wherein the second patterned layer includes a patterned conductive layer and an insulating spacer along the sidewalls of the patterned conductive layer.

1995. For example, the gate stacks on the STI region of the GK107 GPU include a conductive gate surrounded by sidewall spacers. *See GK107 Report* at 30, Figure 2.2.3.

1996. The conductive gate of the gate stacks on the STI region of the GK107 GPU are formed by a patterning process.

1997. The sidewall spacers are on the left and right sides of the conductive gate of the gate stacks of the GK107 GPU. *See, e.g., GK107 Report* at 30, Figure 2.2.3.

1998. The sidewall spacers of the GK107 GPU include silicon nitride.

1999. The sidewall spacers of the GK107 GPU are electrically insulating.

2000. As shown below, the **GK107 Report** depicts the conductive gate and insulating sidewall spacers formed on the STI regions of the substrate.




Figure 2.2.3: Shallow trench isolation, logic region, TEM cross-section.

GK107 Report at 30, Figure 2.2.3.

2001. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 18(f) of the '902 Patent**.

2002. As another example, the gate stacks on the STI region of the Tegra 250 SOC include a polysilicon gate and spacers. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

2003. The polysilicon gate of the Tegra 250 SOC is formed by a patterning process.

2004. The spacers are on the left and right sides of the polysilicon gate of the gate stacks of the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 49, Figure 2.4.3.

2005. **Tegra 250 Report** depicts the polysilicon gate and spacers formed on the STI regions of the substrate.




Figure 2.4.3: Lower portion of PMD and dense S/D contacts; SEM cross-section.

Tegra 250 Report at 49, Figure 2.4.3.

2006. The **Tegra 250 Report** describes that the composition of the spacers. *See Tegra 250 Report* at xiii.

2007. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(f) of the '902 Patent**.

2008. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(f) of the '902 Patent**.

2009. "**Claim 18(g) of the '902 Patent**" recites "wherein the patterned conductive layer does not extend over the active region of the substrate."

2010. The **'902 Accused Products** are made by a process wherein the patterned conductive layer on the field region on the substrate does not extend over the active region of the substrate.

2011. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active PMOS and NMOS regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

2012. The **GK107 Report** depicts that the conductive gates on the STI region do not extend over the active regions of the substrate. *See GK107 Report* at 30, Figure 2.2.3.

2013. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 18(g) of the '902 Patent**.

2014. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the STI regions such that they do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

2015. The **Tegra 250 Report** depicts that the polysilicon gates on the STI region do not extend over the active regions of the substrate. *See Tegra 250 Report* at 49, Figure 2.4.3.

2016. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(g) of the '902 Patent**.

2017. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(g) of the '902 Patent**.

2018. “**Claim 18(h) of the '902 Patent**” recites “wherein the patterned conductive layer is a dummy pattern electrically isolated from the substrate and circuits thereon.”

2019. The **'902 Accused Products** are made by a process wherein the patterned conductive layer is electrically isolated from the substrate and circuits on the substrate and is a dummy pattern.

2020. For example, the conductive gates of the GK107 GPU on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. *See* **GK107 Report** at 30, Figure 2.2.3.

2021. The conductive gates of the GK107 GPU on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

2022. The conductive gates of the GK107 GPU on the STI regions are dummy patterns. *See* **GK107 Report** at 33.

2023. The **GK107 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See* **GK107 Report** at 30 and 33, Figures 2.2.3 and 2.3.1; *see also* **Process Images Report** at 21.

2024. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 18(h) of the '902 Patent**.

2025. As another example, the polysilicon gates of the Tegra 250 SOC on the STI regions of the substrate are wholly formed over the electrically insulated STI regions. **Tegra 250 Report** at 49, Figure 2.4.3.

2026. The polysilicon gates of the Tegra 250 SOC on the STI regions are therefore electrically isolated from the active regions of the substrate and circuits on the active regions.

2027. The polysilicon gates of the Tegra 250 SOC on the STI regions are dummy patterns. *See Tegra 250 Report* at 26.

2028. The **Tegra 250 Report** and **Process Images Report** depict the dummy patterns on the STI region, electrically isolated from the substrate and circuits on the substrate. *See Tegra 250 Report* at 26 and 27, Figures 2.3.1 and 2.3.2; **Process Images Report** at 8.

2029. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(h) of the '902 Patent**.

2030. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(h) of the '902 Patent**.

2031. "**Claim 18(i) of the '902 Patent**" recites "forming an insulating layer covering said substrate, said field isolation layer, and said first and second patterned layers."

2032. The **'902 Accused Products** are made by a process that includes forming an insulating layer covering the first patterned layer, the second patterned layer, the substrate, and the field isolation layer.

2033. For example, the GK107 GPU is made by a process that includes depositing insulating PMD oxide. *See GK107 Report* at 36, Figure 2.3.1.4.

2034. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 18(i) of the '902 Patent**.

2035. As another example, the Tegra 250 is made by a process that includes a liner. *See Tegra 250 Report* at 29, Figure 2.3.5.

2036. The **Tegra 250 Report** depicts the liner. *See Tegra 250 Report* at 29, Figure 2.3.5.

2037. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(i) of the '902 Patent**.

2038. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(i) of the '902 Patent**.

2039. "**Claim 18(j) of the '902 Patent**" recites "forming a contact hole in said insulating layer wherein said contact hole exposes a portion of said active region between said first and second patterned layers."

2040. The **'902 Accused Products** are made by a process that includes forming a contact hole in the insulating layer exposing a portion of the active region between the first patterned layer and the second patterned layer.

2041. For example, the GK107 GPU is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI region and the PMOS and NMOS transistors on the active regions. *See, e.g., GK107 Report* at 62, Figure 2.4.5.

2042. As shown below, the **GK107 Report** depicts the contacts on the active region of the substrate between the dummy gate stacks on the STI regions and the NMOS and PMOS transistors on the active regions.




Figure 2.4.5: Contacts to NMOS S/D, TEM cross-section.

GK107 Report at 62, Figure 2.4.5. (red annotation added).

2043. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOC**s meet the limitations of **Claim 18(j) of the '902 Patent**.

2044. As another example, the Tegra 250 SOC is made by a process that includes forming a contact hole on the active region of the substrate between the dummy gate stacks on the STI region and the transistors on the active regions. *See, e.g., Tegra 250 Report* at 27, Figure 2.3.2.

2045. As shown below, the **Tegra 250 Report** depicts the contacts on the active region of the substrate between the dummy gate stacks on the STI regions and the transistors on the active regions.




Figure 2.3.2: Overview image of logic transistors; SEM cross-section.

Tegra 250 Report at 27, Figure 2.3.2. (red annotation added).

2046. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(j) of the '902 Patent**.

2047. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(j) of the '902 Patent**.

2048. “**Claim 18(k) of the '902 Patent**” recites “wherein the first patterned layer comprises a patterned gate layer and an insulating gate spacer along a sidewall thereof.”

2049. The **'902 Accused Products** are made by a process wherein the first patterned layer includes a patterned gate layer and an insulating spacer along a sidewall.

2050. For example, the PMOS and NMOS transistors of the GK107 GPU on the active regions of the substrate include a conductive gate surrounded by sidewall spacers. *See GK107 Report* at 36, Figure 2.3.1.4.

2051. The conductive gate of the transistors of the GK107 GPU are formed by a patterning process.

2052. The sidewall spacers are on the left and right sides of the conductive gate of the PMOS and NMOS transistors in the GK107 GPU. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

2053. The sidewall spacers of the GK107 GPU include silicon nitride.

2054. The sidewall spacers of the GK107 GPU are electrically insulating.

2055. As shown below, the **GK107 Report** depicts the conductive gate and insulating sidewall spacers for transistors in the GK107 GPU on the active region of the substrate.




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.

GK107 Report at 36, Figure 2.3.1.4.

2056. As described in the preceding paragraphs, the **Accused 28 nm GPUs** and **Accused 28 nm SOCs** meet the limitations of **Claim 18(k) of the '902 Patent**.

2057. As another example, the transistors of the Tegra 250 SOC on the active regions of the substrate include a polysilicon gate and spacers. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

2058. The polysilicon gates of the Tegra 250 SOC are formed by a patterning process.

2059. The spacers are on the left and right sides of the polysilicon gate of the Tegra 250 SOC. *See, e.g., Tegra 250 Report* at 29, Figure 2.3.5.

2060. As shown below, the **Tegra 250 Report** depicts the polysilicon gate and spacers formed on the active regions of the substrate.




Figure 2.3.5: Logic transistor; TEM cross-section.

Tegra 250 Report at 29, Figure 2.3.5.

2061. As described in the preceding paragraphs, the **Accused 40 nm GPUs** and **Accused 40 nm SOCs** meet the limitations of **Claim 18(k) of the '902 Patent**.

2062. As further described in the preceding paragraphs, the **'902 Accused Products** meet the limitations of **Claim 18(k) of the '902 Patent**.

2063. NVIDIA and Velocity have directly infringed and continue to directly infringe the **'902 Patent** by importing, offering to sell, selling, or using the **'902 Accused Products**.

2064. As described above, the **'902 Accused Products** are made by a process that infringes the **'902 Patent**.

2065. NVIDIA uses the **'902 Accused Products**.

2066. For example, NVIDIA uses its own products for demonstration, testing, and development purposes.

2067. NVIDIA has used and continues to use its own products, including the **'902 Accused Products**, such as for demonstration, testing, and development purposes.

2068. For example and without limitation, NVIDIA has used and continues to use the GM107, Tesla K10, NVIDIA Shield Tablet, NVIDIA Shield Portable, GeForce GTX 760, and GeForce GTX 770.

2069. NVIDIA also has used and continues to use third-party products, including the **'902 Accused Products**, such as for demonstration, testing, and development purposes.

2070. When NVIDIA uses the **'902 Accused Products**, such as for demonstration, testing, or development purposes, such use directly infringes the **'902 Patent**.

2071. NVIDIA has in the past and continues to import, offer to sell, and sell its own products made by processes that infringe the **'902 Patent**, including the **'902 Accused Products**.

2072. For example, NVIDIA has in the past and continues to import, offer to sell, and sell the GM107, Tesla K10, NVIDIA Shield Tablet, NVIDIA Shield Portable, GeForce GTX 760, and GeForce GTX 770.

2073. NVIDIA sells the **'902 Accused Products** to end users and customers.

2074. For example, NVIDIA sells the NVIDIA Shield Tablet and NVIDIA Shield Portable via its online store at <http://store.nvidia.com>.

2075. NVIDIA sells the **'902 Accused Products** to intermediate customers.

2076. For example, NVIDIA sells its products to distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers. **NVIDIA 2014 Form 10-K** at 8.

2077. The products NVIDIA sells to intermediate customers, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, include one or more of the **'902 Accused Products**.

2078. For example, NVIDIA sells the GeForce GTX 760 and GeForce GTX 770 to Best Buy, which sells these cards to end users and customers.

2079. When NVIDIA imports into the United States or offers for sale or sells in the United States the **'902 Accused Products**, such activity directly infringes the **'902 Patent**.

2080. Velocity sells computers that incorporate an **Accused 28 nm GPU**.

2081. For example, Velocity sells a computer called the Raptor Signature Edition.

2082. The Raptor Signature Edition computer can be configured to include a variety of NVIDIA GPUs, including a variety of the **Accused 28 nm GPUs**.

2083. For example, the Raptor Signature Edition computer can be configured to include an NVIDIA GeForce GTX Titan Black.

2084. Velocity has used and continues to use its own products made by processes that infringe the **'902 Patent**, including the **'902 Accused Products**, such as for testing and development purposes.

2085. For example and without limitation, Velocity has used and continues to use the Cruz Tablet L510 incorporating the NVIDIA Tegra 250 SOC.

2086. For example and without limitation, Velocity has used and continues to use computers that incorporate an **Accused 28 nm GPU**.

2087. When Velocity uses the **'902 Accused Products**, such as for demonstration, testing, or development purposes, such use directly infringes the **'902 Patent**.

2088. Velocity has in the past and continues to import, offer for sale, and sell its own products made by processes that infringe the **'902 Patent**, including the **'902 Accused Products**.

2089. For example and without limitation, Velocity has in the past and continues to import, offer for sale, and sell the Cruz Tablet L510 incorporating the NVIDIA Tegra 250 SOC

2090. For example and without limitation, Velocity has in the past and continues to import, offer for sale, and sell computers that incorporate an **Accused 28 nm GPU**.

2091. When Velocity imports, offers for sale, or sells the **'902 Accused Products**, such activity directly infringes the **'902 Patent**.

2092. In 2013, Samsung and NVIDIA were engaged in negotiations regarding rights to various patents.

2093. On or around August 7, 2013, Samsung sent infringement claim charts and other documents to NVIDIA that specifically highlighted certain exemplary patents, claims, and products.

2094. The claim charts Samsung sent to NVIDIA on or about August 7, 2013 included claim charts relating to the **'902 Patent**.

2095. NVIDIA has had actual knowledge of the **'902 Patent** since at least as early as August 7, 2013.

2096. Velocity has had actual knowledge of the **'902 Patent** at least as of November 11, 2014, the date the original complaint in this action was served.

2097. NVIDIA and Velocity indirectly infringe the **'902 Patent** by inducing infringement by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(b), in this District and elsewhere in the United States.

2098. Direct infringement is the result of activities performed by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, by, for example, importing, offering to sell, selling, or using the **'902 Accused Products**.

2099. NVIDIA and Velocity have induced and continue to induce infringement of the **'902 Patent** by intending that others infringe the **'902 Patent** by importing, offering to sell, selling, or using the **'902 Accused Products**. NVIDIA and Velocity designed the **'902 Accused Products** such that they would each infringe one or more claims of the **'902 Patent**.

2100. NVIDIA and Velocity provide the **'902 Accused Products** to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users.

2101. By providing **'902 Accused Products** to others, NVIDIA and Velocity intend for **'902 Accused Products** to be imported into the United States or offered for sale, sold, or used in the United States.

2102. NVIDIA and Velocity also provide others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers,

motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, instructions, user guides, and technical specifications.

2103. When others follow such instructions, user guides, and/or other design documentation, they directly infringe one or more claims of the **'902 Patent**.

2104. NVIDIA and Velocity know that by providing such instructions, user guides, and/or other design documentation, others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, follow those instructions, user guides, and other design documentation, and directly infringe one or more claims of the **'902 Patent**. NVIDIA and Velocity thus know that their actions actively induce infringement.

2105. NVIDIA sells the **'902 Accused Products** directly to customers and end users.

2106. For example, NVIDIA sells the NVIDIA Shield Tablet and NVIDIA Shield Portable via its online store at <http://store.nvidia.com>.

2107. Through such sales, NVIDIA specifically intends that others, such as customers and end users, will infringe one or more claims of the **'902 Patent**, for example by use of the **'902 Accused Products**.

2108. NVIDIA provides the **'902 Accused Products**, and reference designs for the **'902 Accused Products**, to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, and resellers.

2109. Through such activity, NVIDIA specifically intends that others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original

design manufacturers, motherboard manufacturers, add-in board manufacturers, and resellers, customers, and end users, infringe the **'902 Patent** by importing into the United States and offering to sell, selling, or using in the United States the **'902 Accused Products**.

2110. Through its manufacture (either directly, or through contract manufacturing facilities) and/or sale of the infringing products, NVIDIA specifically intends that others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, will infringe one or more claims of the **'902 Patent**.

2111. NVIDIA specifically intends for others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'902 Patent** in the United States.

2112. For example, NVIDIA advertised its “NVIDIA Authorized Board Partner Program ensures an exceptional customer experience when purchasing graphics cards and motherboards manufactured by partners that make use of NVIDIA’s latest technologies.”
NVIDIA PartnerForce Info.

2113. NVIDIA claims that its “NVIDIA Authorized Board Partners offer the latest technologies from NVIDIA” and lists six Authorized Board Partners in the United States, six in Canada, and none in any other countries. *See* **NVIDIA PartnerForce Info**. NVIDIA also lists ten distributors in the United States (including “pre and post sales technical support”) and seven in Canada but none in any other country. *See* **NVIDIA PartnerForce Info**.

2114. As an additional example, as of October 31, 2014, NVIDIA also advertised the “NVIDIA PartnerForce Program” which is “a sales and marketing program for value-added

resellers, system builders, etailers and retailers who sell NVIDIA based components or systems.”

NVIDIA PartnerForce Program.

2115. NVIDIA advertises in the United States and in this judicial district that “[b]y joining the PartnerForce Program, you may leverage the world-class brands and technology platforms from NVIDIA, sales and technical support from our experienced team and hundreds of turnkey sales and marketing tools.” **NVIDIA PartnerForce Program.**

2116. NVIDIA also provides marketing materials and templates, including retail display items, web banners, copy blocks, logos, product shots, email and web page templates, to members of the “NVIDIA PartnerForce Program” which includes value-added resellers, system builders, etailers and retailers inducing them to offer to sell and the **'902 Accused Products**. *See NVIDIA PartnerForce Program.*

2117. At the Ford Event and Conference Center, in Dearborn, Michigan, NVIDIA showcased its latest processor technologies, which power everything from the CAD software that designers use to style cars to the infotainment systems that drivers use to map their trips and listen to music, in an effort to encourage domestic car manufacturers to include infringing technology in vehicles manufactured and sold in the United States. *See Traveling The Road To Silicon Motown.*

2118. As described in the preceding paragraphs, NVIDIA specifically targets the United States market for **'902 Accused Products** actively induces others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'902 Patent**.

2119. Velocity sells computers that incorporate an **Accused 28 nm GPU**.

2120. For example, Velocity sells a computer called the Raptor Signature Edition.

2121. The Raptor Signature Edition computer can be configured to include a variety of NVIDIA GPUs, including a variety of the **Accused 28 nm GPUs**.

2122. For example, the Raptor Signature Edition computer can be configured to include an NVIDIA GeForce GTX Titan Black.

2123. Velocity provides the **'902 Accused Products** to others, such as resellers, customers, and end users.

2124. Through such activity, NVIDIA specifically intends that others, such as resellers, customers, and end users, infringe the **'902 Patent** by importing into the United States and offering to sell, selling, or using in the United States the **'902 Accused Products**.

2125. Velocity specifically targets the United States market for its products listed above and actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'902 Patent** in the United States.

2126. For example, Velocity sells products via its website, <http://www.velocitymicro.com>, to customers in the United States. Such products include the **'902 Accused Products**, such as, for example, computers that are configured to include an NVIDIA GeForce GTX Titan Black.

2127. The customers infringe the **'902 Patent** by using the **'902 Accused Products** purchased from Velocity.

2128. Velocity also sells computers to resellers. Such computers include the **'902 Accused Products**, such as, for example, computers that are configured to include an **Accused 28 nm GPU**.

2129. For example, Velocity products have been sold at nearly every major electronics retailer in North America including Best Buy, Circuit City, RadioShack, Costco, Sears, Target, and many others. *See Velocity Company History.*

2130. Velocity introduces products and services that infringe the Asserted Patents intending that they would be used in this Judicial District and elsewhere in the United States.

2131. The resellers infringe the **'902 Patent** by importing, offering selling, selling, or using the **'902 Accused Products**.

2132. The resellers also sell the **'902 Accused Products** to customers and end users, who infringe the **'902 Patent** by using those products.

2133. As described in the preceding paragraphs, Velocity specifically targets the United States market for **'902 Accused Products** actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'902 Patent**.

2134. NVIDIA's infringement of the **'902 Patent** is willful and deliberate, entitling Samsung to increased damages under 35 U.S.C. § 284 and to attorneys' fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

2135. Samsung has no adequate remedy at law for NVIDIA's and Velocity's infringement of the **'902 Patent** and is suffering irreparable harm, requiring permanent injunctive relief.

FOURTH CLAIM FOR RELIEF
INFRINGEMENT OF U.S. PATENT NO. 6,819,602
(AGAINST NVIDIA AND VELOCITY)

2136. Each of the above listed paragraphs is incorporated herein by reference and adopted, as if fully set forth again.

2137. The **'602 Patent** was filed on October 23, 2002, issued on November 16, 2004, and is entitled "Multimode Data Buffer and Method for Controlling Propagation Delay Time." The **'602 Patent** is generally directed to a data strobe buffer circuit.

2138. The **'602 Patent** was assigned to SEC, and SEC continues to hold all rights, title, and interest in the **'602 Patent**. A true and correct copy of the **'602 Patent** is attached hereto as Exhibit D.

2139. For the **'602 Accused Products**, all **'602 Accused GPUs** and **Accused SOCs** infringe the **'602 Patent** in substantially the same way.

2140. For example, all NVIDIA SOCs and GPUs are designed to operate in conjunction with memory, such as DDR3 or GDDR2, that is manufactured according to the standards set by JEDEC.

2141. The **'602 Accused Products** include memory controllers.

2142. The **'602 Accused Products** include memory controllers that have data buffers.

2143. The **Accused SOCs** and **'602 Accused GPUs** are designed to support memory that includes an input data strobe feature supporting single-ended and differential signaling, meaning that all the **'602 Accused Products** are designed to operate in substantially the same way.

2144. For example, the Tegra 2 includes a memory controller and an external memory controller ("EMC"). *See, e.g., Tegra 2 TRM* at 8.

2145. For example, the Tegra 3 includes a memory controller and an EMC. *See, e.g., Tegra 3 TRM* at 10.

2146. For example, the Tegra 4 includes a memory controller and two EMCs. *See, e.g., Tegra 4 TRM* at 12.

2147. For example, the Tegra K1 includes a memory controller. *See, e.g., Tegra K1 TRM* at 12.

2148. “**Claim 1(a) of the ’602 Patent**” recites “[a] data strobe input buffer.”

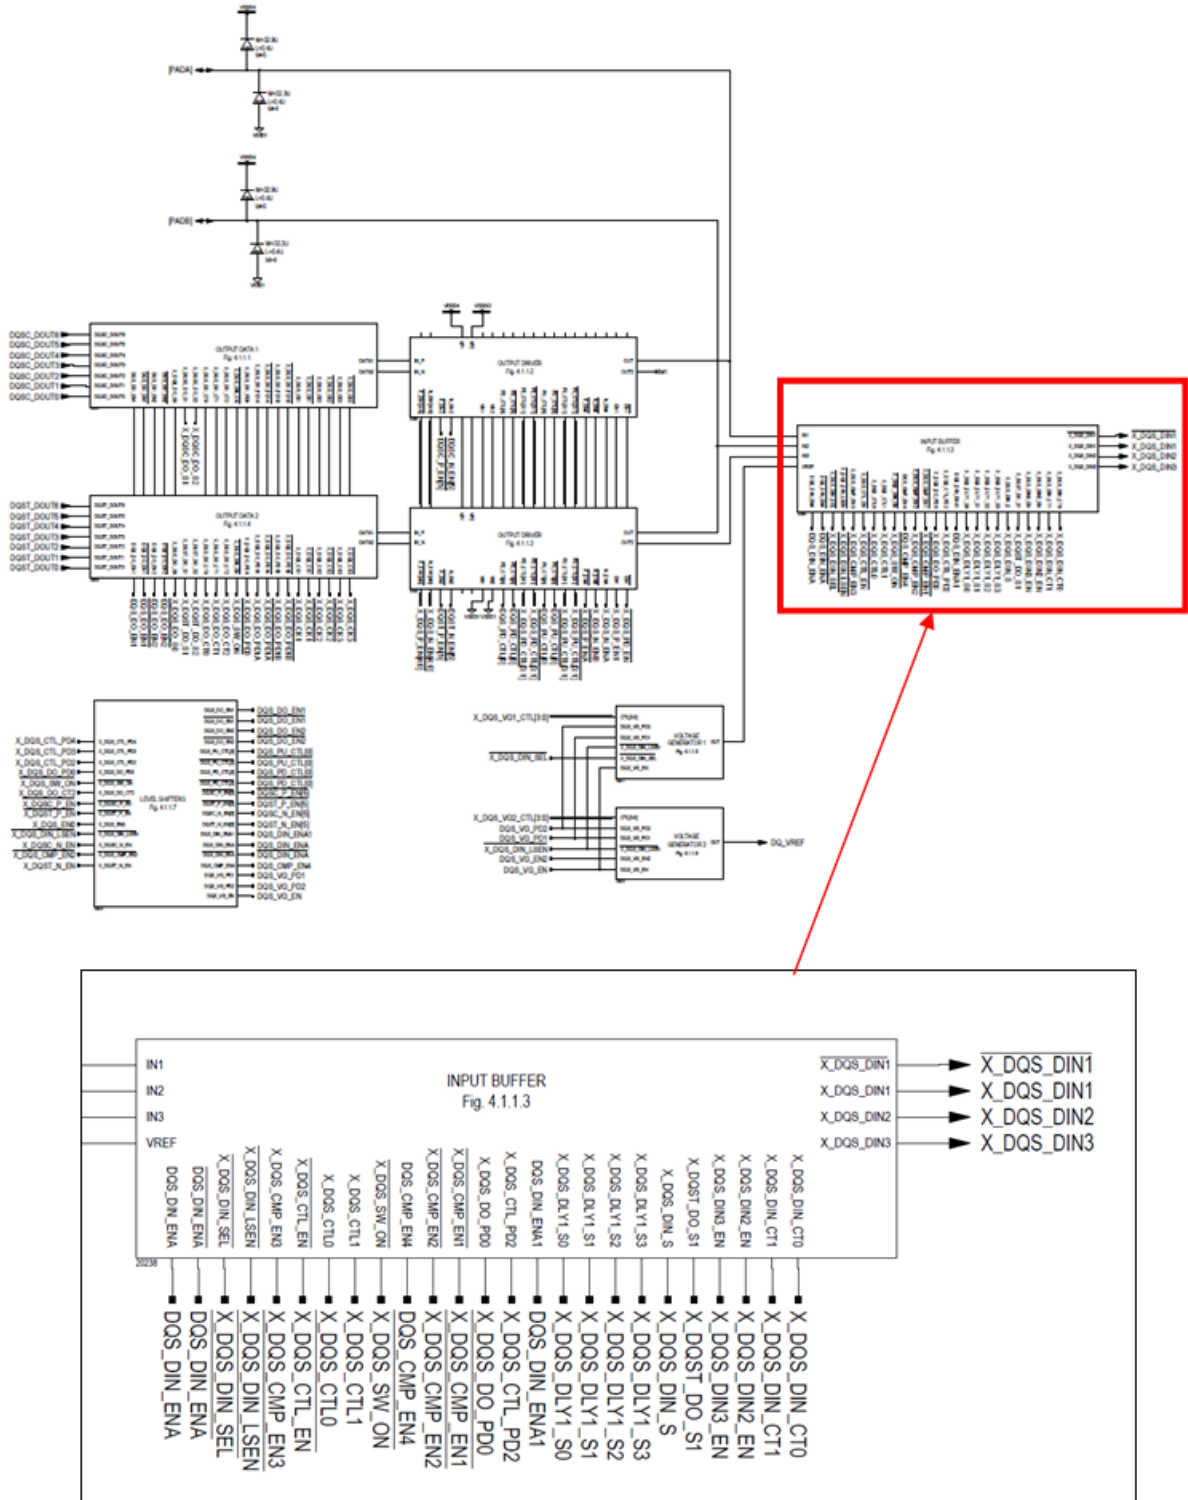
2149. All **’602 Accused Products** include a data strobe input buffer.

2150. For example, the **’602 Accused Products** practice **Claim 1(a) of the ’602 Patent** because they are designed to operate in conjunction with memory having input data strobes that support single-ended and differential signaling and thus require memory interface circuitry that includes a data strobe input buffer.

2151. For example, the Tegra 3 includes a memory controller and an EMC. **Tegra 3 TRM** at 10.

2152. For example, the Tegra 3 includes a memory interface. *See, e.g., Tegra 3 Report* at 10.

2153. As shown in the figures below, the memory interface of the Tegra 3 includes a data strobe input buffer.



Tegra 3 Report at 21 (red annotations added).

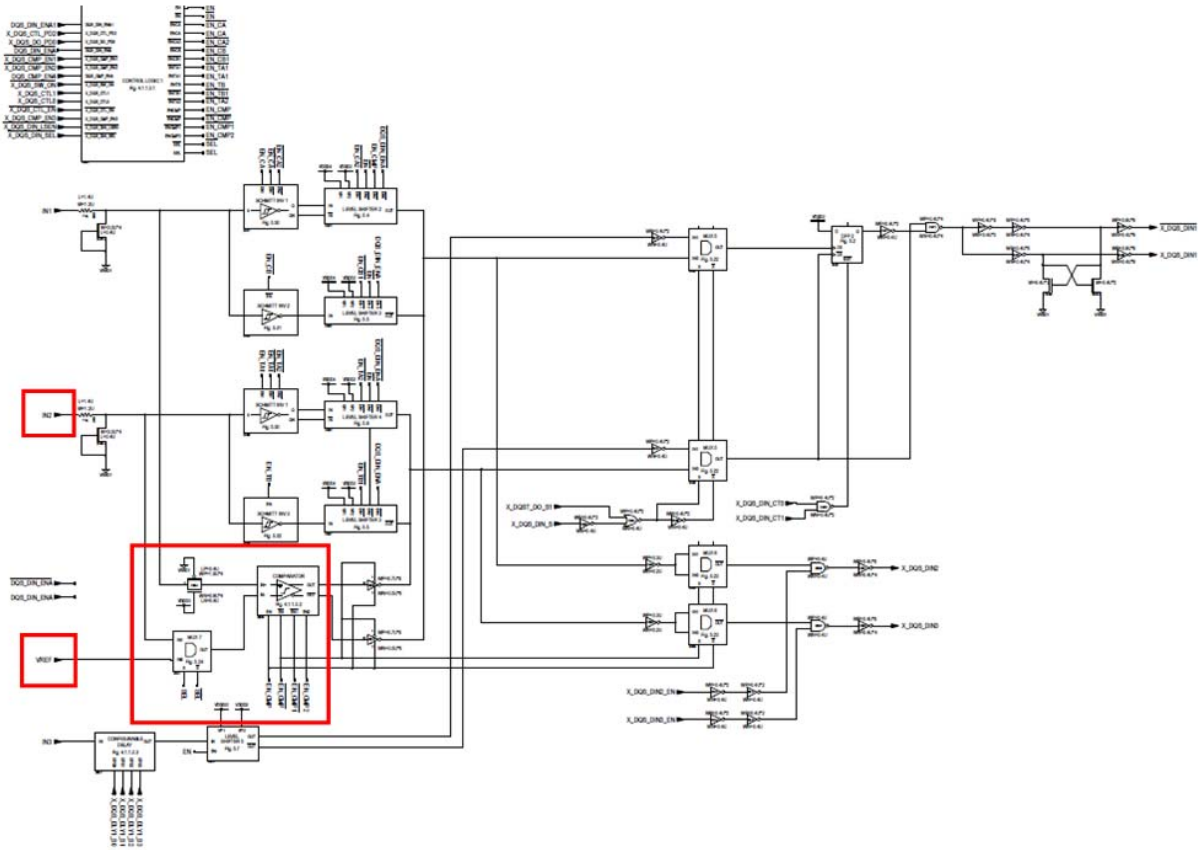
2154. As described in the preceding paragraphs, the **'602 Accused Products** include a data strobe input buffer

2155. “**Claim 1(b) of the '602 Patent**” recites “a differential amplifier circuit including at least two switches for passing an inverse data strobe signal or a reference voltage, respectively, depending on a level of a control signal, and a differential amplifier for receiving a data strobe signal and either the inverse data strobe signal or the reference voltage and outputting a differentially amplified signal.”

2156. All **'602 Accused Products** include a differential amplifier that includes at least two switches for passing an inverse data strobe signal or a reference voltage depending on the level of a control signal.

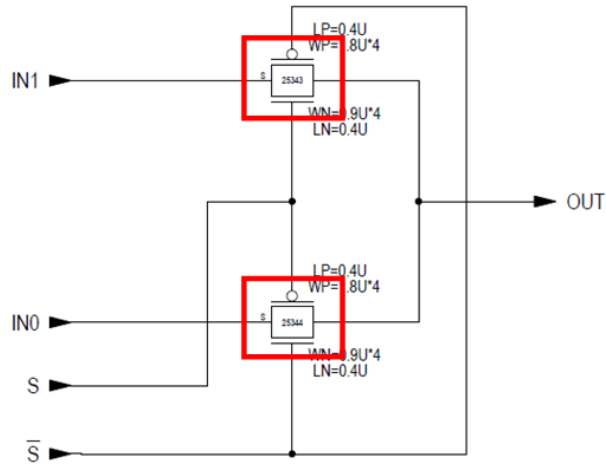
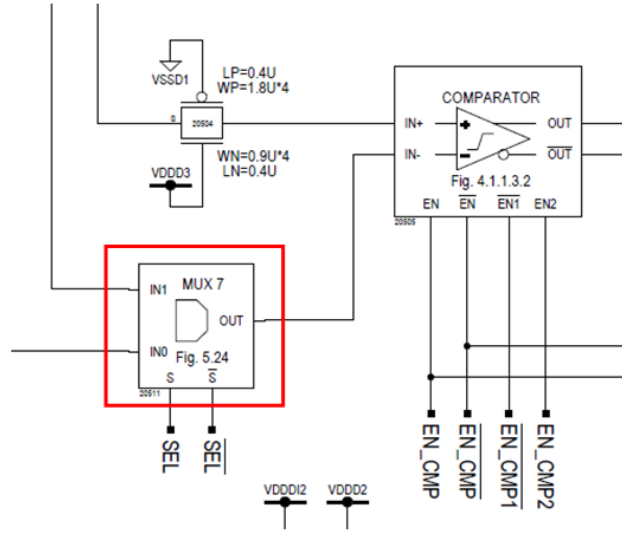
2157. All **'602 Accused Products** also include a differential amplifier that receives a data strobe signal and either an inverse data strobe signal or a reference voltage and outputs a differentially amplified signal.

2158. For example, as shown in the figure below, the Tegra 3 includes an input buffer with a differential amplifier circuit.




Tegra 3 Report at 44 (red annotations added).

2159. As shown in the figures below, the differential amplifier circuit includes two switches in MUX 7 for passing an inverse data strobe signal IN2 or a reference voltage VREF.



Tegra 3 Report at 44 (red annotations added).



Tegra 3 Report at 21 (red annotations added).

2160. The Tegra 3 is designed for use in conjunction with at least LP-DDR2 or DDR3 memory. *See, e.g., Tegra 3 TRM* at 10, 437.

2161. According to the JEDEC standard, DDR3 memory that operates in conjunction with the Tegra 3 utilizes an inverse data signal (which is a differential data strobe signal).

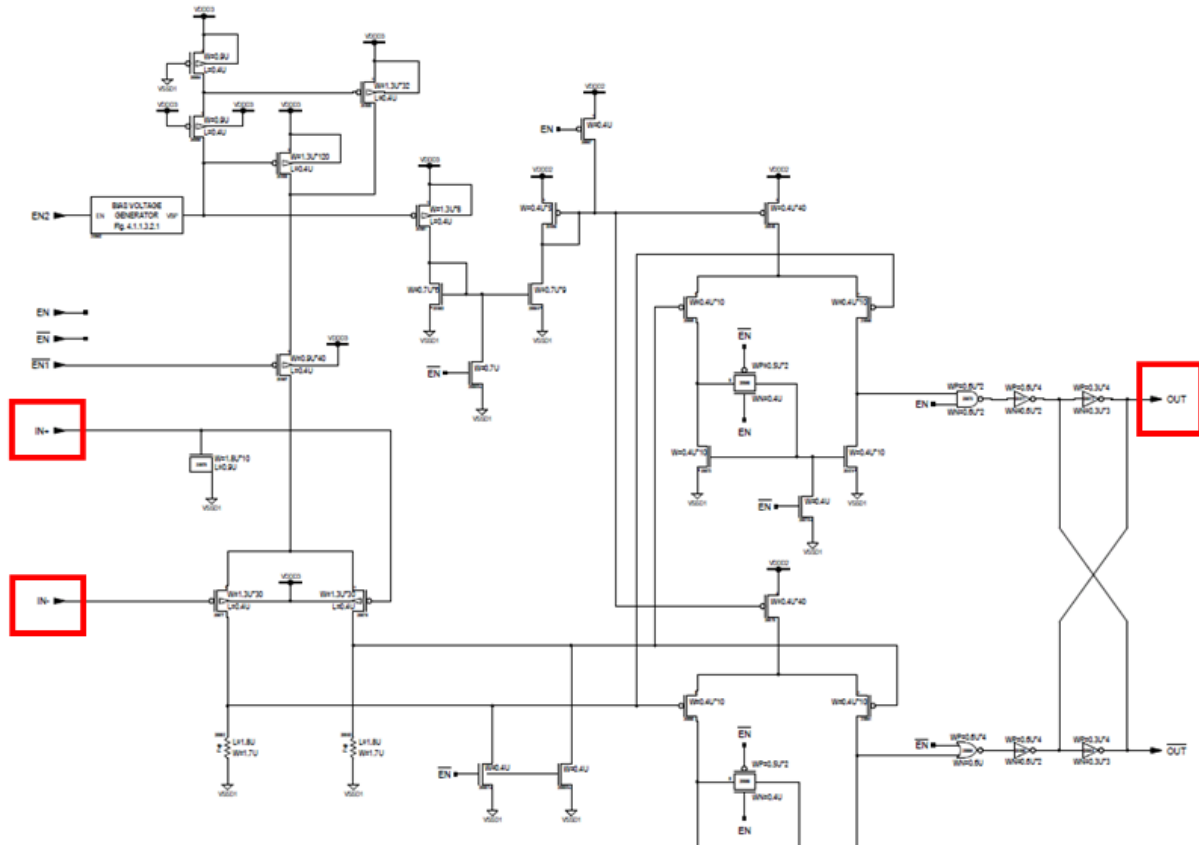
2162. For example, in DDR3 memory, the data strobe is differential. *See, e.g., JEDEC Standard No. 209-2F* at 13.

2163. Thus, the system must provide two signals, the DQS_t and DQS_c, for read and write.

2164. The DQS_c signal is the inverse of the DQS_t signal.

2165. The Tegra 3 uses a control signal S and /S. *See, e.g., Tegra 3 Report* at 44.

2166. As shown in the figure below, the Tegra 3 includes a comparator that receives a data strobe signal and either the inverse data strobe signal or the reference voltage and outputs a differentially amplified signal.



Tegra 3 Report at 46 (red annotations added).

2167. For example, in single mode, the reference voltage VREF is applied to a switch, and S and /S are low and high, respectively. *See, e.g., Tegra 3 Report* at 44.

2168. In dual mode, IN0 is provided to a second switch, and S and /S are high and low, respectively. *See, e.g., Tegra 3 Report* at 44; **Tegra 3 TRM** at 544.

2169. As described in the preceding paragraphs, the **'602 Accused Products** meet the limitations of **Claim 1(b) of the '602 Patent**.

2170. **“Claim 3(a) of the '602 Patent”** recites “[t]he data strobe input buffer of claim 1, wherein the data strobe input buffer is operable in both a single mode and a dual mode, wherein in said single mode, the reference voltage is applied to a first of the at least two switches and the level of the control signal is a first logic state and in said dual mode, the inverse data strobe

signal is provided to a second of the at least two switches and the level of the control signal is a second logic state.”

2171. All **'602 Accused Products** include a data strobe input buffer that is operable in both a single mode and a dual mode.

2172. All **'602 Accused Products** also include a data strobe input buffer that, when operating in single mode, the reference voltage is applied to a first of the at least two switches and the level of the control signal is a first logic state .

2173. All **'602 Accused Products** also include a data strobe input buffer that, when operating in dual mode, the inverse data strobe signal is provided to a second of the at least two switches and the level of the control signal is a second logic state.

2174. For example, in single mode, the reference voltage VREF is applied to a switch, and S and /S are low and high, respectively. *See, e.g., Tegra 3 Report* at 44.

2175. In dual mode, IN0 is provided to a second switch, and S and /S are high and low, respectively. *See, e.g., Tegra 3 Report* at 44.

2176. As described in the preceding paragraphs, the **'602 Accused Products** meet the limitations of **Claim 3(a) of the '602 Patent**.

2177. “**Claim 26(a) of the '602 Patent**” recites “[a] method of controlling propagation delay time of a semiconductor memory.”

2178. All **'602 Accused Products** practice a method of controlling the propagation delay time of the data input buffer.

2179. For example, the Tegra 3 includes a memory controller and an EMC. *See, e.g., Tegra 3 TRM* at 10.

2180. The memory controller controls the propagation delay time of the data input buffer. *See, e.g., Tegra 3 TRM* at 10.

2181. As described in the preceding paragraphs, the **'602 Accused Products** practice a method of controlling propagation delay time of a semiconductor memory.

2182. "**Claim 26(b) of the '602 Patent**" recites "receiving an inverse data strobe signal or a reference voltage, respectively, depending on a level of a control signal."

2183. All **'602 Accused Products** receive either an inverse data strobe signal or a reference voltage (depending on the level of a control signal).

2184. For example, the Tegra 3 includes a memory interface that includes a data strobe input buffer. *See, e.g., Tegra 3 Report* at 10, 21.

2185. For example, the Tegra 3 includes an input buffer with a differential amplifier circuit. *See, e.g., Tegra 3 Report* at 44.

2186. The differential amplifier circuit includes two switches in MUX 7 for passing an inverse data strobe signal IN2 or a reference voltage VREF. *See, e.g., Tegra 3 Report* at 44.

2187. DDR3 that operates in conjunction with the Tegra 3 utilizes an inverse data signal, that is, a differential data strobe signal. *See, e.g., JEDEC Standard No. 209-2F* at 13.

2188. The Tegra 3 uses a control signal S and /S. *See, e.g., Tegra 3 Report* at 44.

2189. As described in the preceding paragraphs, the **'602 Accused Products** meet the limitations of **Claim 26(b) of the '602 Patent**.

2190. "**Claim 26(c) of the '602 Patent**" recites "receiving a data strobe signal."

2191. All **'602 Accused Products** receive a data strobe signal.

2192. According to the JEDEC standard, DDR3 memory that operates in conjunction with the Tegra 3 utilizes an inverse data signal (which is a differential data strobe signal). *See, e.g., JEDEC Standard No. 209-2F* at 13.

2193. For example, in DDR3 memory, the data strobe is differential. *See, e.g., JEDEC Standard No. 209-2F* at 13.

2194. Thus, the system must provide two signals, the DQS_t and DQS_c, for read and write. *See, e.g., JEDEC Standard No. 209-2F* at 13.

2195. The DQS_c signal is the inverse of the DQS_t signal.

2196. As described in the preceding paragraphs, the **'602 Accused Products** meet the limitations of **Claim 26(c) of the '602 Patent**.

2197. "**Claim 26(d) of the '602 Patent**" recites "amplifying and outputting at least two different differentially amplified data strobe signals."

2198. All **'602 Accused Products** amplify and output at least two different differentially amplified data strobe signals.

2199. For example, the Tegra 3 includes a comparator that receives a data strobe signal and either the inverse data strobe signal or the reference voltage and outputs differentially amplified signals. *See, e.g., Tegra 3 Report* at 45.

2200. As described in the preceding paragraphs, the **'602 Accused Products** meet the limitations of **Claim 26(d) of the '602 Patent**.

2201. "**Claim 27 of the '602 Patent**" recites "[t]he method of claim 26, wherein in a single mode, the reference voltage is received and a level of the control signal is a first logic state and in a dual mode, the inverse data strobe signal is received and the level of the control signal is a second logic state."

2202. All **'602 Accused Products** receive a reference voltage where the level of the control signal is a first logic state when in a single mode.

2203. All **'602 Accused Products** receive an inverse data strobe signal where the level of the control signal is a second logic state when in a dual mode.

2204. For example, the Tegra 3 uses a control signal S and /S. *See, e.g., Tegra 3 Report* at 44.

2205. In single mode, the reference voltage VREF is applied to a switch, and S and /S are low and high, respectively. *See, e.g., Tegra 3 Report* at 44.

2206. In dual mode, IN0 is provided to a second switch, and S and /S are high and low, respectively. *See, e.g., Tegra 3 TRM* at 544.

2207. As described in the preceding paragraphs, the **'602 Accused Products** meet the limitations of **Claim 27 of the '602 Patent**.

2208. “**Claim 28 of the '602 Patent**” recites [t]he method of claim 27, wherein the control signal is received from an external source.”

2209. All **'602 Accused Products** include a data strobe input buffer where the control signal is received from an external source.

2210. For example, the Tegra 3 uses a control signal S and /S, which are received from a source outside the input buffer. *See, e.g., Tegra 3 Report* at 44.

2211. As described in the preceding paragraphs, the **'602 Accused Products** meet the limitations of **Claim 28 of the '602 Patent**.

2212. “**Claim 29(a) of the '602 Patent**” recites “[t]he method of claim 27.”

2213. As described in the preceding paragraphs, all **'602 Accused Products** use the method of claim 1 and thus meet the limitations of **Claim 29(a) of the '602 Patent**.

2214. “**Claim 29(b) of the ’602 Patent**” recites “receiving an external command and an address and generating the control signal, wherein a level of the control signal determines an operation mode of the semiconductor memory.”

2215. All **’602 Accused Products** receive an external command and an address and generate a control signal.

2216. The level of the control signal (the control signal described in the previous paragraph) determines an operational mode of the semiconductor memory in the **’602 Accused Products**.

2217. For example, the Tegra 3 receives a control signal S and /S from outside the SDRAM through, for example, the FBIO_CFG5 configuration register. *See, e.g., Tegra 3 Report* at 44.

2218. In single mode, the reference voltage VREF is applied to a switch, and S and /S are low and high, respectively. *See, e.g., Tegra 3 Report* at 44.

2219. In dual mode, IN0 is provided to a second switch, and S and /S are high and low, respectively. *See, e.g., Tegra 3 TRM* at 544.

2220. The Tegra 3 includes a mode to select single ended or differential data strobes. *See, e.g., Tegra 3 TRM* at 544.

2221. As described in the preceding paragraphs, the **’602 Accused Products** meet the limitations of **Claim 29(b) of the ’602 Patent**.

2222. NVIDIA has directly infringed and continue to directly infringe the **’602 Patent** by making, using, offering to sell, selling, and or importing the **’602 Accused Products**.

2223. NVIDIA uses the **’602 Accused Products**.

2224. For example, NVIDIA uses the **'602 Accused Products** for demonstration, testing, development, and configuration purposes.

2225. NVIDIA has used and continues to use its own products, including the **'602 Accused Products**, such as for demonstration, testing, development, and configuration purposes.

2226. For example and without limitation, NVIDIA has used and continues to use the GeForce G100 GPU.

2227. NVIDIA also has used and continues to use third-party products, including the **'602 Accused Products**, such as for demonstration, testing, and development purposes.

2228. When NVIDIA uses the **'602 Accused Products**, such as for testing or development purposes, such use directly infringes the **'602 Patent**.

2229. NVIDIA has in the past and continues to import, offer to sell, and sell its own products, including the **'602 Accused Products**.

2230. For example, NVIDIA has in the past and continues to import, offer to sell, and sell the GeForce G100 GPU.

2231. NVIDIA sells the **'602 Accused Products** to end users and customers.

2232. For example, NVIDIA advertises the GeForce G100 GPU on its website. *See, e.g., GeForce G100 GPU Page.*

2233. NVIDIA sells the **'602 Accused Products** to intermediate customers.

2234. For example, NVIDIA sells its products to distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers. **NVIDIA 2014 Form 10-K** at 8.

2235. The products NVIDIA sells to intermediate customers, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design

manufacturers, motherboard manufacturers, and add-in board manufacturers, include one or more of the **'602 Accused Products**.

2236. For example, graphics boards that incorporate the GeForce G100 GPU are advertised for sale on third party websites, such as **GeForce G100 Purchase Page**.

2237. When NVIDIA imports into the United States or offers for sale or sells in the United States the **'602 Accused Products**, such activity directly infringes the **'602 Patent**.

2238. Velocity has directly infringed and continue to directly infringe the **'602 Patent** by making, using, offering to sell, selling, and or importing the **'602 Accused Products**.

2239. Velocity sells computers that incorporate the **'602 Accused GPUs**.

2240. For example, Velocity sells the Vector Z35, which includes one of the **'602 Accused GPUs**.

2241. The Vector Z35 incorporates the NVIDIA GeForce 9500 GT, which can be configured to operate with GDDR2 memory.

2242. Velocity has used and continues to use its own products, including the **'602 Accused Products**, such as for demonstration, testing, development, and configuration purposes.

2243. For example and without limitation, Velocity has used and continues to use the Vector Z35.

2244. When Velocity uses the **'602 Accused Products**, such as for testing or development purposes, such use directly infringes the **'602 Patent**.

2245. Velocity has in the past and continues to import, offer for sale, and sell its own products, including the **'602 Accused Products**.

2246. When Velocity imports, offers for sale, or sells the **'602 Accused Products**, such activity directly infringes the **'602 Patent**.

2247. NVIDIA has had actual knowledge of the **'602 Patent** at least as of November 12, 2014, the date the original complaint in this action was served.

2248. Velocity has had actual knowledge of the **'602 Patent** at least as of November 11, 2014, the date the original complaint in this action was served.

2249. NVIDIA indirectly infringes the **'602 Patent** by inducing infringement by others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(b), in this District and elsewhere in the United States.

2250. Direct infringement is the result of activities performed by others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, by, for example, making, using, offering to sell, selling, and or importing the **'602 Accused Products**.

2251. NVIDIA and Velocity have induced and continue to induce infringement of the **'602 Patent** by intending that others infringe the **'602 Patent** by using the **'602 Accused Products**. NVIDIA and Velocity designed the **'602 Accused Products** such that they would each infringe one or more claims of the **'602 Patent** when used.

2252. NVIDIA and Velocity provide the **'602 Accused Products** to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users.

2253. By providing **'602 Accused Products** to others, NVIDIA and Velocity intend for **'602 Accused Products** to be used in the United States.

2254. NVIDIA and Velocity also provide others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, instructions, user guides, and technical specifications.

2255. When others follow such instructions, user guides, and/or other design documentation, they directly infringe one or more claims of the **'602 Patent**.

2256. NVIDIA and Velocity know that by providing such instructions, user guides, and/or other design documentation, others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, follow those instructions, user guides, and other design documentation, and directly infringe one or more claims of the **'602 Patent**.

2257. NVIDIA and Velocity thus know that their actions actively induce infringement.

2258. NVIDIA sells the **'602 Accused Products** directly to customers and end users.

2259. For example, NVIDIA advertises the **GeForce G100 GPU**. *See, e.g., GeForce G100 GPU Page*.

2260. Graphics boards that incorporate the **GeForce G100 GPU** are advertised for sale on third party websites. *See, e.g., GeForce G100 Purchase Page*.

2261. Through such sales, NVIDIA specifically intends that others, such as customers and end users, will infringe one or more claims of the **'602 Patent**.

2262. NVIDIA provides the **'602 Accused Products**, and technical specifications for the **'602 Accused Products**, to others, such as distributors, wholesalers, retailers, original

equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users.

2263. Through such activity, NVIDIA specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, infringe the **'602 Patent** by making, using, offering to sell, and selling in the United States, and importing into the United States the **'602 Accused Products**.

2264. Through its manufacture (either directly, or through third-parties) and/or sale of the infringing products, NVIDIA specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, will infringe one or more claims of the **'602 Patent**.

2265. NVIDIA specifically intends for others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'602 Patent** in the United States.

2266. For example, NVIDIA advertised its "NVIDIA Authorized Board Partner Program ensures an exceptional customer experience when purchasing graphics cards and motherboards manufactured by partners that make use of NVIDIA's latest technologies." **NVIDIA PartnerForce Info**.

2267. NVIDIA claims that its "NVIDIA Authorized Board Partners offer the latest technologies from NVIDIA" and lists six Authorized Board Partners in the United States, six in Canada, and none in any other countries. *See* **NVIDIA PartnerForce Info**. NVIDIA also lists

ten distributors in the United States (including “pre and post sales technical support”) and seven in Canada but none in any other country. *See* **NVIDIA PartnerForce Info.**

2268. As an additional example, as of October 31, 2014, NVIDIA also advertised the “NVIDIA PartnerForce Program” which is “a sales and marketing program for value-added resellers, system builders, etailers and retailers who sell NVIDIA based components or systems.” **NVIDIA PartnerForce Program.**

2269. NVIDIA advertises in the United States and in this judicial district that “[b]y joining the PartnerForce Program, you may leverage the world-class brands and technology platforms from NVIDIA, sales and technical support from our experienced team and hundreds of turnkey sales and marketing tools.” **NVIDIA PartnerForce Program.**

2270. NVIDIA also provides marketing materials and templates, including retail display items, web banners, copy blocks, logos, product shots, email and web page templates, to members of the “NVIDIA PartnerForce Program” which includes value-added resellers, system builders, and retailers inducing them to offer to sell and the **'602 Accused Products**. *See* **NVIDIA PartnerForce Program.**

2271. At the Ford Event and Conference Center, in Dearborn, Michigan, NVIDIA showcased its latest processor technologies, which power everything from the CAD software that designers use to style cars to the infotainment systems that drivers use to map their trips and listen to music, in an effort to encourage domestic car manufacturers to include infringing technology in vehicles manufactured and sold in the United States. *See* **Traveling The Road To Silicon Motown.**

2272. As described in the preceding paragraphs, NVIDIA specifically targets the United States market for **'602 Accused Products** actively induces others, such as distributors, original

equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'602 Patent**.

2273. Velocity sells the **'602 Accused Products** directly to customers and end users.

2274. For example, Velocity advertised the Vector Z35, which incorporates the NVIDIA GeForce 9500 GT Graphics Card, on its website. *See, e.g., Configure Your Z35.*

2275. Through such sales, Velocity specifically intends that others, such as customers and end users, will infringe one or more claims of the **'602 Patent**.

2276. Velocity provides the **'602 Accused Products**, and technical specifications for the **'602 Accused Products**, to others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users.

2277. Through such activity, Velocity specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, infringe the **'602 Patent** by making, using, offering to sell, and selling in the United States, and importing into the United States the **'602 Accused Products**.

2278. Through its manufacture (either directly, or through third-parties) and/or sale of the infringing products, Velocity specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, motherboard manufacturers, resellers, customers, and end users, will infringe one or more claims of the **'602 Patent**.

2279. NVIDIA and Velocity indirectly infringe the **'602 Patent** by contributing to infringement by others, such as distributors, wholesalers, retailers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(c), in this District and elsewhere in the United States.

2280. Direct infringement is the result of activities performed by the distributors, wholesalers, retailers, resellers, customers, and end users of the infringing products.

2281. For example, the **'602 Accused Products** allow for a method of controlling the propagation delay time of a semiconductor memory.

2282. When the infringing products are used as intended by distributors, wholesalers, retailers, resellers, customers, and end users, the **'602 Accused Products** necessarily control the propagation delay time of a semiconductor memory in an infringing manner.

2283. The infringing products cannot operate in an acceptable manner without controlling the propagation delay time of a semiconductor memory as claimed in the **'602 Patent**.

2284. From the facts set forth above, it is evident that NVIDIA and Velocity knew that the ability to control the propagation delay time of a semiconductor memory in the infringing products is especially made or especially adapted to operate in the products of NVIDIA's distributors, wholesalers, retailers, resellers, customers, and end users, and is not a staple article or commodity of commerce and that its infringing use is required for operation of the infringing products. Any other use would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental.

2285. NVIDIA's infringing products, with the ability to control the propagation delay time of a semiconductor memory, are each a material part of the invention of the **'602 Patent** and are especially made or adapted to infringe one or more claims of the **'602 Patent**.

2286. Because the use of the **'602 Accused Products** necessarily infringes one or more claims of the **'602 Patent**, NVIDIA's sales of its infringing products have no substantial non-infringing uses.

2287. Accordingly, NVIDIA offers to sell, or sells a component, material, or apparatus for use in practicing one or more claims of the **'602 Patent** knowing the same to be especially made or especially adapted for use in an infringement of such patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

FIFTH CLAIM FOR RELIEF
INFRINGEMENT OF U.S. PATENT NO. 8,252,675
(AGAINST NVIDIA AND VELOCITY)

2288. Each of the above listed paragraphs is incorporated herein by reference and adopted, as if fully set forth again.

2289. The **'675 Patent** was filed on November 9, 2010, issued on August 28, 2012, and is entitled "Methods Of Forming CMOS Transistors With High Conductivity Gate Electrodes." The **'675 Patent** is generally directed to methods of manufacturing a metal-oxide-semiconductor ("MOS") transistor.

2290. The **'675 Patent** was assigned to SEC, and SEC continues to hold all rights, title, and interest in the **'675 Patent**. A true and correct copy of the **'675 Patent** is attached hereto as Exhibit E.

2291. For the **'675 Accused Products**, all **Accused 28 nm GPUs** and **Accused 28 nm SOCs** infringe the **'675 Patent** in substantially the same way.

2292. The **'675 Accused Products** are manufactured using one of TSMC's 28 nanometer high-k, metal gate processes ("**HKMG Processes**").

2293. The TSMC 28 nanometer **HKMG Processes** used to manufacture NVIDIA's 28 nanometer products are:

- 28HP (high performance with HKMG),
- 28HPL (low power with HKMG),
- 28HPM (high performance for mobile computing), or
- 28HPC (high performance compact).

See TSMC 28HPC Process.

2294. TSMC states that “[t]he 28nm high performance (HP) process is the first option to use high-k metal gate process technology.” **TSMC 28nm Technology.**

2295. TSMC uses the 28HP (high performance with HKMG) process in fabricating certain **'675 Accused Products.**

2296. TSMC states that the 28HPL process “adopts the same gate stack as HP technology.” **TSMC 28nm Technology.**

2297. TSMC uses the 28HPL (low power with HKMG) process in fabricating certain **'675 Accused Products.**

2298. TSMC states that the 28HPC (high performance compact) process is a “compact version of 28HPM.” **TSMC 28HPC Process.**

2299. TSMC uses the 28HPM (high performance for mobile computing) process in fabricating certain **'675 Accused Products.**

2300. TSMC uses the 28HPC process in fabricating certain **'675 Accused Products.**

2301. TSMC states that its **HKMG Processes** use a “replacement metal gate” technology. *See History on Side of Gate-Last High-k Approach.*

2302. The **Accused 28 nm GPUs** and **Accused 28 nm SOCs** include transistors formed by depositing one or more polysilicon gates that are replaced by metal layers.

2303. The **Accused 28 nm GPUs** and **Accused 28 nm SOCs** include spacers surrounding the metal gates. *See, e.g., GK107 Report* at 36 and 47, Figures 2.3.1.4 and 2.3.1.13.

2304. The GK107 GPU is one of the **Accused 28 nm GPUs**.

2305. The GK107 GPU is made by TSMC’s 28HP fabrication process. *See GK107 Report* at ix; *see also NVIDIA GeForce GTX 660 Features* (listing the “Kepler GPU Architecture” as a feature).

2306. NVIDIA states that “Kepler is based on 28-nanometer (nm) process technology and succeeds the 40-nm NVIDIA Fermi architecture, which was first introduced into the market in March 2010.” **NVIDIA Kepler Press Release**.

2307. The **GK107 Report** accurately depicts the gate stack for TSMC’s 28 nanometer **HKMG Processes**:




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.




Figure 2.3.1.9: Logic PMOS transistor gate stack, TEM cross-section.

GK107 Report at 36 and 42, Figure 2.3.1.4 and 2.3.1.9.

2308. “**Claim 1(a) of the ’675 Patent**” recites “[a] method of forming an insulated-gate transistor.”

2309. The **’675 Accused Products** are made by a process that includes forming insulated gate transistors.

2310. For example, the GK107 GPU includes one or more NMOS-type insulated transistors.

2311. The GK107 GPU includes one or more PMOS-type insulated transistors.

2312. The GK107 GPU includes one or more input/output (“I/O”) insulated transistors.

2313. The **GK107 Report** depicts topographical and cross sectional views of PMOS, NMOS, and I/O transistors. *See, e.g., GK107 Report* at 33, 35, and 41, Figures 2.3.1.1, 2.3.1.3, and 2.3.1.8.

2314. As described in the preceding paragraphs, the **’675 Accused Products** are made by a process that includes forming an insulated-gate transistor.

2315. “**Claim 1(b) of the ’675 Patent**” recites “forming a gate insulating layer on a substrate.”

2316. The **’675 Accused Products** are made by a process that includes forming a gate insulating layer for transistors on a substrate.

2317. For example, the GK107 GPU includes insulating silicon oxide on the substrate for PMOS, NMOS, and I/O transistors.

2318. The GK107 GPU additionally includes insulating hafnium oxide on the silicon oxide for PMOS, NMOS, and I/O transistors.

2319. The silicon oxide and hafnium oxide on the silicon oxide are insulating materials.

2320. The silicon oxide and hafnium oxide for PMOS, NMOS, and I/O transistors of the GK107 GPU form a gate oxide. *See, e.g., GK107 Report* at 101 and 106, Figures 3.2.6 and 3.2.11.

2321. As shown below, the **GK107 Report** depicts the silicon oxide and hafnium oxide on the substrate for PMOS, NMOS, and I/O transistors.




Figure 2.3.1.6: Logic NMOS transistor, lattice fringe image, TEM cross-section.




Figure 2.3.1.12: Logic PMOS transistor, lattice fringe image, TEM cross-section.




Figure 2.3.2.3: I/O transistor with HV gate oxide, TEM cross-section.

GK107 Report at 38, 45, and 53, Figures 2.3.1.6, 2.3.1.12, and 2.3.2.3. (red annotations added).

2322. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(b) of the '675 Patent**.

2323. **“Claim 1(c) of the '675 Patent”** recites “forming a metal buffer gate electrode layer on the gate insulating layer.”

2324. The **'675 Accused Products** are made by a process that includes depositing, on the insulating layer, a metal buffer gate electrode layer.

2325. For example, the **GK107 GPU** includes a titanium nitride layer on the hafnium oxide for PMOS, NMOS, and I/O transistors.

2326. The titanium nitride layer is an electrically conductive layer.

2327. As shown below, the **GK107 Report** depicts the titanium nitride layers on the hafnium oxide for PMOS, NMOS, and I/O transistors.




Figure 2.3.1.6: Logic NMOS transistor, lattice fringe image, TEM cross-section.




Figure 2.3.1.12: Logic PMOS transistor, lattice fringe image, TEM cross-section.




Figure 2.3.2.3: I/O transistor with HV gate oxide, TEM cross-section.

GK107 Report at 38, 45, and 53, Figures 2.3.1.6, 2.3.1.12, and 2.3.2.3. (red annotations added).

2328. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(c) of the '675 Patent**.

2329. “**Claim 1(d) of the '675 Patent**” recites “forming a dummy gate electrode layer on the buffer gate electrode layer, said dummy gate electrode layer and said buffer gate electrode layer comprising different materials.”

2330. The **'675 Accused Products** are made by a process that includes forming a dummy gate on the buffer gate electrode layer. The dummy gate comprises a different material than the buffer gate electrode layer.

2331. For example, NMOS, PMOS, and I/O transistors in the GK107 GPU are formed by a process that includes depositing a dummy polysilicon gate on the titanium nitride layer.

2332. The dummy polysilicon and the titanium nitride layer comprise different materials.

2333. The **GK107 Report** describes the formation of dummy polysilicon gates in the GK107 GPU. *See GK107 Report* at 15.

2334. As described above, TSMC uses a gate-last deposition method for the high-k/metal gate stack of its 28 nm transistors. *See History on Side of Gate-Last High-k Approach.*

2335. TSMC claims its gate-last deposition method was guided by history. *See History on Side of Gate-Last High-k Approach.*

2336. TSMC's gate-last deposition method uses a replacement metal gate (RMG) technique. *See History on Side of Gate-Last High-k Approach.*

2337. TSMC's replacement metal gate technique is a process that includes forming a polysilicon gate for transistors and replacing the polysilicon gate with metal gates. *See GK107 Report* at 15.

2338. The replacement of polysilicon gates with metal gates has the advantage that metal gates are introduced only at the end of the Front-End-of-Line (FEOL) processing flow so that the metal is not exposed to the high temperature budget during junction activation anneal. *See GK107 Report* at 15.

2339. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(d) of the '675 Patent**.

2340. "**Claim 1(e) of the '675 Patent**" recites "patterning the dummy gate electrode layer and the buffer gate electrode layer in sequence to define a buffer gate electrode on the gate insulating layer and a dummy gate electrode on the buffer gate electrode."

2341. The **'675 Accused Products** are made by a process that includes patterning the buffer gate electrode layer and the dummy gate in sequence on the insulating layer, thereby

defining a buffer gate electrode on the insulating layer and defining a dummy gate on the buffer gate electrode.

2342. For example, transistors in the GK107 GPU are formed by a process that includes sequentially patterning the polysilicon dummy gate on the titanium nitride layer.

2343. The deposition of the polysilicon in the GK107 forms a sacrificial dummy gate.

See **GK107 Report** at x.

2344. Transistors in the GK107 GPU are additionally formed by patterning the titanium nitride layer on the insulating silicon oxide and insulating hafnium oxide.

2345. The patterning of the titanium nitride layer on the insulating silicon oxide and hafnium oxide in the GK107 defines a buffer gate electrode.

2346. The patterning of the titanium nitride layer and the polysilicon layer in sequence in the GK107 defines a dummy gate electrode.

2347. The **GK107 Report** describes the patterning process for the polysilicon layer in the GK107 GPU. *See* **GK107 Report** at x.

2348. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(e) of the '675 Patent**.

2349. "**Claim 1(f) of the '675 Patent**" recites "forming electrically insulating spacers on sidewalls of the dummy gate electrode and on sidewalls of the buffer gate electrode."

2350. The **'675 Accused Products** are made by a process that includes depositing insulating spacers on sidewalls of the buffer gate electrode and the dummy gate.

2351. For example, the GK107 GPU includes sidewalls spacers formed out of silicon nitride. *See* **GK107 Report** at x.

2352. The silicon nitride sidewall spacers are electrically insulating.

2353. The sidewall spacers of the GK107 GPU are formed on the left and right sides of the titanium nitride and the dummy polysilicon gate for NMOS, PMOS, and I/O transistors. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

2354. As shown below, the **GK107 Report** depicts sidewall spacers for transistors in the GK107 GPU.




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.

GK107 Report at 36, Figure 2.3.1.4; *see also GK107 Report* at 47, Figure 2.3.1.13.

2355. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(f) of the '675 Patent**.

2356. “**Claim 1(g) of the '675 Patent**” recites “covering the spacers and the dummy gate electrode with an electrically insulating mold layer.”

2357. The **'675 Accused Products** are made by a process that includes depositing an electrically insulating mold layer covering the dummy gate and sidewall spacers.

2358. For example, the process of forming transistors in the GK107 GPU includes depositing insulating mold material covering the sidewall spacers. *See GK107 Report* at x.

2359. The insulating mold material in the GK107 GPU is additionally deposited such that it covers the dummy gate. *See* **GK107 Report** at x.

2360. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(g) of the '675 Patent**.

2361. "**Claim 1(h) of the '675 Patent**" recites "removing an upper portion of the mold layer to expose an upper surface of the dummy gate electrode."

2362. The **'675 Accused Products** are made by a process that includes exposing the upper surface of the dummy gate by removing the insulating mold layer.

2363. For example, the process of forming transistors in the GK107 GPU includes removing a portion of the mold material to expose the upper surface of the dummy polysilicon gate. *See, e.g.,* **GK107 Report** at x.

2364. By removing the mold material and exposing the dummy polysilicon gate, the dummy polysilicon gate may be removed. *See, e.g.,* **GK107 Report** at x.

2365. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(h) of the '675 Patent**.

2366. "**Claim 1(i) of the '675 Patent**" recites "removing the dummy gate electrode from between the spacers by selectively etching back the dummy gate electrode using the mold layer and the spacers as an etching mask."

2367. The **'675 Accused Products** are made by process that includes removing the exposed dummy gate by selective etching using the mold layer and sidewall spacers as an etching mask.

2368. For example, transistors in the GK107 GPU are formed by a process that includes removing the exposed polysilicon gate. *See* **GK107 Report** at 15, 39.

2369. The polysilicon gate of the GK107 GPU is removed by using the sidewall spacers and the insulating mold material as an etching mask.

2370. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(i) of the '675 Patent**.

2371. "**Claim 1(j) of the '675 Patent**" recites "depositing a first metal layer onto an upper surface of the mold layer and onto inner sidewalls of the spacers and onto an upper surface of the buffer gate electrode."

2372. The **'675 Accused Products** are made by a process that includes depositing a first metal layer onto the upper surface of the buffer gate electrode, on the inner sidewalls of the sidewall spacers, and onto the insulating mold layer.

2373. For example, the removal of the dummy gate of the GK107 GPU exposes the top surface of the titanium nitride layer and exposes inner sidewalls of the sidewall spacers.

2374. The transistors of the GK107 GPU include a first metal layer that is deposited above the exposed top surface of the titanium nitride.

2375. A portion of the first first metal layer is deposited on exposed inner walls of the sidewall spacers.

2376. The first metal layer is also deposited on insulating mold material.

2377. The deposited first metal layer includes a portion having a "U" shaped cross-section.

2378. As shown below, the **GK107 Report** depicts the first metal layer of transistors in the GK107 GPU.




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.

GK107 Report at 36, Figure 2.3.1.4; *see also id.* at 42, 53, Figures 2.3.1.9, and 2.3.2.3.

2379. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(j) of the '675 Patent**.

2380. “**Claim 1(k) of the '675 Patent**” recites “filling a space between the inner sidewalls of the spacers by depositing a second metal layer onto a portion of the first metal layer extending between the inner sidewalls of the spacers to thereby define a metal gate electrode comprising a composite of the second metal layer, a portion of the first metal layer having a U-shaped cross-section and the buffer gate electrode.”

2381. The **'675 Accused Products** are made by a process that includes depositing a second metal layer onto the first metal layer to fill a space between the sidewall spacers. The filled space defines a composite metal gate electrode including the buffer gate electrode, a portion of the first metal layer having a U-shaped cross section, and the second metal layer.

2382. For example, transistors in the GK107 GPU include a second metal layer formed by deposition above the first metal layer and into the space between the sidewall spacers.

2383. After filling the gate with the second metal layer, transistors of the GK107 GPU comprise a composite including the second metal layer, a portion of the first metal layer having a U-shaped cross-section, and the buffer gate electrode.

2384. As shown below, the **GK107 Report** depicts the second metal layer for transistors of the GK107 GPU.




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.

GK107 Report at 36, Figure 2.3.1.4.

2385. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 1(k) of the '675 Patent**.

2386. **“Claim 2 of the '675 Patent”** recites “[t]he method of claim 1, wherein the second metal layer comprises aluminum, the first metal layer comprises titanium nitride and the buffer gate electrode comprises titanium nitride.”

2387. The **'675 Accused Products** are made by a process wherein the buffer gate electrode includes titanium nitride, the first metal layer includes titanium nitride, and the second metal layer includes aluminum.

2388. For example, the buffer gate electrode of transistors in the GK107 GPU includes titanium nitride. *See, e.g., GK107 Report* at 37, Figure 2.3.1.5.

2389. For example, the first metal layer of transistors in the GK107 GPU includes titanium nitride. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

2390. For example, the second metal layer of transistors in the GK107 GPU includes aluminum. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

2391. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 2 of the '675 Patent**.

2392. "**Claim 3 of the '675 Patent**" recites "[t]he method of claim 1, wherein the insulated-gate transistor is a PMOS transistor; and wherein the gate insulating layer comprises hafnium oxide."

2393. The **'675 Accused Products** are made by a process that includes forming a PMOS insulated gate transistor having gate insulating layers containing hafnium oxide.

2394. For example, the GK107 GPU includes PMOS transistors that contain insulating layers comprising hafnium oxide. *See GK107 Report* at 45, Figure 2.3.1.12.

2395. As shown below, the **GK107 Report** depicts PMOS transistors containing hafnium oxide.




Figure 2.3.1.12: Logic PMOS transistor, lattice fringe image, TEM cross-section.

GK107 Report at 45, Figure 2.3.1.12. (red annotation added).

2396. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 3 of the '675 Patent**.

2397. **"Claim 4 of the '675 Patent"** recites "[t]he method of claim 1, wherein the dummy gate electrode comprises polysilicon."

2398. The **'675 Accused Products** are made by a process that includes forming a dummy gate electrode comprising polysilicon.

2399. For example, transistors in the GK107 GPU are formed by a process that includes depositing a dummy polysilicon gate on the titanium nitride layer.

2400. The **GK107 Report** describes the formation of dummy polysilicon electrodes in the GK107 GPU. *See GK107 Report* at 15.

2401. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 4 of the '675 Patent**.

2402. “**Claim 5 of the ’675 Patent**” recites “[t]he method of claim 1, wherein the buffer gate electrode comprises titanium nitride or tantalum nitride.”

2403. The **’675 Accused Products** are made by a process that includes forming buffer gate electrodes that comprise titanium nitride or tantalum nitride.

2404. For example, the GK107 GPU includes a titanium nitride layer on the hafnium oxide for PMOS, NMOS, and I/O transistors.

2405. The **GK107 Report** depicts the titanium nitride layers for PMOS, NMOS, and I/O transistors. *See GK107 Report* at 38, 45, and 53, Figures 2.3.1.6, 2.3.1.12, and 2.3.2.3.

2406. As described in the preceding paragraphs, the **’675 Accused Products** meet the limitations of **Claim 5 of the ’675 Patent**.

2407. “**Claim 6(a) of the ’675 Patent**” recites “[a] method of forming an integrated circuit device.”

2408. The **’675 Accused Products** are made by a process that includes a method of forming an integrated circuit devices.

2409. For example, the GK107 GPU is an integrated circuit device.

2410. The GK107 GPU includes hundreds of millions of NMOS and PMOS insulated gate transistors, and many I/O transistors, all of which form an integrated circuit device.

2411. The **GK107 Report** depicts the die of the GK107 GPU. *See GK107 Report* at 6, Figure 1.2.2.

2412. As described in the preceding paragraphs, the **’675 Accused Products** are made by a process that includes a method of forming an integrated circuit device.

2413. “**Claim 6(b) of the ’675 Patent**” recites “forming a gate insulating layer on a substrate.”

2414. The '675 **Accused Products** are made by a process that includes forming a gate insulating layer for transistors on a substrate.

2415. For example, the GK107 GPU includes insulating silicon oxide on the substrate for PMOS, NMOS, and I/O transistors.

2416. The GK107 GPU additionally includes insulating hafnium oxide on the silicon oxide for PMOS, NMOS, and I/O transistors.

2417. The silicon oxide and hafnium oxide on the silicon oxide are insulating materials.

2418. The silicon oxide and hafnium oxide for PMOS, NMOS, and I/O transistors of the GK107 GPU form a gate oxide. *See, e.g., GK107 Report* at 101, 106.

2419. As shown below, the **GK107 Report** depicts the silicon oxide and hafnium oxide on the substrate for PMOS, NMOS, and I/O transistors.




Figure 2.3.1.6: Logic NMOS transistor, lattice fringe image, TEM cross-section.




Figure 2.3.1.12: Logic PMOS transistor, lattice fringe image, TEM cross-section.




Figure 2.3.2.3: I/O transistor with HV gate oxide, TEM cross-section.

GK107 Report at 38, 45, and 53, Figures 2.3.1.6, 2.3.1.12, and 2.3.2.3. (red annotations added).

2420. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 6(b) of the '675 Patent**.

2421. “**Claim 6(c) of the '675 Patent**” recites “forming a first metal gate electrode layer on the gate insulating layer.”

2422. The **'675 Accused Products** are made by a process that includes depositing, on the insulating layer, a first metal gate electrode.

2423. For example, PMOS, NMOS, and I/O transistors in the GK107 GPU include a titanium nitride layer on the hafnium oxide.

2424. The titanium nitride layer is an electrically conductive layer.

2425. As shown below, the **GK107 Report** depicts the titanium nitride layers on the hafnium oxide for PMOS, NMOS, and I/O transistors.




Figure 2.3.1.6: Logic NMOS transistor, lattice fringe image, TEM cross-section.




Figure 2.3.1.12: Logic PMOS transistor, lattice fringe image, TEM cross-section.




Figure 2.3.2.3: I/O transistor with HV gate oxide, TEM cross-section.

GK107 Report at 38, 45, and 53, Figures 2.3.1.6, 2.3.1.12, and 2.3.2.3. (red annotations added).

2426. As described in the preceding paragraphs, the '675 Accused Products meet the limitations of Claim 6(c) of the '675 Patent.

2427. “**Claim 6(d) of the ’675 Patent**” recites “forming a dummy gate electrode layer on the first metal gate electrode layer, said dummy gate electrode layer and said first metal gate electrode layer comprising different materials.”

2428. The **’675 Accused Products** are made by a process that includes forming a dummy gate on the first metal gate electrode layer. The dummy gate and first metal gate electrode layer of the **’675 Accused Products** comprise different materials.

2429. For example, transistors in the GK107 GPU are formed a process that includes deposition of a dummy polysilicon gate on the titanium nitride layer. *See* **GK107 Report** at x.

2430. The deposition of the polysilicon in the GK107 forms a sacrificial dummy gate. *See* **GK107 Report** at x.

2431. The dummy polysilicon gate and the titanium nitride layer comprise different materials.

2432. The **GK107 Report** describes the formation of dummy polysilicon gates in the GK107 GPU. *See* **GK107 Report** at 15.

2433. As described in the preceding paragraphs, the **’675 Accused Products** meet the limitations of **Claim 6(d) of the ’675 Patent**.

2434. “**Claim 6(e) of the ’675 Patent**” recites “patterning the dummy gate electrode layer and the first metal gate electrode layer in sequence to define a dummy gate electrode on the patterned first metal gate electrode layer.”

2435. The **’675 Accused Products** are made by a process that includes patterning the first metal gate electrode layer and the dummy gate in sequence, thereby defining a dummy gate on the first metal gate electrode layer.

2436. For example, the process of forming transistors in the GK107 GPU includes sequentially patterning the polysilicon dummy gate on the titanium nitride layer.

2437. The patterning of the titanium nitride layer and the polysilicon layer in sequence in the GK107 defines a dummy gate electrode.

2438. The **GK107 Report** describes the patterning process for the polysilicon layer in the GK107 GPU. *See GK107 Report* at x.

2439. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 6(e) of the '675 Patent**.

2440. “**Claim 6(f) of the '675 Patent**” recites “forming electrically insulating spacers on sidewalls of the dummy gate electrode and on sidewalls of the patterned first metal gate electrode layer.”

2441. The **'675 Accused Products** are made by a process that includes forming electrically insulating spacers on sidewalls of the patterned first metal gate electrode layer and the dummy gate electrode.

2442. For example, the GK107 GPU includes sidewalls spacers formed out of silicon nitride. *See GK107 Report* at x.

2443. The silicon nitride sidewall spacers are electrically insulating.

2444. The sidewalls spacers of the GK107 GPU are formed on the left and right sides of the titanium nitride and the dummy polysilicon gate for NMOS, PMOS, and I/O transistors. *See, e.g., GK107 Report* at 36, Figure 2.3.1.4.

2445. As shown below, the **GK107 Report** depicts sidewalls spacers for transistors in the GK107 GPU.




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.




Figure 2.3.1.9: Logic PMOS transistor gate stack, TEM cross-section.

GK107 Report at 36 and 42, Figures 2.3.1.4 and 2.3.1.9; *see also id.* at 53, Figure 2.3.2.3.

2446. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 6(f) of the '675 Patent**.

2447. “**Claim 6(g) of the '675 Patent**” recites “removing the dummy gate electrode from between the spacers by selectively etching back the dummy gate electrode using the spacers as an etching mask.”

2448. The **'675 Accused Products** are made by a process that includes removing the dummy gate by selective etching of the dummy gate using the spacers as an etching mask.

2449. For example, transistors in the GK107 GPU are formed by a process that includes the step of removing the polysilicon gate. *See GK107 Report* at 15, 39.

2450. Polysilicon gates of the GK107 GPU are removed using the sidewall spacers as an etching mask.

2451. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 6(g) of the '675 Patent**.

2452. “**Claim 6(h) of the '675 Patent**” recites “depositing a second metal gate electrode layer onto inner sidewalls of the spacers and onto an upper surface of the patterned first metal gate electrode layer.”

2453. The **'675 Accused Products** are made by a process that includes depositing a second metal gate electrode layer onto the upper surface of the first metal gate electrode and on inner sidewalls of the spacers.

2454. For example, the removal of the dummy gate of the GK107 GPU exposes the top surface of the titanium nitride layer and exposes inner sidewalls of the the sidewall spacers.

2455. The transistors of the GK107 GPU include a second metal layer that is deposited above the exposed top surface of the titanium nitride layer.

2456. A portion of the second metal layer is deposited on exposed inner walls of the sidewall spacers.

2457. The deposited second metal layer of the GK107 GPU includes a portion having a “U” shaped cross-section.

2458. As shown below, the **GK107 Report** depicts the second metal layer for PMOS transistors of the GK107 GPU.




Figure 2.3.1.9: Logic PMOS transistor gate stack, TEM cross-section.

GK107 Report at 42, Figure 2.3.1.9.

2459. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 6(h) of the '675 Patent**.

2460. “**Claim 6(i) of the '675 Patent**” recites “depositing a third metal gate electrode layer onto the second metal gate electrode layer to thereby fill a space between the inner sidewalls of the spacers, said second and third metal gate electrode layers comprising different materials.”

2461. The **'675 Accused Products** are made by a process that includes depositing a third metal layer onto the second metal layer to fill a space between the sidewall spacers with the second and third metal layers comprising different materials.

2462. For example, PMOS transistors in the GK107 GPU include a third metal layer formed by deposition above the second metal layer and into the space between the sidewall spacers.

2463. As shown below, the **GK107 Report** depicts the third metal layer for transistors of the GK107 GPU, with the second and third metal layers comprising different materials.




Figure 2.3.1.9: Logic PMOS transistor gate stack, TEM cross-section.

GK107 Report at 42, Figure 2.3.1.9.

2464. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 6(i) of the '675 Patent**.

2465. “**Claim 6(j) of the '675 Patent**” recites “planarizing the third metal gate electrode layer and the second metal gate electrode layer to thereby define a composite metal gate electrode of a PMOS transistor between the inner sidewalls of the spacers.”

2466. The **'675 Accused Products** are made by a process that includes planarizing the second and third gate electrode layers to define a composite metal gate electrode of a PMOS transistor between the inner sidewalls of the spacers.

2467. For example, PMOS transistors of the GK107 GPU are formed by a process including planarizing the second and third gate electrode layers, thereby defining a composite metal gate electrode.

2468. The gates of the transistors of the GK107 GPU all have the approximately same height as a result of the planarization of the second and third metal gate electrode layers.

2469. As shown below, the **GK107 Report** depicts the transistor formed after planarization of the second and third metal gate electrode layers for PMOS transistors.




Figure 2.3.1.7: Logic PMOS transistors, SEM cross-section.

GK107 Report at 40, Figure 2.3.1.7. (red annotations added); *see also* **GK107 Report** at 41, Figure 2.3.1.8 showing the results of planarization of the second and third metal electrode layers.

2470. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 6(j) of the '675 Patent**.

2471. "**Claim 6(k) of the '675 Patent**" recites "said composite metal gate electrode comprising a portion of the third metal gate electrode layer, a portion of the second metal gate

electrode layer having a U-shaped cross-section and the patterned first metal gate electrode layer.”

2472. The **'675 Accused Products** are made by a process in which composite gate electrodes comprise a portion of the third metal gate electrode layer, a portion of the second metal gate electrode layer having a U-shaped cross-section and the first metal gate electrode layer.

2473. For example, the composite metal gate electrode of PMOS transistors of the GK107 GPU include the patterned titanium nitride layer, a portion of the second metal gate electrode layer with a U-shaped cross-section above the titanium nitride layer, and a portion of the third metal gate electrode layer.

2474. As shown below, the **GK107 Report** depicts the composite PMOS transistors with the patterned titanium nitride layer, portion of the U-shaped second metal gate electrode layer, and portion of the third metal gate electrode layer.




Figure 2.3.1.9: Logic PMOS transistor gate stack, TEM cross-section.

GK107 Report at 40, Figure 2.3.1.9.

2475. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 6(k) of the '675 Patent**.

2476. "**Claim 7 of the '675 Patent**" recites "[t]he method of claim 6, wherein said planarizing comprises planarizing the third metal gate electrode layer and the second metal gate electrode layer in sequence to reveal the portion of the second metal gate electrode layer having a U-shaped cross-section."

2477. In the **'675 Accused Products**, the second and third metal gate electrode layers are planarized in sequence to reveal the portion of the second metal gate electrode layer having a U-shaped cross-section.

2478. For example, the planarization of the GK107 GPU comprises sequentially planarizing the third and second metal gate electrode layers.

2479. Following planarization, the second metal gate electrode layer having a U-shaped cross section is revealed.

2480. The **GK107 Report** depicts the revealed U-shaped cross section of second metal gate electrode layer following planarization. *See GK107 Report* at 40, Figure 2.3.1.9.

2481. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 7 of the '675 Patent**.

2482. "**Claim 8(a) of the '675 Patent**" recites "[t]he method of claim 6, wherein said removing the dummy gate electrode is preceded by covering the spacers and the dummy gate electrode with an electrically insulating mold layer."

2483. The **'675 Accused Products** are made by a process that includes depositing an electrically insulating mold layer covering the dummy gate and sidewall spacers.

2484. For example, the process of forming transistors in the GK107 GPU includes depositing insulating mold material covering the sidewall spacers. *See* **GK107 Report** at x.

2485. The insulating mold material in the GK107 GPU is additionally deposited such that it covers the dummy gate. *See* **GK107 Report** at x.

2486. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 8(a) of the '675 Patent**.

2487. "**Claim 8(b) of the '675 Patent**" recites "wherein said removing the dummy gate electrode is preceded by . . . removing an upper portion of the mold layer to expose an upper surface of the dummy gate electrode."

2488. The **'675 Accused Products** are made by a process that includes exposing the upper surface of the dummy gate by removing the insulating mold layer.

2489. For example, the process of forming transistors in the GK107 GPU includes removing a portion of the mold material to expose the upper surface of the dummy polysilicon gate so that the polysilicon gate may be removed. *See, e.g.,* **GK107 Report** at x.

2490. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 8(b) of the '675 Patent**.

2491. "**Claim 10(a) of the '675 Patent**" recites "[t]he method of claim 6, wherein said patterning the dummy gate electrode layer and the first metal gate electrode layer comprises patterning the dummy gate electrode layer and the first metal gate electrode layer in sequence to define a second dummy gate electrode on the patterned first metal gate electrode layer."

2492. The **'675 Accused Products** are made by a process that includes patterning the dummy gate electrode layer and first metal gate electrode layer in sequence, thus defining a second dummy gate electrode on the patterned first metal gate electrode layer.

2493. For example, the GK107 GPU includes NMOS transistors fabricated with a process that involves patterning the dummy polysilicon gate on the titanium nitride layer. *See* **GK107 Report** at x-xi.

2494. The patterned dummy polysilicon gates define second dummy gate electrodes for NMOS transistors.

2495. The **GK107 Report** describes the patterning process for the polysilicon dummy gate in the GK107 GPU. *See* **GK107 Report** at x.

2496. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 10(a) of the '675 Patent**.

2497. “**Claim 10(b) of the '675 Patent**” recites “wherein said forming electrically insulating spacers on sidewalls of the dummy gate electrode is followed by replacing the second dummy gate electrode with an upper metal gate electrode of an NMOS transistor.”

2498. The **'675 Accused Products** are made by a process that includes replacing the second dummy gate electrode with an upper metal gate electrode of an NMOS transistor.

2499. For example, NMOS transistors of the GK107 GPU are formed by etching the polysilicon and replacing it with an upper metal gate electrode, above the titanium nitride layer.

2500. As shown below, the **GK107 Report** depicts the upper metal gate electrode for an NMOS transistor above the titanium nitride layer.




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.

GK107 Report at 36, Figure 2.3.1.4.

2501. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 10(b) of the '675 Patent**.

2502. “**Claim 11(a) of the '675 Patent**” recites “[t]he method of claim 10, wherein the gate insulating layer comprises a dielectric material selected from a group consisting of hafnium oxide and tantalum oxide.”

2503. The **'675 Accused Products** are made by a process that includes the gate insulating layer comprising a dielectric material from the group consisting of hafnium oxide and tantalum oxide.

2504. For example, the GK107 GPU includes hafnium oxide for PMOS, NMOS, and I/O transistors.

2505. As shown above, the **GK107 Report** depicts the hafnium oxide for PMOS, NMOS, and I/O transistors. See **GK107 Report** at 38, 45, and 53, Figures 2.3.1.6, 2.3.1.12, and 2.3.2.3.

2506. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 11(a) of the '675 Patent**.

2507. “**Claim 11(b) of the '675 Patent**” recites “wherein a gate electrode of the NMOS transistor comprises the upper metal gate electrode and a portion of the patterned first metal gate electrode layer.”

2508. The **'675 Accused Products** are made by a process that includes a gate electrode of NMOS transistors comprising a portion of the patterned first metal gate electrode layer and the upper metal gate electrode.

2509. For example, the GK107 GPU includes NMOS transistors having a gate electrode including the upper metal gate electrode and a portion of the titanium nitride layer.

2510. The **GK107 Report** depicts NMOS transistors with a gate electrode that includes the upper metal gate electrode and a portion of the titanium nitride layer. *See GK107 Report* at 36, Figure 2.3.1.4.

2511. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 11(a) of the '675 Patent**.

2512. “**Claim 12 of the '675 Patent**” recites “[t]he method of claim 11, wherein the portion of third metal gate electrode layer of the PMOS transistor and the upper metal gate electrode of the NMOS transistor comprise different metals.”

2513. The **'675 Accused Products** are made by a process wherein the portion of third metal gate electrode layer of the PMOS transistor and the upper metal gate electrode of the NMOS transistor comprise different metals.

2514. For example, the GK107 GPU includes different metals in the NMOS upper metal gate electrode and the PMOS third metal gate electrode layer.

2515. As shown below, the **GK107 Report** depicts different metals in a side-by-side comparison of the third metal gate electrode layer of the PMOS transistor and the upper metal gate electrode of the NMOS transistor.




Figure 2.3.1.13: Comparison of Logic NMOS and PMOS transistors, TEM images.

GK107 Report at 47, Figure 2.3.1.13.

2516. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 12 of the '675 Patent**.

2517. “**Claim 13 of the '675 Patent**” recites “[t]he method of claim 12, wherein the portion of third metal gate electrode layer of the PMOS transistor comprises titanium nitride and the upper metal gate electrode of the NMOS transistor comprises aluminum.”

2518. The **'675 Accused Products** are made by a process wherein the PMOS transistor includes titanium nitride in the portion of the third metal gate electrode and the NMOS transistor includes aluminum in the upper metal gate electrode.

2519. For example, the GK107 GPU includes titanium nitride in the portion of the third metal gate electrode. *See, e.g., GK 107 Report* at 40, Figure 2.3.1.9.

2520. For example, the GK107 GPU includes aluminum in the upper metal gate electrode. *See, e.g.*, **GK 107 Report** at 36, Figure 2.3.1.4.

2521. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 13 of the '675 Patent**.

2522. "**Claim 14 of the '675 Patent**" recites "[t]he method of claim 12, wherein the gate electrode of the NMOS transistor comprises a composite of aluminum and titanium nitride."

2523. The **'675 Accused Products** are made by a process wherein the NMOS transistor gate electrode comprises a composite of aluminum and titanium nitride.

2524. For example, the GK107 GPU includes titanium rich aluminum titanium nitride in the metal gate electrode of the NMOS transistors. *See, e.g.*, **GK 107 Report** at 36, Figure 2.3.1.4.

2525. As shown below, the **GK107 Report** depicts the titanium rich aluminum titanium nitride of the metal gate electrode of the NMOS transistors.




Figure 2.3.1.4: Logic NMOS transistor gate stack, TEM cross-section.

GK107 Report at 36, Figure 2.3.1.4.

2526. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 14 of the '675 Patent**.

2527. “**Claim 15 of the '675 Patent**” recites “[t]he method of claim 6, wherein the gate insulating layer comprises a dielectric material selected from a group consisting of hafnium oxide and tantalum oxide.”

2528. The **'675 Accused Products** include PMOS insulated gate transistors having gate insulating layers containing hafnium oxide.

2529. For example, the GK107 GPU includes PMOS transistors that contain insulating layers comprising hafnium oxide.

2530. As shown below, the **GK107 Report** depicts PMOS transistors containing hafnium oxide.




Figure 2.3.1.12: Logic PMOS transistor, lattice fringe image, TEM cross-section.

GK107 Report at 45, Figure 2.3.1.12, depicted below. (red annotation added).

2531. As described in the preceding paragraphs, the **'675 Accused Products** meet the limitations of **Claim 15 of the '675 Patent**.

2532. NVIDIA and Velocity have directly infringed and continue to directly infringe the **'675 Patent** by importing, offering to sell, selling, or using the **'675 Accused Products**.

2533. As described above, the **'675 Accused Products** are made by a process that infringes the **'675 Patent**.

2534. NVIDIA uses the **'675 Accused Products**.

2535. For example, NVIDIA uses its own products for demonstration, testing, and development purposes.

2536. NVIDIA has used and continues to use its own products, including the **'675 Accused Products**, such as for demonstration, testing, and development purposes.

2537. For example and without limitation, NVIDIA has used and continues to use the GM107, Tesla K10, NVIDIA Shield Tablet, NVIDIA Shield Portable, GeForce GTX 760, and GeForce GTX 770.

2538. NVIDIA also has used and continues to use third-party products, including the **'675 Accused Products**, such as for demonstration, testing, and development purposes.

2539. When NVIDIA uses the **'675 Accused Products**, such as for demonstration, testing, or development purposes, such use directly infringes the **'675 Patent**.

2540. NVIDIA has in the past and continues to import, offer to sell, and sell its own products made by processes that infringe the **'675 Patent**, including the **'675 Accused Products**.

2541. For example, NVIDIA has in the past and continues to import, offer to sell, and sell the GM107, Tesla K10, NVIDIA Shield Tablet, NVIDIA Shield Portable, GeForce GTX 760, and GeForce GTX 770.

2542. NVIDIA sells the **'675 Accused Products** to end users and customers.

2543. For example, NVIDIA sells the NVIDIA Shield Tablet and NVIDIA Shield Portable via its online store at <http://store.nvidia.com>.

2544. NVIDIA sells the **'675 Accused Products** to intermediate customers.

2545. For example, NVIDIA sells its products to distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers. **NVIDIA 2014 Form 10-K** at 8.

2546. The products NVIDIA sells to intermediate customers, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, include one or more of the **'675 Accused Products**.

2547. For example, NVIDIA sells the GeForce GTX 760 and GeForce GTX 770 to Best Buy, which sells these cards to end users and customers.

2548. When NVIDIA imports into the United States or offers for sale or sells in the United States the **'675 Accused Products**, such activity directly infringes the **'675 Patent**.

2549. Velocity sells computers that incorporate an **Accused 28 nm GPU**.

2550. For example, Velocity sells a computer called the Raptor Signature Edition.

2551. The Raptor Signature Edition computer can be configured to include a variety of NVIDIA GPUs, including a variety of the **Accused 28 nm GPUs**.

2552. For example, the Raptor Signature Edition computer can be configured to include an NVIDIA GeForce GTX Titan Black.

2553. Velocity has used and continues to use its own products made by processes that infringe the **'675 Patent**, including the **'675 Accused Products**, such as for testing and development purposes

2554. For example and without limitation, Velocity has used and continues to use computers that incorporate an **Accused 28 nm GPU**.

2555. When Velocity uses the **'675 Accused Products**, such as for demonstration, testing, or development purposes, such use directly infringes the **'675 Patent**.

2556. Velocity has in the past and continues to import, offer for sale, and sell its own products made by processes that infringe the **'675 Patent**, including the **'675 Accused Products**.

2557. For example and without limitation, Velocity has in the past and continues to import, offer for sale, and sell computers that incorporate an **Accused 28 nm GPU**.

2558. When Velocity imports, offers for sale, or sells the **'675 Accused Products**, such activity directly infringes the **'675 Patent**.

2559. NVIDIA has had actual knowledge of the **'675 Patent** since at least as early as as of November 12, 2014, the date the original complaint in this action was served.

2560. Velocity has had actual knowledge of the **'675 Patent** at least as early as as of November 11, 2014, the date the original complaint in this action was served.

2561. NVIDIA and Velocity indirectly infringe the **'675 Patent** by inducing infringement by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(b), in this District and elsewhere in the United States. Direct infringement is the result of activities performed by others, such as distributors, original equipment manufacturers, contract equipment

manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, by, for example, importing, offering to sell, selling, or using the **'675 Accused Products**.

2562. NVIDIA and Velocity have induced and continue to induce infringement of the **'675 Patent** by intending that others infringe the **'675 Patent** by importing, offering to sell, selling, or using the **'675 Accused Products**. NVIDIA and Velocity designed the **'675 Accused Products** such that they would each infringe one or more claims of the **'675 Patent**.

2563. NVIDIA and Velocity provide the **'675 Accused Products** to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users.

2564. By providing **'675 Accused Products** to others, NVIDIA and Velocity intend for **'675 Accused Products** to be imported into the United States or offered for sale, sold, or used in the United States.

2565. NVIDIA and Velocity also provide others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, instructions, user guides, and technical specifications. When others follow such instructions, user guides, and/or other design documentation, they directly infringe one or more claims of the **'675 Patent**.

2566. NVIDIA and Velocity know that by providing such instructions, user guides, and/or other design documentation, others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard

manufacturers, add-in board manufacturers, resellers, customers, and end users, follow those instructions, user guides, and other design documentation, and directly infringe one or more claims of the **'675 Patent**. NVIDIA and Velocity thus know that their actions actively induce infringement.

2567. NVIDIA sells the **'675 Accused Products** directly to customers and end users.

2568. For example, NVIDIA sells the NVIDIA Shield Tablet and NVIDIA Shield Portable via its online store at <http://store.nvidia.com>.

2569. Through such sales, NVIDIA specifically intends that others, such as customers and end users, will infringe one or more claims of the **'675 Patent**, for example by use of the **'675 Accused Products**.

2570. NVIDIA provides the **'675 Accused Products**, and reference designs for the **'675 Accused Products**, to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, and resellers.

2571. Through such activity, NVIDIA specifically intends that others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, and resellers, customers, and end users, infringe the **'675 Patent** by importing into the United States and offering to sell, selling, or using in the United States the **'675 Accused Products**.

2572. Through its manufacture (either directly, or through contract manufacturing facilities) and/or sale of the infringing products, NVIDIA specifically intends that others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original

design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, will infringe one or more claims of the **'675 Patent**.

2573. NVIDIA specifically intends for others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'675 Patent** in the United States.

2574. For example, NVIDIA advertised its “NVIDIA Authorized Board Partner Program ensures an exceptional customer experience when purchasing graphics cards and motherboards manufactured by partners that make use of NVIDIA’s latest technologies.”

NVIDIA PartnerForce Info.

2575. NVIDIA claims that its “NVIDIA Authorized Board Partners offer the latest technologies from NVIDIA” and lists six Authorized Board Partners in the United States, six in Canada, and none in any other countries. *See* **NVIDIA PartnerForce Info**. NVIDIA also lists ten distributors in the United States (including “pre and post sales technical support”) and seven in Canada but none in any other country. *See* **NVIDIA PartnerForce Info**.

2576. As an additional example, as of October 31, 2014, NVIDIA also advertised the “NVIDIA PartnerForce Program” which is “a sales and marketing program for value-added resellers, system builders, e-tailers and retailers who sell NVIDIA based components or systems.”

NVIDIA PartnerForce Program.

2577. NVIDIA advertises in the United States and in this judicial district that “[b]y joining the PartnerForce Program, you may leverage the world-class brands and technology platforms from NVIDIA, sales and technical support from our experienced team and hundreds of turnkey sales and marketing tools.” **NVIDIA PartnerForce Program**.

2578. NVIDIA also provides marketing materials and templates, including retail display items, web banners, copy blocks, logos, product shots, email and web page templates, to members of the “NVIDIA PartnerForce Program” which includes value-added resellers, system builders, etailers and retailers inducing them to offer to sell and the **'675 Accused Products**. *See NVIDIA PartnerForce Program.*

2579. At the Ford Event and Conference Center, in Dearborn, Michigan, NVIDIA showcased its latest processor technologies, which power everything from the CAD software that designers use to style cars to the infotainment systems that drivers use to map their trips and listen to music, in an effort to encourage domestic car manufacturers to include infringing technology in vehicles manufactured and sold in the United States. *See Traveling The Road To Silicon Motown.*

2580. As described in the preceding paragraphs, NVIDIA specifically targets the United States market for **'675 Accused Products** actively induces others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'675 Patent**.

2581. Velocity sells computers that incorporate an **Accused 28 nm GPU**.

2582. For example, Velocity sells a computer called the Raptor Signature Edition.

2583. The Raptor Signature Edition computer can be configured to include a variety of NVIDIA GPUs, including a variety of the **Accused 28 nm GPUs**.

2584. For example, the Raptor Signature Edition computer can be configured to include an NVIDIA GeForce GTX Titan Black.

2585. Velocity provides the **'675 Accused Products** to others, such as resellers, customers, and end users.

2586. Through such activity, NVIDIA specifically intends that others, such as resellers, customers, and end users, infringe the **'675 Patent** by importing into the United States and offering to sell, selling, or using in the United States the **'675 Accused Products**.

2587. Velocity specifically targets the United States market for its products listed above and actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'675 Patent** in the United States.

2588. For example, Velocity sells products via its website, <http://www.velocitymicro.com>, to customers in the United States. Such products include the **'675 Accused Products**, such as, for example, computers that are configured to include an NVIDIA GeForce GTX Titan Black.

2589. The customers infringe the **'675 Patent** by using the **'675 Accused Products** purchased from Velocity.

2590. Velocity also sells computers to resellers. Such computers include the **'675 Accused Products**, such as, for example, computers that are configured to include an **Accused 28 nm GPU**.

2591. For example, Velocity products have been sold at nearly every major electronics retailer in North America including Best Buy, Circuit City, RadioShack, Costco, Sears, Target, and many others. *See Velocity Company History*.

2592. Velocity introduces products and services that infringe the Asserted Patents intending that they would be used in this Judicial District and elsewhere in the United States.

2593. The resellers infringe the **'675 Patent** by importing, offering selling, selling, or using the **'675 Accused Products**.

2594. The resellers also sell the **'675 Accused Products** to customers and end users, who infringe the **'675 Patent** by using those products.

2595. As described in the preceding paragraphs, Velocity specifically targets the United States market for **'675 Accused Products** actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'675 Patent**.

2596. Samsung has no adequate remedy at law for NVIDIA's and Velocity's infringement of the **'675 Patent** and is suffering irreparable harm, requiring permanent injunctive relief.

SIXTH CLAIM FOR RELIEF
INFRINGEMENT OF U.S. PATENT NO. 6,804,724
(AGAINST NVIDIA AND VELOCITY)

2597. Each of the above listed paragraphs is incorporated herein by reference, and adopted, as if fully set forth again.

2598. The **'724 Patent** was filed on April 29, 1999, issued on October 12, 2004, and is entitled "Analog/Digital Display Adapter and a Computer System Having the Same." The **'724 Patent** is generally directed to computer display adapters.

2599. The **'724 Patent** was assigned to SEC, and SEC and continues to hold all rights, title, and interest in the **'724 Patent**. A true and correct copy of the **'724 Patent** is attached hereto as Exhibit F.

2600. All laptop computers with a DisplayPort port connected via a DisplayPort cable to a DisplayPort monitor with power infringe the **'724 Patent**.

2601. All **'724 Accused Products**, when connected via DisplayPort using a laptop computer's external DisplayPort or Thunderbolt port, infringe the **'724 Patent** in substantially the same way.

2602. For example, all **'724 Accused Velocity Laptops** practice elements of the **'724 Patent** in substantially the same way.

2603. All **'724 Accused Velocity Laptops** include similar flat panel LCD panels that fold onto the main body of the laptop.

2604. All **'724 Accused Velocity Laptops** include similar video controllers that provide digital video data to the LCD panel.

2605. All **'724 Accused Velocity Laptops** support DisplayPort and include an external DisplayPort port.

2606. DisplayPort support requires a DisplayPort transmitter, which is a digital transmitter that transmits parallel digital video signals.

2607. All **'724 Accused Velocity Laptops** include a DisplayPort transmitter.

2608. All **'724 Accused Velocity Laptops** receive and detect hot plug detect signals from external monitors.

2609. For further example, all **'724 Accused Mobile GPUs** practice elements of the **'724 Patent** in substantially the same way.

2610. All **'724 Accused Mobile GPUs** include video controllers that provide support for display technologies.

2611. All **'724 Accused Mobile GPUs** include a DisplayPort transmitter.

2612. All **'724 Accused Mobile GPUs** detect hot plug detect signals in similar ways.

2613. For example, all **'724 Accused Mobile GPUs** use the Device Control Block 4.0 Specification. *See, e.g., DCB 4.0*: “DCB 4.0 is used with Core6, Core7, and Core8 VBIOS (G80+).”

2614. According to the **DCB 4.0**, the hot plug detect signal is linked to a GPIO pin on all **'724 Accused Mobile GPUs**.

2615. “**Claim 6(a) of the '724 Patent**” recites “[a] portable computer system.”

2616. Laptop computers are portable computer systems.

2617. All **'724 Accused Velocity Laptops** are portable computer systems.

2618. As described in the preceding paragraphs, all **'724 Accused Velocity Laptops** include a portable computer system.

2619. All **'724 Accused Mobile GPUs** are intended to be included in laptop computers.

2620. As described in the preceding paragraphs, all laptop computers containing **'724 Accused Mobile GPUs** are portable computer systems.

2621. “**Claim 6(b) of the '724 Patent**” recites “a LCD panel hingedly attached to a main body of the portable computer.”

2622. All **'724 Accused Velocity Laptops** include an LCD panel that is hingedly attaches to the portable computer’s main body.

2623. For example, the NoteMagix M15 includes an LCD panel that attaches to and can fold toward the laptop’s main body.

2624. As described in the preceding paragraphs, **Claim 6(b) of the '724 Patent** is met by all **'724 Accused Velocity Laptops**.

2625. All **'724 Accused Mobile GPUs** are intended to be included in laptop computers.

2626. As described in the preceding paragraphs, **Claim 6(b) of the '724 Patent** is met by all laptop computers containing **'724 Accused Mobile GPUs**.

2627. "**Claim 6(c) of the '724 Patent**" recites "a video controller for providing digital video data to the LCD panel."

2628. All **'724 Accused Velocity Laptops** and all **'724 Accused Mobile GPUs** include a video controller that can provide digital video data to the LCD panel.

2629. For example, all **'724 Accused Velocity Laptops** include a **'724 Accused Mobile GPU**, which include a video controller.

2630. For example, the NoteMagix M15 Ultra includes a GeForce GTS 250M.

2631. The GeForce GTS 250M includes a video controller.

2632. When used as intended in a laptop computer, all **'724 Accused Mobile GPUs'** video controllers provide digital video data to the laptop's LCD panel.

2633. For example, the GeForce GTS 250M supports LVDS. **GeForce GTS 250M Specifications.**

2634. LVDS is a technology that can communicate digital video data to a laptop's LCD panel.

2635. In the NoteMagix M15 Ultra, the video controller included in the GeForce GTS 250M provides digital video data to the M15's LCD panel.

2636. The digital video data discussed in the preceding paragraph is provided to the M15's LCD panel via LVDS.

2637. As described in the preceding paragraphs, **Claim 6(c) of the '724 Patent** is met by all **'724 Accused Velocity Laptops**.

2638. As described in the preceding paragraphs, **Claim 6(c) of the '724 Patent** is met by all **'724 Accused Mobile GPUs**.

2639. "**Claim 6(d) of the '724 Patent**" recites "a first external video port connecting a digital cable to an external digital monitor."

2640. All **'724 Accused Velocity Laptops** include an external video port that can connect a digital cable to an external digital monitor.

2641. For example, the NoteMagix M15 Ultra includes a DisplayPort port configured to connect a DisplayPort cable to an external digital monitor.

2642. The NoteMagix M15 Ultra's DisplayPort port is an external video port.

2643. A DisplayPort cable is a digital cable.

2644. When used as intended, the **'724 Accused Velocity Laptops'** video port connects a digital cable to an external digital monitor.

2645. For example, when used as intended, the NoteMagix M15 Ultra's DisplayPort port connects a DisplayPort cable to an external digital monitor.

2646. All **'724 Accused Mobile GPUs** support DisplayPort and are intended to be included in laptop computers, such as the NoteMagix M15 Ultra, that support DisplayPort.

2647. As described in the preceding paragraphs, **Claim 6(d) of the '724 Patent** is met by all **'724 Accused Velocity Laptops**.

2648. As described in the preceding paragraphs, **Claim 6(d) of the '724 Patent** is met by laptop computers containing **'724 Accused Mobile GPUs**.

2649. "**Claim 6(e) of the '724 Patent**" recites "a digital transmitter sending parallel digital video data to said external digital monitor."

2650. All **'724 Accused Velocity Laptops** and all **'724 Accused Mobile GPUs** powered and connected to a powered digital monitor include a digital transmitter sending digital video data to the external monitor.

2651. For example, the NoteMagix M15 Ultra includes a GeForce GTS 250M, which supports DisplayPort. The GeForce GTS 250M therefore includes a DisplayPort digital transmitter for sending parallel digital video data to an external digital monitor. *See, e.g., DisplayPort Standard* at 29: “The capabilities of the DisplayPort transmitter and receiver, and the quality of the channel (or a cable) will determine whether the link rate is set to 2.7 Gbps or 1.62 Gbps,” *DisplayPort Standard* at 34: “DisplayPort supports three options for the number of Main Link lanes and two options for Main Link data rate per lane as follows: • 4, 2, or 1 lanes • 2.7 Gbps or 1.62 Gbps per lane[.]”

2652. Continuing the example of the previous two paragraphs, the digital transmitter in the GeForce GTS 250M of the NoteMagix M15 Ultra sends parallel digital video data to the external monitor when, for example, all devices have power and are connected.

2653. As described in the preceding paragraphs, **Claim 6(e) of the '724 Patent** is met by all **'724 Accused Velocity Laptops**.

2654. As described in the preceding paragraphs, **Claim 6(e) of the '724 Patent** is met by laptop computers containing **'724 Accused Mobile GPUs**.

2655. “**Claim 6(f) of the '724 Patent**” recites “said digital monitor comprises a means for generating a cable sensing signal to be sent to said first external video port over the digital cable, thereby informing the video controller of the digital cable connection state of said first external port.”

2656. All **'724 Accused Velocity Laptops** and all **'724 Accused Mobile GPUs** connected via a DisplayPort cable to a DisplayPort-capable monitor include a digital monitor that generates a cable sensing signal to be sent to the external video port over the digital cable and to inform the video controller of the digital cable connection state of the external video port.

2657. For example, when an **Accused Velocity Micro Laptop** connects via a DisplayPort cable to a DisplayPort-capable monitor, the DisplayPort-capable monitor generates a hot plug detect signal and sends it to the DisplayPort port of the **Accused Velocity Micro Laptop**.

2658. In all **'724 Accused Velocity Laptops**, the included **Accused Mobile GPU** connects to the DisplayPort port.

2659. For example, when the NoteMagix M15 Ultra connects via a DisplayPort cable to a DisplayPort-capable monitor, the DisplayPort-capable monitor generates a hot plug detect signal and sends it to the DisplayPort port of the NoteMagix M15 Ultra, where it is received by the GeForce GTS 250M.

2660. The DisplayPort hot plug detect signal indicates to a video controller that a DisplayPort monitor is connected and has power. *See, e.g., DisplayPort Standard* at 117: “Whenever the Hot Plug Detect signal is active (the connectors are plugged in and the Sink Device has at least a “trickle” AC power), AUX CH services must be available.”

2661. For example, in the example described in Paragraph 2659, the hot plug detect signal informs the GeForce GTS 250M’s video controller of the digital cable connection state of the DisplayPort port.

2662. For example, when the NoteMagix M15 Ultra is powered on and connected via a DisplayPort cable to a DisplayPort-capable monitor, a sensor on the GeForce GTS 250M detects

the hot plug detect signal. The hot plug detect signal is linked to a GPIO pin on the GeForce GTS 250M. *See, e.g., DCB 4.0* at Pin Function ID List (red annotation added):

Function (15:8)

This lists the function of each GPIO pin. Here's a list of the function numbers and a short description of each:

- 0 = LCD0 backlight: Backlight control. LCD0 corresponds to the LCD0 defined in the LCD ID field in the Connector Table.
- 1 = LCD0 power: Panel Power control. LCD0 corresponds to the LCD0 defined in the LCD ID field in the Connector Table.
- 2 = LCD0 Power Status: Panel Power status. LCD0 corresponds to the LCD0 defined in the LCD ID field in the Connector Table.
- 3 = VSYNC: Alternate VSync signal using GPIO pin.
- 4 = VSEL0: Voltage Select Bit 0
- 5 = VSEL1: Voltage Select Bit 1
- 6 = VSEL2: Voltage Select Bit 2
- 7 = Hotplug A: 1st Hotplug signal
- 8 = Hotplug B: 2nd Hotplug signal
- . . .
- 80 = DVI to DAC connector switch. This GPIO allows for DAC 0 (TV) to be selected to route to the DVI Connector when the GPIO is set to the logical OFF state. When the GPIO is set to logical ON state, DAC 1 (CRT) will be routed to the DVI connector.
- 81 = Hotplug C: 3rd Hotplug signal
- 82 = Hotplug D: 4th Hotplug signal
- 83 = DisplayPort to DVI dongle present C, when this GPIO asserts, we need to configure DisplayPort encoder to output TMDS signal.
- 84 = DisplayPort to DVI dongle present D, when this GPIO asserts, we need to configure DisplayPort encoder to output TMDS signal.
- 85 = Maxim Max6305 or compatible external reset controller. Enabled is Active Low so init value should be Active High [No inversions]
- 86 = Active display LED to indicate the GPU with active display in SLI mode.
- 87 = SPDIF input.
- 88 = TOSLINK input.
- 89 = SPDIF/TOSLINK Select. When GPIO is set LOW, SPDIF is selected. When GPIO is set HI, TOSLINK is selected.
- 90 = DPAUX/I2C select A. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 91 = DPAUX/I2C select B. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 92 = DPAUX/I2C select C. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 93 = DPAUX/I2C select D. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 94 = Hotplug E: 5th Hotplug signal
- 95 = Hotplug F: 6th Hotplug signal
- 96 = Hotplug G: 7th Hotplug signal
- 99 = GPIO External Device 1 Interrupt - Used to surface an interrupt from a GPIO external device

2663. As described in the preceding paragraphs, **Claim 6(f) of the '724 Patent** is met by all **'724 Accused Velocity Laptops** and all **'724 Accused Mobile GPUs**.

2664. “**Claim 6(g) of the '724 Patent**” recites “said system further comprising a monitor power sensor detecting a presence of power applied to the external digital monitor, whereby a display enable signal is generated in the video controller and is sent to the transmitter to enable the digital video signals to be sent to the external digital monitor when the presence of power applied to the external digital monitor is detected.”

2665. All **'724 Accused Velocity Laptops** and all **'724 Accused Mobile GPUs** powered on and connected via a DisplayPort cable to a powered DisplayPort-capable monitor include a monitor power sensor that detects whether the external digital monitor has power and accordingly enables video signal transmission via a display enable signal generated in the video controller.

2666. For example, when the NoteMagix M15 Ultra is powered on and connected via a DisplayPort cable to a DisplayPort-capable monitor, a sensor on the GeForce GTS 250M detects the hot plug detect signal.

2667. The hot plug detect signal is linked to a GPIO pin on the GeForce GTS 250M.
See, e.g., DCB 4.0 at Pin Function ID List (red annotation added):

Function (15:8)

This lists the function of each GPIO pin. Here's a list of the function numbers and a short description of each:

- 0 = LCD0 backlight: Backlight control. LCD0 corresponds to the LCD0 defined in the LCD ID field in the Connector Table.
- 1 = LCD0 power: Panel Power control. LCD0 corresponds to the LCD0 defined in the LCD ID field in the Connector Table.
- 2 = LCD0 Power Status: Panel Power status. LCD0 corresponds to the LCD0 defined in the LCD ID field in the Connector Table.
- 3 = VSYNC: Alternate VSync signal using GPIO pin.
- 4 = VSEL0: Voltage Select Bit 0
- 5 = VSEL1: Voltage Select Bit 1
- 6 = VSEL2: Voltage Select Bit 2
- 7 = Hotplug A: 1st Hotplug signal
- 8 = Hotplug B: 2nd Hotplug signal
- ...
- 80 = DVI to DAC connector switch. This GPIO allows for DAC 0 (TV) to be selected to route to the DVI Connector when the GPIO is set to the logical OFF state. When the GPIO is set to logical ON state, DAC 1 (CRT) will be routed to the DVI connector.
- 81 = Hotplug C: 3rd Hotplug signal
- 82 = Hotplug D: 4th Hotplug signal
- 83 = DisplayPort to DVI dongle present C, when this GPIO asserts, we need to configure DisplayPort encoder to output TMDS signal.
- 84 = DisplayPort to DVI dongle present D, when this GPIO asserts, we need to configure DisplayPort encoder to output TMDS signal.
- 85 = Maxim Max6305 or compatible external reset controller. Enabled is Active Low so init value should be Active High [No inversions]
- 86 = Active display LED to indicate the GPU with active display in SLI mode.
- 87 = SPDIF input.
- 88 = TOSLINK input.
- 89 = SPDIF/TOSLINK Select. When GPIO is set LOW, SPDIF is selected. When GPIO is set HI, TOSLINK is selected.
- 90 = DPAUX/I2C select A. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 91 = DPAUX/I2C select B. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 92 = DPAUX/I2C select C. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 93 = DPAUX/I2C select D. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 94 = Hotplug E: 5th Hotplug signal
- 95 = Hotplug F: 6th Hotplug signal
- 96 = Hotplug G: 7th Hotplug signal
- 99 = GPIO External Device 1 Interrupt - Used to surface an interrupt from a GPIO external device

2668. The hot plug detect signal discussed in the preceding paragraph indicates the presence of power applied to the external digital monitor. *See, e.g., DisplayPort Standard* at 117: “Whenever the Hot Plug Detect signal is active (the connectors are plugged in and the Sink Device has at least a “trickle” AC power), AUX CH services must be available.”

2669. Continuing the example in paragraph 2666, the GeForce GTS 250M’s video controller generates a display enable signal to enable video signal transmission from the DisplayPort transmitter. This display enable signal is generated in response to an active hot plug detect signal. *See, e.g., DisplayPort Standard* at 118: “Upon detecting an IRQ Hot Plug Detect signal toggle, the Source Link Policy Maker must read the Receiver Capability field in the DPCD of the Sink Device and configure the link accordingly, using Link Training procedure as

described in Section 2.5.3.3 and Section 3.5.1.3”; **DisplayPort Standard** at 138: “After Link Training is successfully completed (which means that DisplayPort receiver is synchronized to the incoming Main Link data and before the transport of a main video stream starts, the Source Main Link transmitter must be sending the ‘idle pattern’”)

2670. As described in the preceding paragraphs, **Claim 6(g) of the ’724 Patent** is met by all **’724 Accused Velocity Laptops** and all **’724 Accused Mobile GPUs**.

2671. “**Claim 7 of the ’724 Patent**” recites, “The portable computer system of claim 6, further comprising a power supply control circuit switching on/off of supply voltages of the transmitter based on said cable sensing signal fed from the external digital monitor such that the supply voltages are supplied to the transmitter only if said digital cable is connected to said first external video port.”

2672. All **’724 Accused Velocity Laptops** and all **’724 Accused Mobile GPUs** further include a power supply control circuit that switches power based on the cable sensing signal from the external digital monitor.

2673. For example, the NoteMagix M15 Ultra includes a GeForce GTS 250M, which in turn includes a power supply control circuit that enables or disables power to the DisplayPort transmitter based on whether the hot plug detect signal is active.

2674. As described in the preceding paragraphs, **Claim 7 of the ’724 Patent** is met by all **’724 Accused Velocity Laptops** and all laptops containing **’724 Accused Mobile GPUs**.

2675. “**Claim 8 of the ’724 Patent**” recites, “The portable computer system of 7, further comprising a power supply control circuit switching on/off of supply voltages of the transmitter based on a signal fed from the external digital monitor causing power to be applied to the transmitter only power is applied to said external digital monitor.”

2676. All **'724 Accused Velocity Laptops** and all **'724 Accused Mobile GPUs** further include a power supply control circuit that switches power based on the cable sensing signal from the external digital monitor.

2677. For example, the NoteMagix M15 Ultra includes a GeForce GTS 250M, which in turn includes a power supply control circuit that enables or disables power to the DisplayPort transmitter based on whether the hot plug detect signal is active.

2678. As described in the preceding paragraphs, **Claim 8 of the '724 Patent** is met by all **'724 Accused Velocity Laptops** and all laptops containing **'724 Accused Mobile GPUs**.

2679. "**Claim 9(a) of the '724 Patent**" recites "[a] portable computer system."

2680. Laptop computers are a type of portable computer system.

2681. All **'724 Accused Mobile GPUs** are intended to be included in laptop computers.

2682. For example, the ASUS G750JX, which includes a GeForce GTX 770M, is a laptop computer.

2683. As described in the preceding paragraphs, all laptop computers containing **'724 Accused Mobile GPUs** are portable computer systems.

2684. "**Claim 9(b) of the '724 Patent**" recites "a LCD panel hingedly attached to a main body of the portable computer."

2685. All laptop computers containing **'724 Accused Mobile GPUs with Analog Output** include an LCD that hingedly attaches to the portable computer's main body.

2686. For example, the ASUS G750JX, which contains a GeForce GTX 770M, includes a flat panel LCD that attaches to and can fold toward the laptop's main body.

2687. All **'724 Accused Mobile GPUs** are intended to be included in laptop computers.

2688. All laptop computers include an LCD that hingedly attaches to the portable computer's main body.

2689. As described in the preceding paragraphs, **Claim 9(b) of the '724 Patent** is met by laptop computers containing **'724 Accused Mobile GPUs**.

2690. "**Claim 9(c) of the '724 Patent**" recites "a video controller for providing digital video data to the LCD panel."

2691. All **'724 Accused Mobile GPUs** include a video controller that can provide digital video data to the LCD panel.

2692. For example, the GeForce GTX 770M includes a video controller.

2693. When used as intended in a laptop computer, all **'724 Accused Mobile GPUs'** video controllers provide digital video data to the laptop's LCD panel.

2694. For example, the GeForce GTX 770M supports LVDS. **GeForce GTX 770M Specifications.**

2695. LVDS is a technology that can communicate digital video data to a laptop's LCD panel.

2696. For example, in the ASUS G750JX, the video controller included in the GeForce GTX 770M provides digital video data to the LCD panel.

2697. The digital video data discussed in the preceding paragraph is provided to the LCD panel via LVDS.

2698. As described in the preceding paragraphs, **Claim 9(c) of the '724 Patent** is met by all **'724 Accused Mobile GPUs**.

2699. "**Claim 9(d) of the '724 Patent**" recites "a first external video port connecting a digital cable to an external digital monitor."

2700. All **'724 Accused Velocity Laptops** include an external video port that can connect a digital cable to an external digital monitor.

2701. For example, the ASUS G750JX includes a Thunderbolt port configured to connect a DisplayPort cable to an external digital monitor.

2702. Thunderbolt is interoperable with DisplayPort 1.1a, and a mini-DisplayPort-to-DisplayPort cable can connect a DisplayPort-capable monitor directly to a Thunderbolt port.

2703. The ASUS G750JX's Thunderbolt port is an external video port.

2704. A mini-DisplayPort-to-DisplayPort cable is a digital cable.

2705. When used as intended, the video ports of laptop computers containing **'724 Accused Mobile GPUs with Analog Output** connect a digital cable to an external digital monitor.

2706. For example, when used as intended, the ASUS G750JX's Thunderbolt port connects via DisplayPort to an external digital monitor.

2707. All **'724 Accused Mobile GPUs** support DisplayPort and are intended to be included in laptop computers, such as the ASUS G750JX, that support DisplayPort.

2708. As described in the preceding paragraphs, **Claim 9(d) of the '724 Patent** is met by laptop computers containing **'724 Accused Mobile GPUs with Analog Output**.

2709. "**Claim 9(e) of the '724 Patent**" recites "a digital transmitter sending parallel digital video data to said external digital monitor."

2710. All **'724 Accused Mobile GPUs with Analog Output** powered and connected to a powered digital monitor include a digital transmitter sending digital video data to the external monitor.

2711. For example, the ASUS G750JX includes a GeForce GTX 770M, which supports DisplayPort. The GeForce GTX 770M therefore includes a DisplayPort digital transmitter for sending parallel digital video data to an external digital monitor. *See, e.g., DisplayPort Standard* at 29: “The capabilities of the DisplayPort transmitter and receiver, and the quality of the channel (or a cable) will determine whether the link rate is set to 2.7 Gbps or 1.62 Gbps,” *DisplayPort Standard* at 34: “DisplayPort supports three options for the number of Main Link lanes and two options for Main Link data rate per lane as follows: • 4, 2, or 1 lanes • 2.7 Gbps or 1.62 Gbps per lane[.]”

2712. Continuing the example of the previous two paragraphs, the digital transmitter in the GeForce GTX 770M of the ASUS G750JX sends parallel digital video data to the external monitor when, for example, all devices have power and are connected.

2713. As described in the preceding paragraphs, **Claim 9(e) of the '724 Patent** is met by laptop computers containing **'724 Accused Mobile GPUs with Analog Output**.

2714. “**Claim 9(f) of the '724 Patent**” recites “the video controller further generates analog video signals to be sent to a second external video port and then to an external analog monitor, said external analog monitor being connected to said second external video port by an analog cable.”

2715. All **'724 Accused Mobile GPUs with Analog Output** connected via a DVI-I, VGA, or other analog cable to an external monitor capable of receiving analog signals include a video controller that generates analog signals to send via a cable to an analog monitor.

2716. For example, when the ASUS G750JX VGA port connects with a VGA cable to an external analog monitor, the video controller on the GeForce GTX 770M generates analog video signals to be sent to the external analog monitor.

2717. As described in the preceding paragraphs, **Claim 9(f) of the '724 Patent** is met by laptop computers containing **'724 Accused Mobile GPUs with Analog Output**.

2718. “**Claim 9(g) of the '724 Patent**” recites “said system further comprising a monitor power sensor detecting a presence of power applied to the external digital monitor, whereby a display enable signal is generated in the video controller and is sent to the transmitter to enable the digital video signals to be sent to the external digital monitor when the presence of power applied to the external digital monitor is detected.”

2719. All **'724 Accused Mobile GPUs with Analog Output** powered on and connected via a DisplayPort cable to a powered DisplayPort-capable monitor include a monitor power sensor that detects whether the external digital monitor has power and accordingly enables video signal transmission via a display enable signal generated in the video controller.

2720. For example, when the ASUS G750JX is powered on and connected via DisplayPort to a DisplayPort-capable monitor, a sensor on the GeForce GTX 770M detects the hot plug detect signal.

2721. The hot plug detect signal is linked to a GPIO pin on the GeForce GTX 770M. *See, e.g., DCB 4.0* at Pin Function ID List (red annotation added):

Function (15:8)

This lists the function of each GPIO pin. Here's a list of the function numbers and a short description of each:

- 0 = LCD0 backlight: Backlight control. LCD0 corresponds to the LCD0 defined in the LCD ID field in the Connector Table.
- 1 = LCD0 power: Panel Power control. LCD0 corresponds to the LCD0 defined in the LCD ID field in the Connector Table.
- 2 = LCD0 Power Status: Panel Power status. LCD0 corresponds to the LCD0 defined in the LCD ID field in the Connector Table.
- 3 = VSYNC: Alternate VSync signal using GPIO pin.
- 4 = VSEL0: Voltage Select Bit 0
- 5 = VSEL1: Voltage Select Bit 1
- 6 = VSEL2: Voltage Select Bit 2
- 7 = Hotplug A: 1st Hotplug signal
- 8 = Hotplug B: 2nd Hotplug signal
- ...
- 80 = DVI to DAC connector switch. This GPIO allows for DAC 0 (TV) to be selected to route to the DVI Connector when the GPIO is set to the logical OFF state. When the GPIO is set to logical ON state, DAC 1 (CRT) will be routed to the DVI connector.
- 81 = Hotplug C: 3rd Hotplug signal
- 82 = Hotplug D: 4th Hotplug signal
- 83 = DisplayPort to DVI dongle present C, when this GPIO asserts, we need to configure DisplayPort encoder to output TMDS signal.
- 84 = DisplayPort to DVI dongle present D, when this GPIO asserts, we need to configure DisplayPort encoder to output TMDS signal.
- 85 = Maxim Max6305 or compatible external reset controller. Enabled is Active Low so init value should be Active High [No inversions]
- 86 = Active display LED to indicate the GPU with active display in SLI mode.
- 87 = SPDIF input.
- 88 = TOSLINK input.
- 89 = SPDIF/TOSLINK Select. When GPIO is set LOW, SPDIF is selected. When GPIO is set HI, TOSLINK is selected.
- 90 = DPAUX/I2C select A. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 91 = DPAUX/I2C select B. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 92 = DPAUX/I2C select C. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 93 = DPAUX/I2C select D. When this GPIO is set to Logical ON state, DPAUX will be selected. Logical OFF state selects I2C.
- 94 = Hotplug E: 5th Hotplug signal
- 95 = Hotplug F: 6th Hotplug signal
- 96 = Hotplug G: 7th Hotplug signal
- 99 = GPIO External Device 1 Interrupt - Used to surface an interrupt from a GPIO external device

2722. The hot plug detect signal discussed in the preceding paragraph indicates the presence of power applied to the external digital monitor. *See, e.g., DisplayPort Standard* at 117: “Whenever the Hot Plug Detect signal is active (the connectors are plugged in and the Sink Device has at least a “trickle” AC power), AUX CH services must be available.”

2723. Continuing the example in paragraph 2720, the GeForce GTX 770M’s video controller generates a display enable signal to enable video signal transmission from the DisplayPort transmitter. This display enable signal is generated in response to an active hot plug detect signal. *See, e.g., DisplayPort Standard* at 118: “Upon detecting an IRQ Hot Plug Detect signal toggle, the Source Link Policy Maker must read the Receiver Capability field in the DPCD of the Sink Device and configure the link accordingly, using Link Training procedure as

described in Section 2.5.3.3 and Section 3.5.1.3”; **DisplayPort Standard** at 138: “After Link Training is successfully completed (which means that DisplayPort receiver is synchronized to the incoming Main Link data and before the transport of a main video stream starts, the Source Main Link transmitter must be sending the ‘idle pattern’ . . .”).

2724. As described in the preceding paragraph, **Claim 9(g) of the ’724 Patent** is met by all **’724 Accused Mobile GPUs with Analog Output** powered on and connected via a DisplayPort cable to a powered DisplayPort-capable monitor.

2725. “**Claim 10 of the ’724 Patent**” recites, “The portable computer system of claim 9, further comprising a monitor cable sensor detecting a presence/absence of a connection of said monitor cables to said first and said second external video ports, respectively, whereby a display enable signals are generated in the video controller and sent to corresponding transmitters when a connection between a corresponding monitor cable and a corresponding external video port is detected.”

2726. All laptops containing **’724 Accused Mobile GPUs with Analog Output** powered on and connected via a DisplayPort cable to a powered DisplayPort-capable monitor and connected via a DVI-I, VGA, or other analog cable to an external monitor include a sensor detecting the monitor cable connections in the first and second video ports, with the video controller generating a display enable signal to send to corresponding transmitters when a connection is detected.

2727. For example, the ASUS G750JX includes a GeForce GTX 770M. The GeForce GTX 770M detects the load on analog output pins. *See DCB 4.0:*

- 12 = DAC 1 Select: DAC 1 mux select that allows us to switch between using the CRT (Off state) or TV (On State) filters on the board.
- 13 = DAC 1 Alternate Load Detect: When the DAC 1 is not currently switched to a device that needs detection, this GPIO pin can be used to detect the alternate load on the green channel.

2728. The GeForce GTX 770M also detects the hotplug signal associated with DisplayPort.

2729. Load and hotplug signal detection is associated with the presence or absence of monitor cable connections.

2730. The GeForce GTX 770M generates display enable signals when it detects a connection between a monitor cable and the corresponding video port.

2731. As described in the preceding paragraph, **Claim 10 of the '724 Patent** is met by laptops containing **'724 Accused Mobile GPUs with Analog Output** powered on and connected via a DisplayPort cable to a powered DisplayPort-capable monitor and connected via a DVI-I, VGA, or other analog cable to an external monitor.

2732. NVIDIA and Velocity have directly infringed and continue to directly infringe the **'724 Patent** by making, using, offering to sell, selling, or importing infringing systems using the **'724 Accused Products**.

2733. NVIDIA assembles its own products into infringing systems, such as a laptop computer powered on and connected via DisplayPort to a powered external monitor, such as for demonstration, testing, and development purposes.

2734. For example, NVIDIA assembles its own products into infringing systems for demonstration, testing, and development purposes.

2735. NVIDIA has assembled and continues to assemble its own products, including the **'724 Accused Mobile GPUs**, into infringing systems, such as a laptop computer powered on and connected via DisplayPort to a powered external monitor, such as for demonstration, testing, and development purposes.

2736. For example and without limitation, NVIDIA has assembled and continues to assemble the GeForce GTX 770M into infringing systems.

2737. NVIDIA also has used and continues to use third-party products that include the **'724 Accused Mobile GPUs** in infringing systems, such as for demonstration, testing, and development purposes.

2738. When NVIDIA assembles the **'724 Accused Products** into infringing systems, such as for demonstration, testing, or development purposes, NVIDIA directly infringes the **'724 Patent**.

2739. NVIDIA uses these infringing systems, such as for demonstration, testing, and development purposes.

2740. When NVIDIA uses the infringing systems described in the previous paragraphs, such as for demonstration, testing, or development purposes, NVIDIA directly infringes the **'724 Patent**.

2741. Velocity assembles laptop computers that incorporate **'724 Accused Mobile GPUs**.

2742. For example, Velocity makes a laptop computer called the NoteMagix M15.

2743. The NoteMagix M15 includes a **'724 Accused Mobile GPU**, the GeForce GTX 770M.

2744. Velocity assembles its own products into infringing systems in this District.

2745. For example, Velocity assembles its own products into infringing systems in this District for demonstration, testing, and development purposes.

2746. Velocity has assembled and continues to assemble its own products, including the **'724 Accused Velocity Laptops**, into infringing systems, such as a laptop computer powered on

and connected via DisplayPort to a powered external monitor, such as for demonstration, testing, and development purposes.

2747. For example and without limitation, Velocity has assembled and continues to assemble the NoteMagix M15 incorporating the GeForce GTX 770M into a laptop computer powered on and connected via DisplayPort to a powered external monitor.

2748. For example and without limitation, Velocity has assembled and continues to assemble laptop computers that incorporate **'724 Accused Mobile GPUs** into infringing systems.

2749. When Velocity assembles the **'724 Accused Products** into infringing systems, such as for demonstration, testing, or development purposes, such assembly directly infringes the **'724 Patent**.

2750. Velocity uses these infringing systems, such as for demonstration, testing, and development purposes.

2751. When Velocity uses the infringing systems described in the previous paragraphs, such as for demonstration, testing, or development purposes, Velocity directly infringes the **'724 Patent**.

2752. NVIDIA has had actual knowledge of the **'724 Patent** since at least November 12, 2014.

2753. Velocity has had actual knowledge of the **'724 Patent** since at least November 11, 2014.

2754. NVIDIA and Velocity indirectly infringe the **'724 Patent** by inducing infringement by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(b),

in this District and elsewhere in the United States. Direct infringement is the result of activities performed by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, by, for example, assembling the **'724 Accused Products** into infringing systems as intended by NVIDIA and by Velocity.

2755. NVIDIA and Velocity have induced and continue to induce infringement of the **'724 Patent** by intending that others infringe the **'724 Patent** by assembling the **'724 Accused Products** into infringing systems.

2756. NVIDIA and Velocity designed the **'724 Accused Products** such that they would each infringe one or more claims of the **'724 Patent** when assembled, as intended, into an infringing system.

2757. NVIDIA and Velocity provide the **'724 Accused Products** to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users.

2758. By providing **'724 Accused Products** to others, NVIDIA and Velocity intend for **'724 Accused Products** to be assembled into infringing systems in the United States.

2759. By providing **'724 Accused Products** to others, NVIDIA and Velocity intend for **'724 Accused Products** to be used in infringing systems in the United States.

2760. NVIDIA and Velocity also provide others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, instructions, user guides, and technical specifications. When others follow such instructions,

user guides, and/or other design documentation, they directly infringe one or more claims of the **'724 Patent**. NVIDIA and Velocity know that by providing such instructions, user guides, and/or other design documentation, others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, follow those instructions, user guides, and other design documentation, and directly infringe one or more claims of the **'724 Patent**. NVIDIA and Velocity thus know that their actions actively induce infringement.

2761. NVIDIA sells or has sold **'724 Accused Mobile GPUs** in this District. Through sales of **'724 Accused Mobile GPUs**, NVIDIA specifically intends that others, such as customers and end users, will infringe one or more claims of the **'724 Patent**.

2762. NVIDIA provides the **'724 Accused Products**, and reference designs for the **'724 Accused Products**, to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, and resellers.

2763. NVIDIA also instructs others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, and resellers, to assemble the **'724 Accused Mobile GPUs** into infringing systems.

2764. For example, NVIDIA provides instructions on its website that tell the reader to “[m]ake sure both monitors are connected to the NVIDIA based graphics card” in order to setup multiple monitors. **NVIDIA Multiple Display Instructions**.

2765. When an external monitor is connected via DisplayPort to the NVIDIA based graphics card in a powered-on laptop computer in the manner NVIDIA instructs, such a system infringes the **'724 Patent**.

2766. Through such instruction, NVIDIA specifically intends that others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, and resellers, customers, and end users, infringe the **'724 Patent** by assembling in the United States the **'724 Accused Products** into infringing systems.

2767. Through its manufacture (either directly, or through contract manufacturing facilities) and/or sale of the infringing products, NVIDIA specifically intends that others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, will infringe one or more claims of the **'724 Patent**.

2768. NVIDIA specifically intends for others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **'724 Patent** in the United States.

2769. For example, NVIDIA advertised its “NVIDIA Authorized Board Partner Program ensures an exceptional customer experience when purchasing graphics cards and motherboards manufactured by partners that make use of NVIDIA’s latest technologies.”
NVIDIA PartnerForce Info.

2770. NVIDIA claims that its “NVIDIA Authorized Board Partners offer the latest technologies from NVIDIA” and lists six Authorized Board Partners in the United States, six in

Canada, and none in any other countries. *See* **NVIDIA PartnerForce Info**. NVIDIA also lists ten distributors in the United States (including “pre and post sales technical support”) and seven in Canada but none in any other country. *See* **NVIDIA PartnerForce Info**.

2771. As an additional example, as of October 31, 2014, NVIDIA also advertised the “NVIDIA PartnerForce Program” which is “a sales and marketing program for value-added resellers, system builders, etailers and retailers who sell NVIDIA based components or systems.”

NVIDIA PartnerForce Program.

2772. NVIDIA advertises in the United States and in this judicial district that “[b]y joining the PartnerForce Program, you may leverage the world-class brands and technology platforms from NVIDIA, sales and technical support from our experienced team and hundreds of turnkey sales and marketing tools.” **NVIDIA PartnerForce Program.**

2773. NVIDIA also provides marketing materials and templates, including retail display items, web banners, copy blocks, logos, product shots, email and web page templates, to members of the “NVIDIA PartnerForce Program” which includes value-added resellers, system builders, etailers and retailers inducing them to offer to sell the **’724 Accused Mobile GPUs**.

See **NVIDIA PartnerForce Program.**

2774. As described in the preceding paragraphs, NVIDIA specifically targets the United States market for **’724 Accused Mobile GPUs** and actively induces others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board manufacturers, resellers, customers, and end users, to directly infringe one or more claims of the **’724 Patent**.

2775. Velocity sells **’724 Accused Velocity Laptops**, which incorporate **’724 Accused Mobile GPUs**.

2776. For example, Velocity sells a laptop computer called the NoteMagix M15.

2777. The NoteMagix M15 includes a **'724 Accused Mobile GPU**, the GeForce GTX 770M.

2778. Velocity provides the **'724 Accused Velocity Laptops** to others, such as resellers, customers, and end users.

2779. Through such activity, Velocity specifically intends that others, such as resellers, customers, and end users, infringe the **'724 Patent** by assembling the **'724 Accused Velocity Laptops** into infringing systems in the United States.

2780. Velocity specifically targets the United States market for its products listed above and actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'724 Patent** in the United States.

2781. For example, Velocity sells products via its website, <http://www.velocitymicro.com>, to customers in the United States. Such products include the **'724 Accused Velocity Laptops**, such as, for example, laptop computers that are configured to include a GeForce GTX 770M.

2782. For further example, Velocity at one time offered for sale the NoteMagix M15 configured to include an external monitor.

2783. The customers infringe the **'724 Patent** by assembling the **'724 Accused Velocity Laptops** into infringing systems, such as a laptop computer powered on and connected via DisplayPort to a powered external monitor.

2784. Velocity also sells to resellers the **'724 Accused Velocity Laptops**, such as, for example, **'724 Accused Velocity Laptops** that are configured to include a GeForce GTX 770M.

2785. For example, Velocity products have been sold at nearly every major electronics retailer in North America including Best Buy, Circuit City, RadioShack, Costco, Sears, Target, and many others. *See Velocity Company History.*

2786. Velocity introduces products and services that infringe the Asserted Patents intending their assembly into infringing systems in this Judicial District and elsewhere in the United States.

2787. The resellers infringe the **'724 Patent** by assembling the **'724 Accused Products** into infringing systems.

2788. The resellers also sell the **'724 Accused Products** to customers and end users, who infringe the **'724 Patent** by assembling the **'724 Accused Products** into infringing systems.

2789. As described in the preceding paragraphs, Velocity specifically targets the United States market for **'724 Accused Products** actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'724 Patent**.

2790. NVIDIA and Velocity indirectly infringe the **'724 Patent** by contributing to infringement by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(c), in this District and elsewhere in the United States.

2791. Direct infringement is the result of assembly of infringing systems by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users.

2792. For example, end users and resellers assemble the **'724 Accused Products** to connect via DisplayPort to a DisplayPort-capable monitor.

2793. The DisplayPort features of the **'724 Accused Products** are a material component of the patented machine claimed by the **'724 Patent**.

2794. NVIDIA and Velocity knew the **'724 Accused Products'** DisplayPort features are a material component of the patented machine claimed by the **'724 Patent**.

2795. The DisplayPort features of the **'724 Accused Products** are not a staple article suitable for substantial non-infringing use.

2796. NVIDIA and Velocity knew the DisplayPort features of the infringing products were especially made or especially adapted to operate in the infringing systems of NVIDIA's and Velocity's distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, and is not a staple article or commodity of commerce and that its infringing use is required for operation of the infringing products. Any other use would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental.

2797. For example, NVIDIA and Velocity knew users would connect the **'724 Accused Products** to an external monitor, resulting in an infringing system.

2798. For further example, NVIDIA and Velocity also knew the **'724 Accused Products'** DisplayPort features are especially adapted to connect via a DisplayPort cable to a DisplayPort-capable external monitor.

2799. Accordingly, NVIDIA and Velocity offer to sell or sell a component of the patented machine claimed by the **'724 Patent**, knowing the same to be especially made or

especially adapted for use in an infringement of such patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

2800. Samsung has no adequate remedy at law for NVIDIA's and Velocity's infringement of the **'724 Patent** and is suffering irreparable harm, requiring permanent injunctive relief.

SEVENTH CLAIM FOR RELIEF
INFRINGEMENT OF U.S. PATENT NO. 7,073,054
(AGAINST VELOCITY)

2801. Each of the above listed paragraphs is incorporated herein by reference and adopted, as if fully set forth again.

2802. The **'054 Patent** was filed on August 16, 2002, issued on July 4, 2006, and is entitled "Computer System And Method For Booting Up The Same." The **'054 Patent** is generally directed to a computer system with a hard disk drive and non-volatile memory and a method for reducing the booting time of the computer.

2803. The **'054 Patent** was assigned to SEC, and SEC continues to hold all rights, title, and interest in the **'054 Patent**. A true and correct copy of the **'054 Patent** is attached hereto as Exhibit G.

2804. All **'054 Accused Products** infringe the **'054 Patent** in substantially the same way. All **'054 Accused Products** are computer systems that contain a main memory, a control unit or controller capable of reading a booting program, and a hybrid hard drive further containing a motor-driven hard disk and a nonvolatile storage unit.

2805. Commercially-available hybrid hard drives store, read, load, and transmit a booting program in substantially the same way, or in ways that are functionally equivalent.

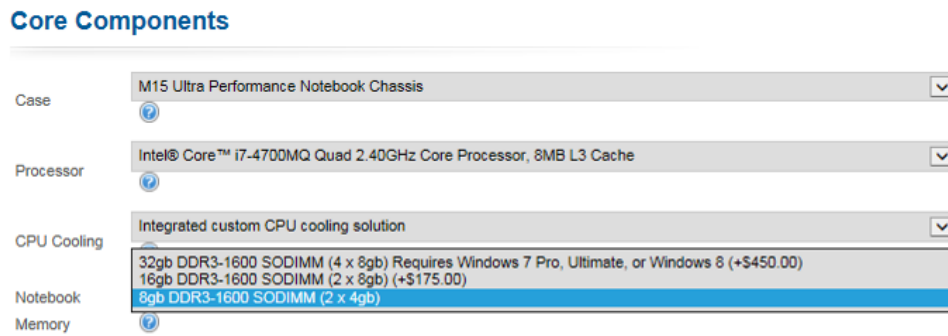
2806. "**Claim 1(a) of the '054 Patent**" recites "[a] computer system."

2807. All **'054 Accused Products** are computer systems.

2808. **Claim 1(b) of the '054 Patent** recites “a main memory.”

2809. All **'054 Accused Products** include a main memory.

2810. For example, the **NoteMagix M15 HHD** includes a main memory. As shown below, Velocity offers three options for main memory of its **NoteMagix M15 HHD** computer system: 8 GB DDR3-1600 SODIMM, 16 GB DDR3-1600 SODIMM, and 32GB DDR3-1600 SODIMM.



Archived NoteMagix M15 HHD Webpage.

2811. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a main memory.

2812. As a result, the **'054 Accused Products** meet the limitations of **Claim 1(b) of the '054 Patent**.


2813. “**Claim 1(c) of the '054 Patent**” recites “a disk drive in communication with the main memory.”

2814. All **'054 Accused Products** include a disk drive in communication with the main memory.

2815. For example, the **NoteMagix M15 HHD** includes a disk drive in communication with the main memory.

2816. As shown below, Velocity offers the **NoteMagix M15 HHD** with a hybrid hard drive configuration, listed as “750gb 7200rpm Mobile Hard Drive with 8gb SSD Cache.” Thus, the hybrid hard drive of the **NoteMagix M15 HHD** is a disk drive in communication with the main memory.

Storage



Archived NoteMagix M15 HHD Webpage.

2817. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a main memory with a disk drive in communication with the main memory.


2818. As a result, the **'054 Accused Products** meet the limitations of **Claim 1(c) of the '054 Patent**.

2819. “**Claim 1(d) of the '054 Patent**” recites “a driving motor driving a disk of the disk drive.”

2820. All **'054 Accused Products** include a driving motor driving a disk of the disk drive.

2821. For example, the **NoteMagix M15 HHD** includes a driving motor driving a disk of the disk drive. As shown below, the **Note Magix M15** advertised that the hybrid hard drive has a 750GB disk with a rotational speed of 7200 rotations per minute. Thus, the hybrid hard drive of the **NoteMagix M15 HHD** includes a driving motor driving a disk of the disk drive.

Storage



Archived NoteMagix M15 HHD Webpage.

2822. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a driving motor driving a disk of a disk drive.

2823. As a result, the **'054 Accused Products** meet the limitations of **Claim 1(d) of the '054 Patent**.

2824. "**Claim 1(e) of the '054 Patent**" recites a "non-volatile storage unit storing a booting program of an operating system."

2825. All **'054 Accused Products** include a non-volatile storage unit storing a booting program of an operating system.

2826. For example, the **NoteMagix M15 HHD** includes a non-volatile storage unit storing a booting program of an operating system. As shown below, the **NoteMagix M15 HHD** includes an "8GB SSD Cache." The 8GB SSD (solid state drive) cache is a non-volatile storage unit.

Storage



Archived NoteMagix M15 HHD Webpage.

2827. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides "Solid State Drive options for quick boot up and fast response."

This Velocity Micro® **NoteMagix™ M15** system includes:

- 4th Gen (Haswell) Intel® Core® i7 processor
- Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
- 15.6" 1080p LED backlit display
- Solid State Drive options for quick boot up and fast response
- Wireless connectivity with Bluetooth and 802.11b/g/n
- Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
- Rigorous quality control testing
- Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

2828. To provide the “quick boot up” advertised by Velocity, the booting program of the operating system is stored on the non-volatile storage unit (the 8GB SSD cache), not the motor-driven disk of the hybrid hard drive.

2829. A motor-driven hard disk requires time to speed up to its normal operating speed. The **NoteMagix M15 HHD** achieves the “quick boot up” advertised by Velocity because the booting program of the operating system is stored on the non-volatile storage unit (the 8GB of SSD cache).

2830. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that stores a booting program of an operating system on a non-volatile storage unit of a disk drive.

2831. As a result, the **'054 Accused Products** meet the limitations of **Claim 1(e) of the '054 Patent**.

2832. “**Claim 1(f) of the '054 Patent**” recites “a control unit reading the booting program stored in the non-volatile storage unit and loading the booting program onto the main memory before the driving motor reaches a normal speed as power is supplied to the computer system.”

2833. All **'054 Accused Products** include a control unit that reads the booting program stored in the non-volatile storage unit and loads the booting program onto the main memory before the driving motor reaches a normal speed as power is supplied to the computer system.

2834. For example, the **NoteMagix M15 HHD** includes a control unit reading the booting program stored in the non-volatile storage unit and loading the booting program onto the main memory before the driving motor reaches a normal speed as power is supplied to the computer system.

2835. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides “Solid State Drive options for quick boot up and fast response.” A motor-driven hard disk requires time to speed up to its normal operating speed. To provide the “quick boot up” advertised by Velocity, a control unit must read the booting program stored on the non-volatile storage unit (the 8GB SSD cache) and load the booting program onto the main memory before the motor-driven disk of the hybrid hard drive reaches its normal operating speed. The control unit will read and load the booting program as the power is supplied to the computer system.

This Velocity Micro® **NoteMagix™ M15** system includes:

- 4th Gen (Haswell) Intel® Core® i7 processor
- Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
- 15.6" 1080p LED backlit display
- Solid State Drive options for quick boot up and fast response
- Wireless connectivity with Bluetooth and 802.11b/g/n
- Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
- Rigorous quality control testing
- Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

2836. Velocity made, sold, or offered for sale since November 4, 2008 a computer system having a control unit that loads a booting program from a non-volatile storage unit of a disk drive to the main memory.

2837. Velocity made, sold, or offered for sale since November 4, 2008 a computer system having a control unit that loads a booting program from a non-volatile storage unit of a disk drive to the main memory before the motor-driven disk of the disk drive reaches a normal operating speed.

2838. As a result, the **'054 Accused Products** meet the limitations of **Claim 1(f) of the '054 Patent**.

2839. "**Claim 2 of the '054 Patent**" recites "[t]he computer system according to claim 1, wherein the control unit stores the booting program in the non-volatile storage unit from the disk when installing the operating system."

2840. All **'054 Accused Products** also include a control unit that may store the booting program on the non-volatile storage unit of the hybrid hard disk when installing the operating system.

2841. For example, the **NoteMagix M15 HHD** includes a control unit that stores the booting program on the disk drive.

2842. The control unit discussed in the preceding paragraph may store the booting program on the non-volatile storage unit of the disk drive, after reading the booting program from the disk of the disk drive.

2843. During the installation process, the control unit stores the booting program on the non-volatile storage unit (the 8GB SSD cache) from the disk of the disk drive to ensure quick boot up times.

2844. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a control unit that stores a booting program in the non-volatile storage unit of the hard drive during installation of the operating system.


2845. As a result, the **'054 Accused Products** meet the limitations of **Claim 2 of the '054 Patent**.

2846. “**Claim 3 of the '054 Patent**” recites “[t]he computer system according to claim 1, wherein the non-volatile storage unit is a flash memory.”

2847. For all **'054 Accused Products**, the non-volatile storage unit is a flash memory.

2848. For example, as shown below, the **NoteMagix M15 HHD** offers a hybrid hard drive option with 8GB SSD cache, which is a flash memory.

Storage



Archived NoteMagix M15 HHD Webpage.

2849. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a disk drive having a non-volatile storage unit that is flash memory.

2850. As a result, the **'054 Accused Products** meet the limitations of **Claim 3 of the '054 Patent**.

2851. “**Claim 4 of the '054 Patent**” recites “[a] method of booting up a computer system.”

2852. All **'054 Accused Products** perform a method of booting up a computer system.

2853. For example, as shown below, the **NoteMagix M15 HHD** performs a method of booting up a computer system.

2854. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides “Solid State Drive options for quick boot up and fast response.”

This Velocity Micro® **NoteMagix™ M15** system includes:

- 4th Gen (Haswell) Intel® Core® i7 processor
- Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
- 15.6" 1080p LED backlit display
- Solid State Drive options for quick boot up and fast response
- Wireless connectivity with Bluetooth and 802.11b/g/n
- Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
- Rigorous quality control testing
- Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

2855. To provide the “quick boot up” advertised by Velocity, the **NoteMagix M15 HHD** must perform a method of booting up a computer system.

2856. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that performs a method of booting up a computer system.

2857. As a result, the **'054 Accused Products** practice a method of booting up a computer system.

2858. “**Claim 4(b) of the '054 Patent**” recites “reading a booting program of an operating system from a disk of the disk drive.”

2859. The **'054 Accused Products** read a booting program of an operating system from a disk of the disk drive.

2860. For example, the **NoteMagix M15 HHD** reads a booting program of an operating system from a disk of the disk drive after the booting program is installed on the disk during the installation of the operating system.

2861. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that reads a booting program from a hard disk of a hard drive.


2862. As a result, the **'054 Accused Products** meet the limitations of **Claim 4(b) of the '054 Patent**.

2863. “**Claim 4(c) of the '054 Patent**” recites “storing the booting program in a predetermined storage unit in the disk drive when installing an operating system on the disk of the disk drive.”

2864. The **'054 Accused Products** store the booting program in a predetermined storage unit in the disk drive when installing an operating system on the disk of the disk drive.

2865. For example, as shown below, the **NoteMagix M15 HHD** offers a hybrid hard drive option with 8GB SSD cache, which is a flash memory.

Storage



Archived NoteMagix M15 HHD Webpage.

2866. And, as shown below, Velocity advertised that the **NoteMagix M15 HHD** provides “Solid State Drive options for quick boot up and fast response.”

- This Velocity Micro® **NoteMagix™ M15** system includes:
- 4th Gen (Haswell) Intel® Core® i7 processor
 - Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
 - 15.6" 1080p LED backlit display
 - Solid State Drive options for quick boot up and fast response
 - Wireless connectivity with Bluetooth and 802.11b/g/n
 - Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
 - Rigorous quality control testing
 - Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

2867. To provide the “quick boot up” advertised by Velocity, the **NoteMagix M15 HHD** may store the booting program of the operating system on the predetermined storage unit (for example, the 8 GB SSD cache), while installing an operating system. The **NoteMagix M15 HHD** may choose the predetermined storage unit instead of the motor-driven disk of the hybrid

hard drive, because the predetermined storage unit provides faster boot up than the motor-driven disk, which requires time to speed up to its normal operating speed.

2868. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that stores a booting program on a predetermined storage unit of a hard drive.

2869. As a result, the **'054 Accused Products** meet the limitations of **Claim 4(c) of the '054 Patent**.

2870. "**Claim 4(d) of the '054 Patent**" recites "reading the stored booting program and loading the booting program onto a main memory of the computer system before a driving motor of the disk drive reaches a normal speed as power is supplied to the computer system"

2871. All **'054 Accused Products** include a control unit that reads the booting program stored in the non-volatile storage unit and loads the booting program onto the main memory before the driving motor reaches a normal speed as power is supplied to the computer system.

2872. For example, the **NoteMagix M15 HHD** includes a control unit reading the booting program stored in the non-volatile storage unit and loading the booting program onto the main memory before the driving motor reaches a normal speed as power is supplied to the computer system.

2873. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides "Solid State Drive options for quick boot up and fast response." A motor-driven hard disk requires time to speed up to its normal operating speed. To provide the "quick boot up" advertised by Velocity, a control unit must read the booting program stored on the non-volatile storage unit (the 8GB SSD cache) and load the booting program onto the main memory before the motor-driven disk of the hybrid hard drive reaches its normal operating speed. The control unit will read and load the booting program as the power is supplied to the computer system.

This Velocity Micro® **NoteMagix™ M15** system includes:

- 4th Gen (Haswell) Intel® Core® i7 processor
- Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
- 15.6" 1080p LED backlit display
- Solid State Drive options for quick boot up and fast response
- Wireless connectivity with Bluetooth and 802.11b/g/n
- Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
- Rigorous quality control testing
- Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

2874. Velocity made, sold, or offered for sale since November 4, 2008 a computer system having a control unit that loads a booting program from a non-volatile storage unit of a disk drive to the main memory.

2875. Velocity made, sold, or offered for sale since November 4, 2008 a computer system having a control unit that loads a booting program from a non-volatile storage unit of a disk drive to the main memory before the motor-driven disk of the disk drive reaches a normal operating speed.

2876. As a result, the **'054 Accused Products** meet the limitations of **Claim 4(d) of the '054 Patent**.

2877. "**Claim 4(e) of the '054 Patent**" recites "booting up the computer system with the booting program loaded onto the main memory."

2878. All **'054 Accused Products** boot up the computer system with the booting program loaded onto the main memory.

2879. For example, the **NoteMagix M15 HHD** loads booting program in the main memory, and the CPU reads the booting program from the main memory to boot the computer system.

2880. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 boots up with the booting program loaded onto the main memory.


2881. As a result, the **'054 Accused Products** meet the limitations of **Claim 4(e) of the '054 Patent**.

2882. "**Claim 5 of the '054 Patent**" recites "[t]he method according to claim 4, wherein the storage unit is a non-volatile storage unit."

2883. All **'054 Accused Products** also contain disk drive with a non-volatile storage unit.

2884. For example, as shown below, the **NoteMagix M15 HHD** offers a hybrid hard drive option with 8GB SSD cache, which is a non-volatile storage unit.

Storage



Archived NoteMagix M15 HHD Webpage.

2885. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a disk drive having a non-volatile storage unit.

2886. As a result, the **'054 Accused Products** meet the limitations of **Claim 5 of the '054 Patent**.

2887. "**Claim 6(a) of the '054 Patent**" recites "[a] computer system."

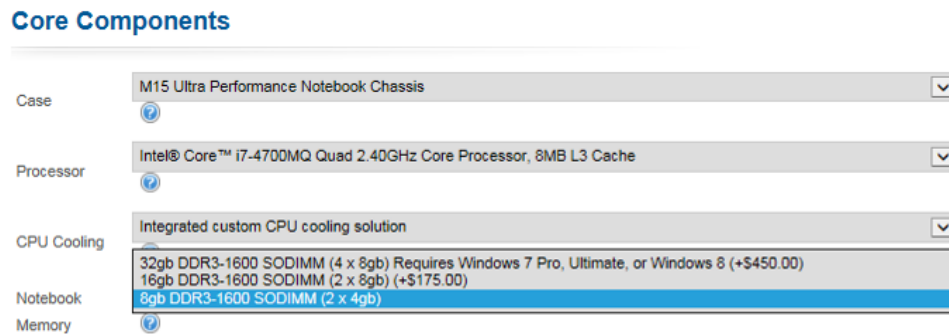
2888. All **'054 Accused Products** are computer systems.

2889. Velocity made, sold, or offered for sale since November 4, 2008 at least one computer system.

2890. "**Claim 6(b) of the '054 Patent**" recites "a main memory."

2891. All **'054 Accused Products** include a main memory.

2892. For example, the **NoteMagix M15 HHD** includes a main memory. As shown below, Velocity offered three options for main memory of its **NoteMagix M15 HHD** computer system: 8 GB DDR3-1600 SODIMM, 16 GB DDR3-1600 SODIMM, and 32GB DDR3-1600 SODIMM.



Archived NoteMagix M15 HHD Webpage.

2893. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a main memory.

2894. As a result, the **'054 Accused Products** meet the limitations of **Claim 6(b) of the '054 Patent**.


2895. "**Claim 6(c) of the '054 Patent**" recites "a disk drive in communication with the main memory."

2896. All **'054 Accused Products** include a disk drive in communication with the main memory.

2897. For example, the **NoteMagix M15 HHD** includes "a disk drive in communication with the main memory."

2898. As shown below, Velocity offers the **NoteMagix M15 HHD** with a hybrid hard drive configuration, listed as “750gb 7200rpm Mobile Hard Drive with 8gb SSD Cache.” Thus, the hybrid hard drive of the **NoteMagix M15 HHD** is a disk drive in communication with the main memory.

Storage



Archived NoteMagix M15 HHD Webpage.

2899. Velocity made, sold, or offers for sale since November 4, 2008 a computer system that includes a main memory with a disk drive in communication with the main memory.


2900. As a result, the **'054 Accused Products** meet the limitations of **Claim 6(c) of the '054 Patent**.

2901. “**Claim 6(d) of the '054 Patent**” recites “a driving motor driving a disk of the disk drive.”

2902. All **'054 Accused Products** include a driving motor driving a disk of the disk drive.

2903. For example, the **NoteMagix M15 HHD** includes a driving motor driving a disk of the disk drive. As shown below, the **Note Magix M15** advertised that the hybrid hard drive has a 750GB disk with a rotational speed of 7200 rotations per minute. Thus, the hybrid hard drive of the **NoteMagix M15 HHD** includes a driving motor driving a disk of the disk drive.

Storage



Archived NoteMagix M15 HHD Webpage.

2904. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a driving motor driving a disk of a disk drive.


2905. As a result, the **'054 Accused Products** meet the limitations of **Claim 6(d) of the '054 Patent**.

2906. **"Claim 6(e) of the '054 Patent"** recites "a storage unit storing a booting program of the computer system."

2907. All **'054 Accused Products** include a storage unit storing a booting program of a computer system.

2908. For example, the **NoteMagix M15 HHD** includes a storage unit storing a booting program of a computer system. As shown below, the **NoteMagix M15 HHD** includes an "8GB SSD Cache." The 8GB SSD (solid state drive) cache is a storage unit.

Storage



Archived NoteMagix M15 HHD Webpage.

2909. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides "Solid State Drive options for quick boot up and fast response."

This Velocity Micro® **NoteMagix™ M15** system includes:

- 4th Gen (Haswell) Intel® Core® i7 processor
- Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
- 15.6" 1080p LED backlit display
- Solid State Drive options for quick boot up and fast response
- Wireless connectivity with Bluetooth and 802.11b/g/n
- Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
- Rigorous quality control testing
- Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

2910. To provide the “quick boot up” advertised by Velocity, the booting program of the computer system is stored on the storage unit (the 8GB SSD cache), not the motor-driven disk of the hybrid hard drive.

2911. A motor-driven hard disk requires time to speed up to its normal operating speed. The **NoteMagix M15 HHD** achieves the “quick boot up” advertised by Velocity because the booting program of the computer system is stored on the storage unit (the 8GB of SSD cache).

2912. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that stores a booting program of a computer system on a storage unit of a disk drive.

2913. As a result, the **'054 Accused Products** meet the limitations of **Claim 6(e) of the '054 Patent**.

2914. “**Claim 6(f) of the '054 Patent**” recites “a control unit loading the booting program stored in the storage unit into the main memory before the driving motor reaches a normal speed when power is supplied to the computer system.”

2915. All **'054 Accused Products** include a control unit that loads the booting program stored in the storage unit into the main memory before the driving motor reaches a normal speed when power is supplied to the computer system.

2916. For example, the **NoteMagix M15 HHD** includes a control unit that loads the booting program stored in the storage unit into the main memory before the driving motor reaches a normal speed when power is supplied to the computer system.

2917. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides “Solid State Drive options for quick boot up and fast response.” A motor-driven hard disk requires time to speed up to its normal operating speed. To provide the “quick boot up” advertised by Velocity, a control unit must load the booting program stored on the non-volatile storage unit (the 8GB SSD cache) onto the main memory before the motor-driven disk of the hybrid hard drive reaches its normal operating speed. The control unit will begin loading the booting program when the power is supplied to the computer system.

- This Velocity Micro® **NoteMagix™ M15** system includes:
- 4th Gen (Haswell) Intel® Core® i7 processor
 - Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
 - 15.6" 1080p LED backlit display
 - Solid State Drive options for quick boot up and fast response
 - Wireless connectivity with Bluetooth and 802.11b/g/n
 - Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
 - Rigorous quality control testing
 - Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

2918. Velocity made, sold, or offered for sale since November 4, 2008 a computer system with a control unit that loads the booting program stored in the storage unit into the main memory before the driving motor reaches a normal speed.


2919. As a result, the **'054 Accused Products** meet the limitations of **Claim 6 of the '054 Patent**.

2920. “**Claim 7 of the '054 Patent**” recites “[t]he computer system of claim 6, wherein the storage unit is non-volatile.”

2921. All **'054 Accused Products** include a hard disk drive with a non-volatile storage unit.

2922. For example, as shown below, the **NoteMagix M15 HHD** offers a hybrid hard drive option with 8GB SSD cache, which is a non-volatile storage unit.

Storage



Archived NoteMagix M15 HHD Webpage.

2923. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a disk drive with a non-volatile storage unit.


2924. As a result, the **'054 Accused Products** meet the limitations of **Claim 7 of the '054 Patent**.

2925. "**Claim 8 of the '054 Patent**" recites "[t]he computer system of claim 7, wherein the non-volatile storage unit is flash memory."

2926. For all **'054 Accused Products**, the non-volatile storage unit is a flash memory.

2927. For example, as shown below, the **NoteMagix M15 HHD** offers a hybrid hard drive option with 8GB SSD cache, which is a flash memory.

Storage



Archived NoteMagix M15 HHD Webpage.

2928. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a disk drive having a non-volatile storage unit that is flash memory.

2929. As a result, the **'054 Accused Products** meet the limitations of **Claim 8 of the '054 Patent**.

2930. “**Claim 9 of the '054 Patent**” recites “[t]he computer system of claim 6, wherein the control unit updates the booting program stored in the storage unit.”

2931. All **'054 Accused Products** also include a control unit that updates the booting program stored in the storage unit.

2932. For example, updates to the booting program of the **NoteMagix M15 HHD** may be installed at the location where the booting program is stored. Because the booting program of **claim 6** is stored on the storage unit, the control unit will update the booting program stored on the storage unit.

2933. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a control unit that updates the booting program stored on the non-volatile storage unit.

2934. As a result, the **'054 Accused Products** meet the limitations of **Claim 9 of the '054 Patent**.

2935. “**Claim 10 of the '054 Patent**” recites “[t]he computer system of claim 6, wherein the control unit stores the booting program in the storage unit during installation of an operating system for the computer system.”

2936. The **'054 Accused Products** also include a control unit that stores the booting program on the storage unit of the hybrid hard disk when installing the operating system.

2937. For example, the **NoteMagix M15 HHD** includes a control unit that stores the booting program on the disk drive.

2938. The control unit discussed in the preceding paragraph stores the booting program on the storage unit of the disk drive (the 8GB SSD cache) during the installation process. Because the storage unit does not need to speed up to normal operating speeds, like a motor-driven disk, storing the booting program on the storage unit proves faster boot up times.


2939. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a controller that stores a booting program in the storage unit of the hard drive.

2940. “**Claim 11(a) of the '054 Patent**” recites “[a] computer hard disk drive having a disk.”

2941. All **'054 Accused Products** include a computer hard disk drive having a disk.

2942. As shown below, Velocity offers the **NoteMagix M15 HHD** with a hybrid hard drive configuration, listed as “750gb 7200rpm Mobile Hard Drive with 8gb SSD Cache.” Thus, the hybrid hard drive of the **NoteMagix M15 HHD** is a computer hard drive having a disk, the 750 disk drive.

Storage



Archived NoteMagix M15 HHD Webpage.

2943. Velocity made, sold, or offered for sale since November 4, 2008 a computer system with a hard disk driving having a disk.

2944. As a result, the **'054 Accused Products** include a computer hard disk drive having a disk.

2945. “**Claim 11(b) of the '054 Patent**” recites “a driving motor driving the disk.”

2946. All **'054 Accused Products** include a driving motor driving a disk of a disk drive.

2947. For example, the **NoteMagix M15 HHD** includes a driving motor driving a disk of a disk drive. As shown below, the **Note Magix M15** advertised that the hybrid hard drive has a 750GB disk with a rotational speed of 7200 rotations per minute. Thus, the hybrid hard drive of the **NoteMagix M15 HHD** includes a driving motor driving a disk of a disk drive.

Storage



Archived NoteMagix M15 HHD Webpage.

2948. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a driving motor driving a disk of a disk drive.


2949. As a result, the **'054 Accused Products** meet the limitations of **Claim 11(b) of the '054 Patent**.

2950. "**Claim 11(c) of the '054 Patent**" recites "a storage unit storing a booting program of a computer system controlling the disk drive."

2951. All **'054 Accused Products** include a storage unit storing a booting program of a computer system controlling the disk drive.

2952. For example, the **NoteMagix M15 HHD** includes a non-volatile storage unit storing a booting program of a computer system controlling the disk drive. As shown below, the **NoteMagix M15 HHD** includes an "8GB SSD Cache." The 8GB SSD (solid state drive) cache is a storage unit.

Storage



Archived NoteMagix M15 HHD Webpage.

2953. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides “Solid State Drive options for quick boot up and fast response.”

- This Velocity Micro® **NoteMagix™ M15** system includes:
- 4th Gen (Haswell) Intel® Core® i7 processor
 - Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
 - 15.6" 1080p LED backlit display
 - Solid State Drive options for quick boot up and fast response
 - Wireless connectivity with Bluetooth and 802.11b/g/n
 - Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
 - Rigorous quality control testing
 - Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

2954. To provide the “quick boot up” advertised by Velocity, the booting program of the computer system controlling the disk drive is stored on the storage unit (the 8GB SSD cache), not the motor-driven disk of the hybrid hard drive.

2955. A motor-driven hard disk requires time to speed up to its normal operating speed. The **NoteMagix M15 HHD** drive achieves the “quick boot up” advertised by Velocity because the booting program of the computer system controlling the disk drive is stored on the storage unit (the 8GB of SSD cache).

2956. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a storage unit storing a booting program of a computer system controlling the disk drive.

2957. As a result, the **'054 Accused Products** meet the limitations of **Claim 11(c) of the '054 Patent**.

2958. “**Claim 11(d) of the '054 Patent**” recites “a controller in communication with the computer system and transmitting the booting program to the computer system before the driving motor reaches a normal speed when power is supplied to the computer system.”

2959. All **'054 Accused Products** include a controller in communication with the computer system and transmitting the booting program to the computer system before the driving motor reaches a normal speed when power is supplied to the computer system.

2960. For example, the **NoteMagix M15 HHD** includes a controller in communication with the computer system and transmitting the booting program to the computer system before the driving motor reaches a normal speed when power is supplied to the computer system.

2961. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides “Solid State Drive options for quick boot up and fast response.” A motor-driven hard disk drive requires time to speed up to its normal operating speed. To provide the “quick boot up” advertised by Velocity, the controller must be in communication with the computer system. The controller must then transmit the booting program from the storage unit (the 8GB SSD cache) to the computer system before the motor-driven disk reaches its normal operating speed.

This Velocity Micro® **NoteMagix™ M15** system includes:

- 4th Gen (Haswell) Intel® Core® i7 processor
- Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
- 15.6" 1080p LED backlit display
- Solid State Drive options for quick boot up and fast response
- Wireless connectivity with Bluetooth and 802.11b/g/n
- Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
- Rigorous quality control testing
- Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

2962. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a controller transmits the booting program to the computer system before the driving motor reaches a normal speed.


2963. As a result, the **'054 Accused Products** meet the limitations of **Claim 11 of the '054 Patent**.

2964. "**Claim 12 of the '054 Patent**" recites "[t]he computer hard disk drive of claim 11, wherein the storage unit is nonvolatile."

2965. All **'054 Accused Products** include a hard disk drive with a non-volatile storage unit.

2966. For example, as shown below, the **NoteMagix M15 HHD** offers a hybrid hard drive option with 8GB SSD cache, which is a non-volatile storage unit.

Storage



Archived NoteMagix M15 HHD Webpage.

2967. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a disk drive with a non-volatile storage unit.


2968. As a result, the **'054 Accused Products** meet the limitations of **Claim 12 of the '054 Patent**.

2969. "**Claim 13 of the '054 Patent**" recites "[t]he computer hard disk drive of claim 12, wherein the non-volatile storage unit is flash memory."

2970. For all **'054 Accused Products**, the non-volatile storage unit is a flash memory.

2971. For example, as shown below, the **NoteMagix M15 HHD** offers a hybrid hard drive option with 8GB SSD cache, which is a flash memory.

Storage



Archived NoteMagix M15 HHD Webpage.

2972. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that includes a disk drive having a non-volatile storage unit that is flash memory.

2973. As a result, the **'054 Accused Products** meet the limitations of **Claim 8 of the '054 Patent**.

2974. "**Claim 14 of the '054 Patent**" recites "[t]he computer hard disk drive of claim 12, wherein the controller updates the booting program stored in the storage unit."

2975. All **'054 Accused Products** also include a controller that updates the booting program stored in the storage unit.

2976. For example, updates to the booting program of the **NoteMagix M15 HHD** may be installed at the location where the booting program is stored. Because the booting program of **claim 6** is stored on the storage unit, the controller will update the booting program stored on the storage unit.

2977. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a controller that updates the booting program stored on the non-volatile storage unit.

2978. As a result, the **'054 Accused Products** meet the limitations of **Claim 12 of the '054 Patent**.

2979. “**Claim 15 of the ’054 Patent**” recites “[t]he computer hard disk drive of claim 14, wherein the controller stores the booting program on the disk and the controller updates the booting program in the storage unit responsive to changes in the booting program stored on the disk.”

2980. The **’054 Accused Products** include a controller that stores the booting program on the disk and updates the booting program in the storage unit responsive to changes in the booting program stored on the disk.

2981. For example, the **NoteMagix M15 HHD** includes a controller that loads updates to the operating system, including the booting program, on the hard disk.

2982. The controller described in the preceding paragraph updates the booting program in the storage unit based on changes to the booting program stored on the disk.

2983. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a controller that stores updates a booting program stored on a storage unit on a hard disk.

2984. As a result, the **’054 Accused Products** meet the limitations of **Claim 15 of the ’054 Patent**.

2985. “**Claim 16 of the ’054 Patent**” recites “[t]he computer hard disk drive of claim 12, wherein the controller stores the booting program in the storage unit during installation of an operating system for the computer system on the disk of the disk drive.”

2986. All **’054 Accused Products** also include a controller that may store the booting program in the storage unit during installation of an operating system for the computer system on the disk of the disk drive.

2987. For example, the **NoteMagix M15 HHD** includes a controller that stores the booting program on the disk drive during installation of the operating system.

2988. During the installation process, the controller discussed in the preceding paragraph stores the booting program on the storage unit of the disk drive (the 8GB SSD cache) to ensure quick boot up times.

2989. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a controller that stores a booting program in the storage unit of the hard drive during installation of the operating system.

2990. As a result, the **'054 Accused Products** meet the limitations of **Claim 16 of the '054 Patent**.

2991. "**Claim 17 of the '054 Patent**" recites "[t]he computer hard disk drive of claim 16, wherein the controller installs the operating system on the disk of the disk drive responsive to commands from the computer system."

2992. All **'054 Accused Products** include a controller that installs the operating system on the disk of a disk drive in response to commands from the computer system.

2993. For example, the **NoteMagix M15 HHD** includes a controller that transmits commands between the computer system and the disk drive.

2994. The controller described in the preceding paragraph responds to commands from the computer system to install the operating system on the disk drive.

2995. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a controller that responds to commands from the computer system to install the operating system on the disk of the hard drive.

2996. As a result, the **'054 Accused Products** meet the limitations of **Claim 17 of the '054 Patent**.

2997. “**Claim 18(a) of the '054 Patent**” recites “[a] controller controlling a computer-system hard disk drive having a disk and in communication with other components of the computer system.”

2998. All **'054 Accused Products** include a controller controlling a computer-system hard disk drive having a disk and in communication with other components of the computer system.

2999. For example, the **NoteMagix M15 HHD** includes a controller controlling a computer-system hard disk drive having a disk and in communication with other components of the computer system.

3000. As shown below, Velocity offers the **NoteMagix M15 HHD** with a hybrid hard drive configuration, listed as “750gb 7200rpm Mobile Hard Drive with 8gb SSD Cache.” The **NoteMagix M15 HHD** thus includes a controller that controls a hard disk (the 750 GB disk drive) and is in communication with other components of the computer system.



Archived NoteMagix M15 HHD Webpage.

3001. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a controller that (a) controls a hard disk drive and (b) communicates with other components of the computer system.

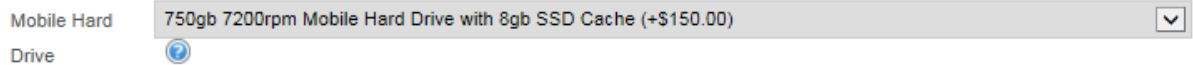
3002. As a result, the **'054 Accused Products** include a controller controlling a computer-system hard disk drive having a disk and in communication with other components of the computer system.

3003. **Claim 18(b) of the '054 Patent** includes the additional limitation of “storing a booting program in a memory in the hard disk drive.”

3004. All **'054 Accused Products** include a controller that stores a booting program in a memory in the hard disk drive.

3005. For example, the **NoteMagix M15 HHD** includes a controller that stores a booting program in a memory in the hard disk drive. As shown below, the **NoteMagix M15 HHD** includes the following hard drive option “750GB 7200rpm Mobile Hard Drive with 8GB SSD Cache.” The 8GB SSD (solid state drive) is a memory in the hard disk drive.

Storage



Mobile Hard Drive 750gb 7200rpm Mobile Hard Drive with 8gb SSD Cache (+\$150.00)

Archived NoteMagix M15 HHD Webpage.

3006. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides “Solid State Drive options for quick boot up and fast response.”

This Velocity Micro® **NoteMagix™ M15** system includes:

- 4th Gen (Haswell) Intel® Core® i7 processor
- Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
- 15.6" 1080p LED backlit display
- Solid State Drive options for quick boot up and fast response
- Wireless connectivity with Bluetooth and 802.11b/g/n
- Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
- Rigorous quality control testing
- Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

3007. To provide the “boot up” advertised by Velocity, the booting program of the operating system is stored in the memory of the hard disk drive, for example, the 8 GB SSD.

3008. Velocity made, sold, or offered for sale since November 4, 2008 a computer system that stores a booting program in the memory of a hard disk drive.

3009. As a result, the **'054 Accused Products** meet the limitations of **Claim 18(b) of the '054 Patent**.

3010. **Claim 18(c) of the '054 Patent** includes the additional limitation of “loading the booting program into a main memory of the computer system before a driving motor driving the disk of the hard disk drive reaches a normal speed as power is supplied to the computer system.”

3011. All **'054 Accused Products** include a controller that loads the booting program into the computer system’s main memory before the motor-driven disk reaches a normal speed as power is supplied to the computer system.

3012. For example, the **NoteMagix M15 HHD** includes a controller that loads the booting program into the computer system’s main memory before the motor-driven disk reaches a normal speed as power is supplied to the computer system.

3013. As shown below, Velocity advertised that the **NoteMagix M15 HHD** provides “Solid State Drive options for quick boot up and fast response.” A motor-driven hard disk requires time to speed up to its normal operating speed. To provide the “quick boot up” advertised by Velocity, a controller must load the booting program stored on the non-volatile storage unit (the 8GB SSD cache) before the motor-driven disk of the hybrid hard drive reaches its normal operating speed. In this way the controller will load the booting program into the main memory as the power is supplied to the computer system.

This Velocity Micro® **NoteMagix™ M15** system includes:

- 4th Gen (Haswell) Intel® Core® i7 processor
- Discrete graphics solution in the NVIDIA® GeForce® GTX 770M
- 15.6" 1080p LED backlit display
- Solid State Drive options for quick boot up and fast response
- Wireless connectivity with Bluetooth and 802.11b/g/n
- Wired connections High def video out, Display port, USB 3.0, eSATA, and Firewire
- Rigorous quality control testing
- Thorough performance benchmarking

Archived NoteMagix M15 HHD Webpage.

3014. Velocity made, sold, or offered for sale since November 4, 2008 a computer system with a controller that loads the booting program stored in the storage unit into the main memory before the driving motor of the hard disk reaches a normal speed.

3015. As a result, the **'054 Accused Products** meet the limitations of **Claim 18 of the '054 Patent**.

3016. “**Claim 19(a) of the '054 Patent**” recites “[t]he controller of claim 18, wherein the process further comprises storing the booting program of the computer system on the disk of the hard disk drive responsive to commands from the other components of the computer system.”

3017. The **'054 Accused Products** include a controller that may store the booting program of the computer system on the disk of the hard drive responsive to commands from other components of the computer system.

3018. For example, the **NoteMagix M15 HHD** includes a controller that may store the booting program on the motor-driven disk of the hard drive when the computer system is installing updates to the booting program.

3019. The controller described in the preceding paragraph stores the booting program on the disk of the hard drive responsive to commands from other components of the computer system, such as the main memory or an external drive.

3020. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a controller that stores the booting program of the computer system on the disk of the hard drive responsive to commands from other components of the computer system.

3021. As a result, the **'054 Accused Products** meet the limitations of **Claim 19(a) of the '054 Patent**.

3022. "**Claim 19(b) of the '054 Patent**" recites "updating the booting program in the memory from the disk."

3023. All **'054 Accused Products** include a controller that may update the booting program on the memory of the disk drive from the motor-driven disk.

3024. For example, the **NoteMagix M15 HHD** includes a controller may update the booting program on the memory of the disk drive (the 8 GB SSD Cache) from the motor-driven disk.

3025. If the updated booting program is first stored on the hard disk of the disk drive, the booting program stored on the memory of the disk drive (the 8 GB SSD Cache) would be updated from the updated booting program on the motor-driven disk.

3026. At least one computer system that Velocity has made, sold, or offered for sale since November 4, 2008 includes a controller that updates the booting program on the memory of the disk drive from the motor-driven disk.

3027. As a result, the **'054 Accused Products** meet the limitations of **Claim 19 of the '054 Patent**.

3028. Velocity has directly infringed and continues to directly infringe the **'054 Patent** by making, using, offering to sell, selling, and or importing the **'054 Accused Products**.

3029. For example, Velocity has used the **'054 Accused Products** for testing, demonstration, development, and configuration purposes.

3030. For example and without limitation, Velocity has used the **NoteMagix M15 HHD**.

3031. When Velocity uses the **'054 Accused Products**, such as for testing, demonstration, development, or configuration purposes, such use directly infringes the **'054 Patent**.

3032. Velocity has in the past and continues to make, offer to sell, sell, and import its own products, including the **'054 Accused Products**.

3033. For example and without limitation, Velocity has in the past and continues to make, offer to sell, sell, and import the **NoteMagix M15 HHD**.

3034. Velocity has sold at least one **'054 Accused Product** in this District.

3035. Velocity sells or offered to sell the **NoteMagix M15 HHD** via its online store at <http://www.velocitymicro.com>.

3036. Velocity sells or offered to sell the **'054 Accused Products** to end users and customers.

3037. Velocity sells or offered to sell the **'054 Accused Products** to intermediate customers including distributors, wholesalers, and retailers.

3038. When Velocity makes, offers to sell, or sells in the United States, or imports into the United States the **'054 Accused Products**, such activity directly infringes the **'054 Patent**.

3039. Velocity has had actual knowledge of the **'054 Patent** at least as of November 11, 2014 original complaint in this action served upon Velocity.

3040. Velocity indirectly infringe the **'054 Patent** by inducing infringement by others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, hard drive manufacturers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(b), in this District and elsewhere in the United States. Direct infringement is the result of activities performed by others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, hard drive manufacturers, resellers, customers, and end users, by, for example, making, using, offering to sell, selling, and or importing the **'054 Accused Products**.

3041. Velocity has induced and continues to induce infringement of the **'054 Patent** by intending that others infringe the **'054 Patent** by making, using, offering to sell, selling, and or importing the **'054 Accused Products**. Velocity designed the **'054 Accused Products** such that they would each infringe one or more claims of the **'054 Patent**.

3042. Velocity provides the **'054 Accused Products** to others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users.

3043. Through such activity, Velocity specifically intends that others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, infringe the **'054 Patent** by making, using, offering to sell, and selling in the United States, and importing into the United States the **'054 Accused Products**.

3044. Velocity also provide others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, hard drive manufacturers, resellers,

customers, and end users, instructions, user guides, and technical specifications. When others follow such instructions, user guides, and/or other design documentation, they directly infringe one or more claims of the **'054 Patent**. Velocity know that by providing such instructions, user guides, and/or other design documentation, others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, hard drive manufacturers, resellers, customers, and end users, follow those instructions, user guides, and other design documentation, and directly infringe one or more claims of the **'054 Patent**. Velocity thus know that its actions actively induce infringement.

3045. Velocity specifically targets the United States market for its products listed above and actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'054 Patent** in the United States.

3046. For example, Velocity sells products via its website, <http://www.velocitymicro.com>, to customers in the United States. Such products include the **'054 Accused Products**, such as, for example, the **NoteMagix M15 HDD**.

3047. The customers infringe the **'054 Patent** by using the **'054 Accused Products** purchased from Velocity.

3048. Velocity also sells computers to resellers, customers, and end users. Such computers include the **'054 Accused Products**, such as, for example, the **NoteMagix M15 HDD**.

3049. For example, Velocity products have been sold at nearly every major electronics retailer in North America including Best Buy, Circuit City, RadioShack, Costco, Sears, Target, and many others. *See Velocity Company History* Velocity introduces products and services

that infringe the Asserted Patents intending that they would be used in this Judicial District and elsewhere in the United States.

3050. The resellers infringe the **'054 Patent** by making, using, offering to sell, selling, and or importing the **'054 Accused Products**.

3051. The resellers also sell the **'054 Accused Products** to customers and end users, who infringe the **'054 Patent** by using those products.

3052. Through such sales, Velocity specifically intends that others, such as customers and end users, will infringe one or more claims of the **'054 Patent**.

3053. Velocity provides the **'054 Accused Products**, and technical specifications for the **'054 Accused Products**, to others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, hard drive manufacturers, resellers, customers, and end users.

3054. Through such activity, Velocity specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, hard drive manufacturers, resellers, customers, and end users, infringe the **'054 Patent** by making, using, offering to sell, and selling in the United States, and importing into the United States the **'054 Accused Products**.

3055. Through its manufacture (either directly, or through third-parties) and/or sale of the infringing products, Velocity specifically intends that others, such as distributors, wholesalers, retailers, original equipment manufacturers, contract equipment manufacturers, hard drive manufacturers, resellers, customers, and end users, will infringe one or more claims of the **'054 Patent**.

3056. As described in the preceding paragraphs, Velocity specifically targets the United States market for **'054 Accused Products** and actively induces others, such as resellers, customers, and end users, to directly infringe one or more claims of the **'054 Patent**.

3057. Velocity indirectly infringes the **'054 Patent** by contributing to infringement by others, such as distributors, wholesalers, retailers, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(c), in this District and elsewhere in the United States. Direct infringement is the result of activities performed by the distributors, wholesalers, retailers, resellers, customers, and end users of the infringing products.

3058. For example, the **'054 Accused Products** allow for controlling the hard drive to reduce the booting time of a computer. When the infringing products are used as intended by distributors, wholesalers, retailers, resellers, customers, and end users, the **'054 Accused Products** necessarily boot up the computer system in an infringing manner. The infringing products cannot boot up the computer system in an acceptable manner without infringing the **'054 Patent**.

3059. From the facts set forth above, it is evident that Velocity knew that the booting feature of the **'054 Accused Products** is especially made or especially adapted to boot up in the infringing manner and that feature is are not a staple article or commodity of commerce and that infringing use is required for the booting operation of the **'054 Accused Products**. Any other use would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental.

3060. The booting feature of the **'054 Accused Products** is a material part of the invention of the **'054 Patent** and are especially made or adapted to infringe one or more claims of the **'054 Patent**. Because the use of the booting feature of the **'054 Accused Products**

necessarily infringes one or more claims of the **'054 Patent**, Velocity's sales of its infringing products have no substantial non-infringing uses.

3061. Accordingly, Velocity offers to sell, or sells a component, material, or apparatus for use in practicing one or more claims of the **'054 Patent** knowing the same to be especially made or especially adapted for use in an infringement of such patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

3062. Samsung has no adequate remedy at law for Velocity's infringement of the **'054 Patent** and is suffering irreparable harm, requiring permanent injunctive relief.

EIGHTH CLAIM FOR RELIEF
INFRINGEMENT OF U.S. PATENT NO. 5,777,854
(AGAINST VELOCITY)

3063. Each of the above listed paragraphs is incorporated herein by reference, and adopted, as if fully set forth again.

3064. The **'854 Patent** was filed on May 30, 1997, issued on July 7, 1998, and is entitled "Integrated Flexible Contacts Grounding System For A Computer System Chassis." The **'854 Patent** is generally directed to a chassis for housing electronic devices with a plurality of flexible contacts between a base and cover of the chassis.

3065. The **'854 Patent** was assigned to SEC, and SEC continues to hold all rights, title, and interest in the **'854 Patent**. A true and correct copy of the **'854 Patent** is attached hereto as Exhibit H.

3066. "**Claim 1(a) of the '854 Patent**" recites "[a] chassis for housing electronic devices."

3067. The **'854 Accused Products** include a chassis for housing electronic devices.

3068. For example, the Vector Z25 includes a case that houses a computer, which is an electronic device.

3069. As a result, the **'854 Accused Products** include a chassis for housing electronic devices.

3070. "**Claim 1(b) of the '854 Patent**" recites "a base formed from a base material."

3071. The **'854 Accused Products** have a base formed from a base material.

3072. For example, the Vector Z25 includes a computer case formed of at least two parts, including a base and a cover.

3073. As a result, the **'854 Accused Products** meet the limitations of **Claim 1(b) of the '854 Patent**.

3074. "**Claim 1(c) of the '854 Patent**" recites "a cover formed from a cover material."

3075. The **'854 Accused Products** have a cover formed from a cover material.

3076. For example, the Vector Z25 includes a computer case formed of at least two parts, including a base and a cover.

3077. As a result, the **'854 Accused Products** meet the limitations of **Claim 1(c) of the '854 Patent**.

3078. "**Claim 1(d) of the '854 Patent**" recites "said cover configured to attach to said base, with certain cover regions contacting certain base regions, and wherein said contact between said base and said cover is less than continuous."

3079. The **'854 Accused Products** have a cover configured to attach to a base.

3080. The **'854 Accused Products** also contain certain cover regions contacting certain base regions.

3081. The contact between the cover and base regions of the **'854 Accused Products** is not continuous.

3082. For example, the Vector Z25 includes a cover configured to attach to the base, and the contact between the base and cover is not continuous.

3083. As a result, the **'854 Accused Products** meet the limitations of **Claim 1(d) of the '854 Patent**.

3084. "**Claim 1(e) of the '854 Patent**" recites "a plurality of flexible contacts integrally formed from the material of at least one of said base regions and said cover regions and shaped to contact the other of said base regions and said cover regions."

3085. The **'854 Accused Products** have more than one flexible contacts integrally formed from the material of the base or the cover.

3086. The flexible contacts in the **'854 Accused Products** are shaped to make contact between the base and cover.

3087. The Vector Z25 includes many flexible contacts made of the same material as the base and designed to contact the cover. These flexible contacts are visible in the **Vector Z25 Chassis Image** (red annotations added):



3088. As a result, the **'854 Accused Products** meet the limitations of **Claim 1(e) of the '854 Patent**.

3089. "**Claim 1(f) of the '854 Patent**" recites "said flexible contacts improving electrical contact between said base regions and said cover regions and reducing the size of any gaps that form between the cover regions and base regions, and said flexible contacts having sufficient flexibility to enable more than one contact to contact said other region."

3090. The flexible contacts in the **'854 Accused Products** improve electrical contact between the base and the cover.

3091. The flexible contacts in the **'854 Accused Products** reduce the size of any gaps that form between the base and the cover.

3092. The flexible contacts in the **'854 Accused Products** have sufficient flexibility to enable more than one contact between the base and the cover.

3093. For example, all **'854 Accused Products'** flexible contacts improve electrical contact between the base and cover. The electrical contact improves because the flexible contacts are formed from the base material. The base and cover are formed from an electrically

conductive material. *See, e.g., Velocity Grounding Instructions*: “Common methods of properly grounding yourself are prioritized in order of effectiveness below (1 being most effective, etc.): . . . 4. ‘Discharging’ yourself by touching a piece of metal or the case.”

3094. As a result, the **'854 Accused Products** meet the limitations of **Claim 1(f) of the '854 Patent**.

3095. “**Claim 2 of the '854 Patent**” recites, “The chassis of claim 1, wherein said improved contact between said base regions and said cover regions is an electrical contact for improving electrical grounding of said cover.”

3096. The **'854 Accused Products** also have an electrical contact for improved contact between the base and the cover.

3097. The improved electrical contact in the **'854 Accused Products** improves the electrical grounding of the cover.

3098. For example, the Vector Z25's flexible contacts are on the base, which is typically grounded, and improve electrical grounding of the cover.

3099. As a result, the **'854 Accused Products** meet the limitations of **Claim 2 of the '854 Patent**.

3100. “**Claim 3 of the '854 Patent**” recites, “The chassis of claim 1, wherein said improved contact between said base regions and said cover regions is a mechanical contact for reducing the size of potential gaps.”

3101. The improved contact between the base and the cover in the **'854 Accused Products** reduces the size of potential gaps.

3102. The improved contact between the base and the cover in the **'854 Accused Products** results from a mechanical contact between the base and eth cover.

3103. For example, the Vector Z25's flexible contacts act as a mechanical contact that reduces the size of potential gaps between the base and cover, which do not have continuous contact between them.

3104. As a result, the **'854 Accused Products** meet the limitations of **Claim 3 of the '854 Patent**.

3105. "**Claim 4 of the '854 Patent**" recites, "The chassis of claim 1, wherein said flexible contacts are integrally formed into said base during manufacturing of a base blank."

3106. The **'854 Accused Products** also include flexible contacts are integrally formed as part of the base.

3107. The flexible contacts in the **'854 Accused Products** are formed during the manufacturing of the base.

3108. For example, the Vector Z25's flexible contacts are formed in the base during manufacturing of the base blank, as shown by the continuity of the material of the base and flexible contacts in the **Vector Z25 Chassis Image**.

3109. As a result, the **'854 Accused Products** meet the limitations of **Claim 4 of the '854 Patent**.

3110. "**Claim 5 of the '854 Patent**" recites, "The chassis of claim 1, wherein said electronic devices comprise a computer system."

3111. The **'854 Accused Products** are designed to be used, and are used, to house a computer system.

3112. For example, the Vector Z25 is a desktop computer system. *See* **Configure Your Z25**.

3113. As a result, the **'854 Accused Products** meet the limitations of **Claim 5 of the '854 Patent**.

3114. “**Claim 6 of the '854 Patent**” recites, “The chassis of claim 1, wherein said chassis is a computer system chassis used for housing a computer system comprising a microprocessor.”

3115. The **'854 Accused Products** also are computer system chassis used for housing a computer system.

3116. The computer systems that are housed in the **'854 Accused Products** include a microprocessor.

3117. For example, the Vector Z25 is a desktop personal computer that includes an “Intel® Core i5-4460 Processor, 4-core @ 3.2GHz (3.4GHz Turbo) 6MB L3 Cache, 84 watts” microprocessor *See, e.g., Configure Your Z25*.

3118. The Intel i5 processor in the Vector Z25 computer is a microprocessor.

3119. As a result, the **'854 Accused Products** meet the limitations of **Claim 6 of the '854 Patent**.

3120. “**Claim 15(a) of the '854 Patent**” recites “[a] chassis for housing electronic device.”

3121. The **'854 Accused Products** are chassis for housing electronic devices.

3122. For example, the Vector Z25 includes a case that houses a computer, which is an electronic device.

3123. As a result, the **'854 Accused Products** include a chassis for housing electronic devices.

3124. “**Claim 15(b) of the ’854 Patent**” recites “a base formed from an electrically conductive material, a cover formed from an electrically conductive material.”

3125. The **’854 Accused Products** have a base formed from an electrically conductive material.

3126. The **’854 Accused Products** have a cover formed from an electrically conductive material.

3127. For example, the Vector Z25 includes a computer case formed of at least two parts, including a base formed from an electrically conductive material and a cover formed from an electrically conductive material.

3128. The Vector Z25 includes an electrically conductive base and cover. *See, e.g., Velocity Grounding Instructions*: “Common methods of properly grounding yourself are prioritized in order of effectiveness below (1 being most effective, etc.): . . . 4. ‘Discharging’ yourself by touching a piece of metal or the case.”

3129. As a result, the **’854 Accused Products** meet the limitations of **Claim 15(b) of the ’854 Patent**.

3130. “**Claim 15(c) of the ’854 Patent**” recites “said base having predetermined base regions of intended contact with said cover when said base and said cover are mutually attached, said cover having predetermined cover regions of intended contact with said base when said base and said cover are mutually attached.”

3131. The **’854 Accused Products** have predetermined base regions designed to contact the cover when the base and cover are attached to each other.

3132. The **’854 Accused Products** also have predetermined cover regions designed to contact the base when the base and cover are attached to each other.

3133. For example, the Vector Z25 includes a base and cover intended to attach to one another at particular areas on each. *See, e.g., Vector Z25 Chassis Image* (red annotations added)



3134. As a result, the **'854 Accused Products** meet the limitations of **Claim 15(c) of the '854 Patent**.

3135. "**Claim 15(d) of the '854 Patent**" recites "a plurality of flexible contacts integrally formed from said electrically conductive material of at least one of said base regions and said cover regions and shaped to contact the other of said base regions and said cover regions, said flexible contacts reducing the size of gaps which form between said base regions and said cover regions wherein said flexible contacts extend across said gaps, said contact thereby increasing contact between said base and said cover and reducing the size of said gaps, said flexible contacts having sufficient flexibility to enable more than one contact to contact said other region."

3136. The **'854 Accused Products** include more than one flexible contact integrally formed from the electrically conductive material of the base and cover regions and shaped to contact the other of the base and cover regions.

3137. The **'854 Accused Products** also include flexible contacts that reduce the size of gaps between the base regions and cover regions.

3138. The flexible contact between the base and the cover in the **'854 Accused Products** also increases the contact between the base and cover regions.

3139. The Vector Z25 includes flexible contacts that improve electrical and mechanical contact to the cover. The electrical contact improves because the flexible contacts are formed from the electrically conductive base material. *See, e.g., Velocity Grounding Instructions:* “Common methods of properly grounding yourself are prioritized in order of effectiveness below (1 being most effective, etc.): . . . 4. ‘Discharging’ yourself by touching a piece of metal or the case.”

3140. The flexible contacts discussed in the preceding paragraph are flexible enough that, if one protrudes more than the others towards the cover, it will flex enough to allow another flexible contact that protrudes less to contact the cover.

3141. As a result, the **'854 Accused Products** meet the limitations of **Claim 15(d) of the '854 Patent**.

3142. “**Claim 16(a) of the '854 Patent**” recites “[a] chassis for housing electronic devices.”

3143. The **'854 Accused Products** are chassis for housing electronic devices.

3144. For example, the Vector Z25 includes a case that houses a computer, which is an electronic device.

3145. As a result, the **'854 Accused Products** include a chassis for housing electronic devices.

3146. **Claim 16(b) of the '854 Patent**” recites “a base section; a cover section, said cover section attaching to said base section such that at least a plurality of cover section regions contact a plurality of base section regions.”

3147. The **'854 Accused Products** include a base section and a cover section.

3148. The base section and the cover section in the **'854 Accused Products** make contact with each other in more than one location.

3149. The Vector Z25 includes a base and cover intended to attach to one another at particular areas on each.

3150. As a result, the **'854 Accused Products** meet the limitations of **Claim 16(b) of the '854 Patent**.

3151. “**Claim 16(c) of the '854 Patent**” recites “a plurality of flexible contacts homogeneously formed from at least one of said base section and said cover section, wherein each of said flexible contacts is separated from said at least one of said base section and said cover section by a slot so that said flexible contact may move outwardly from said at least one of said base section and said cover section to contact the other of said base section and said cover section to provide an electrical connection therebetween to provide continuity between EMF shielding provided by said base section and EMF shielding provided by said cover section.”

3152. The **'854 Accused Products** include flexible contacts formed from the base or cover sections.

3153. Each of the flexible contacts of the **'854 Accused Products** is separated from the base section and the cover section by a slot so that the flexible contact may move outwardly from

the base section and the cover section to contact the other base section and cover section to provide an electrical connection therebetween.

3154. The electrical connection in the **'854 Accused Products** is to provide continuity between EMF shielding provided by the base section and EMF shielding provided by the cover section.

3155. The Vector Z25 includes electrically conductive flexible contacts formed in the base and separated from it by a slot that allows them to protrude to contact the cover. These flexible contacts and slots are visible in the **Vector Z25 Chassis Image** (red annotations added):



3156. The electrical contact provides continuity between EMF shielding provided by the base and cover.

3157. As a result, the **'854 Accused Products** meet the limitations of **Claim 16(c) of the '854 Patent**.

3158. **"Claim 17 of the '854 Patent"** recites, "A chassis as defined in claim 16, wherein said flexible contact has a protrusion partially traversing the width of the flexible contact, said

protrusion further enhancing contact between said flexible contact and the other of said base section and said cover section.”

3159. The flexible contact in the **'854 Accused Products** has a protrusion partially traversing the width of the flexible contact.

3160. The protrusion in the flexible contact in the **'854 Accused Products** enhances contact between the flexible contact and the other of the base section and the cover section.

3161. The Vector Z25 includes electrically conductive flexible contacts that include a protrusion that partially traverses the width of the flexible contact.

3162. As a result, the **'854 Accused Products** meet the limitations of **Claim 17 of the '854 Patent**.

3163. “**Claim 19(a) of the '854 Patent**” recites “[a] chassis for housing electronic devices.”

3164. The **'854 Accused Products** are chassis for housing electronic devices.

3165. For example, the Vector Z25 includes a case that houses a computer, which is an electronic device.

3166. As a result, the **'854 Accused Products** include a chassis for housing electronic devices.

3167. “**Claim 19(b) of the '854 Patent**” recites “a base; a cover.”

3168. The **'854 Accused Products** include a base and a cover.

3169. For example, the Vector Z25 includes a base and cover intended to attach to one another at particular areas on each.

3170. As a result, the **'854 Accused Products** meet the limitations of **Claim 19(b) of the '854 Patent**.

3171. “**Claim 19(c) of the ’854 Patent**” recites “a plurality of flexible contacts formed as part of an original material from which at least one of said base and said cover portion is formed.”

3172. The **’854 Accused Products** include multiple flexible contacts formed as part of an original material of the base or cover.

3173. The flexible contacts in the **’854 Accused Products** are formed as part of the original material from which either the base or the cover is formed.

3174. The Vector Z25 includes many flexible contacts made of the same material as the base and designed to contact the cover. These flexible contacts are visible in the annotated **Vector Z25 Chassis Image**, depicted above.

3175. As a result, the **’854 Accused Products** meet the limitations of **Claim 19(c) of the ’854 Patent**.

3176. “**Claim 19(d) of the ’854 Patent**” recites “a first end of each contact being a continuum of said original material, and a second end of each contact moving with respect to said cover portion, each said contact flexing between said respective first end and said respective second end.”

3177. The **’854 Accused Products** include a first end of each contact being a continuum of the original material.

3178. The **’854 Accused Products** also include a second end of each contact moving with respect to the cover portion.

3179. Each contact in the **’854 Accused Products** flexes between the above described first and second ends.

3180. For example, the Vector Z25's flexible contacts each have one end that is a continuum of the original material that remains fixed, and have another end that flexes to move in a path nearly perpendicular to the plane of the base. *See* **Vector Z25 Chassis Image**.

3181. As a result, the **'854 Accused Products** meet the limitations of **Claim 19(d) of the '854 Patent**.

3182. Velocity has had actual knowledge of the **'854 Patent** at least as early as as of November 11, 2014, the date the original complaint in this action was served.

3183. Velocity has directly infringed and continues to directly infringe one or more of the claims of the **'854 Patent** by making, using, offering to sell, or selling, the **'854 Accused Products**.

3184. Velocity makes the **'854 Accused Products**. For example, Velocity makes the Velocity Vector Z25. *See* **Configure Your Z25** (describing the Z25 as "assembled in Richmond, Virginia, USA").

3185. Velocity uses the **'854 Accused Products**.

3186. For example, Velocity uses the Velocity Z25 for demonstration, testing, and development purposes.

3187. Velocity sells and offers to sell the **'854 Accused Products**.

3188. For example, Velocity sells the Velocity Z25 online. *See* **Configure Your Z25**.

3189. Velocity indirectly infringes the **'854 Patent** by inducing infringement by others, such as distributors, resellers, customers, and end users, in accordance with 35 U.S.C. § 271(b), in this District and elsewhere in the United States. Direct infringement is the result of activities performed by others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, and add-in board

manufacturers, resellers, customers, and end users, by, for example, using, selling, or offering for sale the **'854 Accused Products**.

3190. Velocity has induced and continues to induce infringement of the **'854 Patent** by intending that others infringe the **'854 Patent** by using, selling, or offering for sale the **'854 Accused Products**. Velocity designed the **'854 Accused Products** such that they would each infringe one or more claims of the **'854 Patent**.

3191. Velocity provides the **'854 Accused Products** to others, such as distributors, resellers, customers, and end users.

3192. For example, Velocity products have been sold at nearly every major electronics retailer in North America including Best Buy, Circuit City, RadioShack, Costco, Sears, Target, and many others. *See, e.g., Velocity Company History*.

3193. By providing **'854 Accused Products** to others, Velocity intends for **'854 Accused Products** to be used in the United States.

3194. Velocity also provides others, such as distributors, resellers, customers, and end users, instructions, user guides, and technical specifications. When others follow such instructions, user guides, and/or other design documentation, they directly infringe one or more claims of the **'854 Patent**. Velocity knows that by providing such instructions, user guides, and/or other design documentation, others, such as distributors, original equipment manufacturers, contract equipment manufacturers, original design manufacturers, motherboard manufacturers, add-in board manufacturers, resellers, customers, and end users, follow those instructions, user guides, and other design documentation, and directly infringe one or more claims of the **'854 Patent**. Velocity thus knows that its actions actively induce infringement.

3195. Samsung has no adequate remedy at law for Velocity's infringement of the '**854 Patent** and is suffering irreparable harm, requiring permanent injunctive relief.

NINTH CLAIM FOR RELIEF
FALSE ADVERTISING (AGAINST NVIDIA)
(Va. Code §§ 18.2-216 and 59.1-68.3)

3196. Each of the above listed paragraphs is incorporated herein by reference, and adopted, as if fully set forth again.

3197. Samsung sells in the United States and in this District, among other devices, mobile computing devices that use the Android operating system.

3198. NVIDIA sells in the United States and in this District, among other devices, mobile computing devices under the Shield brand. These devices further include the NVIDIA Shield Tablet. The NVIDIA Shield Tablet also uses the Android operating system and competes directly with Samsung's mobile computing devices.

3199. NVIDIA advertises through its website, with intent to sell to the public or otherwise increase public consumption of the NVIDIA Shield Tablet, that the NVIDIA Shield Tablet has the "world's fastest mobile processor." *See NVIDIA Shield*. The NVIDIA Shield Tablet contains an NVIDIA Tegra K1 SOC.

3200. NVIDIA's claim that the NVIDIA Shield Tablet has the world's fastest mobile processor is a false and misleading statement of fact. Rather, standard benchmarking tools such as Primate Labs' Geekbench 3 benchmark reveal that the NVIDIA Tegra K1 SOC is not the world's fastest mobile processor. The Samsung Exynos 5433 SOC, as used in the Samsung Galaxy Note 4 mobile computing device, scores higher in both the single-core and multi-core Geekbench 3 benchmark tests than the NVIDIA Tegra K1 SOC used in the NVIDIA Shield Tablet. Moreover, the Samsung Exynos 5433 SOC is not the only SOC that is faster than the

NVIDIA Tegra K1 SOC. The Apple A8X SOC, used in the Apple iPad Air 2 tablet, also scores higher in both the single-core and multi-core Geekbench 3 benchmark tests than the NVIDIA Tegra K1 SOC.

3201. NVIDIA's claim is thus false, defamatory, deceptive, and likely to confuse consumers. Furthermore, this claim is designed to influence consumers to not do business with Samsung and to do business with NVIDIA instead.

3202. NVIDIA's actions described above constitute false and misleading advertising in violation of Virginia Code § 18.2-216. Samsung has suffered and is continuing to suffer monetary, economic, and reputational losses as a result thereof and seeks damages under Va. Code § 59.1-68.3.

3203. NVIDIA's actions, as described above, are intended to wrongfully divert sales and business from Samsung to NVIDIA.

3204. NVIDIA's advertisement, as described above, contains a false representation which is untrue, deceptive, and misleading.

3205. NVIDIA has made and distributed this advertisement via its website to consumers in Virginia and in this District.

3206. NVIDIA's conduct has caused damage to Samsung's business, reputation, and goodwill, and Samsung seeks monetary damages in an amount to be determined at trial. Samsung is further entitled to recover attorneys' fees under Virginia Code § 59.1-68.3.

JURY DEMAND

3207. Samsung demands a trial by jury on all issues.

PRAYER FOR RELIEF

WHEREFORE, Samsung respectfully requests entry of a judgment in its favor against Defendants as follows:

- i) A declaration that Defendants have directly and/or indirectly infringed the '158, '938, '902, '602, '675, '724, '054, and '854 Patents;
- ii) An award of damages arising out of Defendants' infringement of the **Asserted Patents** sufficient to compensate Samsung, including enhanced damages pursuant to 35 U.S.C. § 284, as well as any prejudgment and post-judgment interest, in an amount according to proof, such amount not including any damages for:
 - a) NVIDIA's indirect infringement of the '602, '675, and '724 Patents that occurred prior to November 12, 2014, and
 - b) Velocity's indirect infringement of the '158, '938, '902, '602, '675, '724, '054, and '854 Patents that occurred prior to November 11, 2014;
- iii) An award of enhanced damages pursuant to 35 U.S.C. § 284 for NVIDIA's willful infringement of the '158, '938, and '902 Patents;
- iv) An Order permanently enjoining Defendants and their respective officers, agents, employees, and those acting in privity with them, from further infringement, including contributory infringement and/or inducing infringement, of the '158, '902, '675, '724, '054, and '854 Patents;
- v) An award of damages arising out of NVIDIA's false advertising in violation of Va. Code § 18.2-216, pursuant to Va. Code § 59.1-68.3;
- vi) An award of attorneys' fees pursuant to 35 U.S.C. § 285, Virginia Code § 59.1-68.3, and as otherwise permitted by law; and

vii) An award of any such other costs and further relief as the Court may deem just and proper.

Respectfully submitted,

SAMSUNG ELECTRONICS CO., LTD.
SAMSUNG ELECTRONICS AMERICA, INC.
By Counsel

/s/ Robert W. McFarland

Robert W. McFarland (VSB No. 24021)
Sarah K. McConaughy (VSB No. 80674)
McGuireWoods LLP
101 W. Main Street, Suite 9000
Norfolk, Virginia 23510
Telephone: (757) 640-3716
Facsimile: (757) 640-3966
E-mail: rmcfarland@mcguirewoods.com
E-mail: smcconaughey@mcguirewoods.com

Brian C. Riopelle (VSB No. 36454)
McGuireWoods LLP
One James Center 901 E. Cary Street
Richmond, VA 23219
Telephone: (804) 775-1084
Facsimile: (804) 698-2150
E-mail: briopelle@mcguirewoods.com

Darin W. Snyder (Pro Hac Vice)
Alexander B. Parker (Pro Hac Vice)
Elysa Q. Wan (Pro Hac Vice)
O'Melveny & Myers LLP
Two Embarcadero Center, 28th Floor
San Francisco, CA 94111
E-mail: dsnyder@omm.com
E-mail: aparker@omm.com
E-mail: ewan@omm.com

Vision L. Winter (Pro Hac Vice)
Ryan K. Yagura (Pro Hac Vice)
Michael A. Koplow (Pro Hac Vice)
O'Melveny & Myers LLP
400 South Hope Street, 18th Floor
Los Angeles, CA 90071
E-mail: vwinter@omm.com
E-mail: ryagura@omm.com
E-mail: mkoplow@omm.com

Los Angeles, CA 90071

E-mail: ryagura@omm.com

E-mail: vwinter@omm.com

*Attorneys for PLAINTIFFS Samsung Electronics
Co., Ltd. and Samsung Electronics America, Inc.*

CERTIFICATE OF SERVICE

I hereby certify that on April 10, 2015, a true and correct copy of the foregoing was filed electronically using the CM/ECF system. As such, this document was served on all counsel who have consented to electronic service, including as follows:

Dabney Jefferson Carr, IV
Troutman Sanders LLP
1001 Haxall Point
Richmond, VA 23219
Telephone: 804-697-1238
Facsimile: 804-698-5119
Email: Dabney.carr@troutmansanders.com

Clement Joseph Naples
Latham & Watkins LLP
885 Third Avenue, 25th Floor
New York, NY 10022
Telephone: 212-906-1200
Facsimile: 212-906-1201
Email: clement.naples@lw.com

*Counsel for NVIDIA Corporation,
Velocity Micro, Ind., d/b/a Velocity Micro
and Velocity Holdings, LLC*

/s/ Robert W. McFarland
Robert W. McFarland (VSB No. 24021)
Sarah K. McConaughy (VSB No. 80674)
MCGUIREWOODS LLP
101 West Main Street, Suite 9000
Norfolk, Virginia 23510
Telephone: (757) 640-3716
Facsimile: (757) 640-3966
E-mail: rmcfarland@mcguirewoods.com

*Attorneys for PLAINTIFFS Samsung Electronics
Co., Ltd., and Samsung Electronics America, Inc.*