

**UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

ALTAIR LOGIX LLC,

Plaintiff,

v.

MICROSOFT CORPORATION,

Defendant.

CASE NO. 2:18-cv-325

JURY TRIAL DEMANDED

PATENT CASE

**ORIGINAL COMPLAINT FOR PATENT INFRINGEMENT
AGAINST MICROSOFT CORPORATION**

Plaintiff Altair Logix LLC files this Original Complaint for Patent Infringement against Microsoft Corporation, and would respectfully show the Court as follows:

I. THE PARTIES

1. Plaintiff Altair Logix LLC (“Altair Logix” or “Plaintiff”) is a Texas limited liability company with its principal place of business at 15922 Eldorado Pkwy, Suite 500 #1513, Frisco, TX 75035.

2. On information and belief, Defendant Microsoft Corporation (“Defendant”) is a corporation organized and existing under the laws of Washington, with a place of business at 2601 Preston Rd. #1176, Frisco, Texas 75034, in Collin County, Texas.

II. JURISDICTION AND VENUE

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

4. On information and belief, Defendant is subject to this Court’s specific and general personal jurisdiction, pursuant to due process and the Texas Long-Arm Statute, due at least to its

business in this forum, including at least a portion of the infringements alleged herein. Furthermore, Defendant is subject to this Court's specific and general personal jurisdiction because it has a place of business within this District, including at 2601 Preston Rd. #1176, Frisco, Texas 75034, in Collin County, Texas.

5. Without limitation, on information and belief, within this District and state, Defendant has used the patented inventions thereby committing, and continuing to commit, acts of patent infringement alleged herein. In addition, on information and belief, Defendant has derived revenues from its infringing acts occurring within the Eastern District of Texas and Texas. Further, on information and belief, Defendant is subject to the Court's general jurisdiction, including from regularly doing or soliciting business, engaging in other persistent courses of conduct, and deriving substantial revenue from goods and services provided to persons or entities in the Eastern District of Texas and Texas. Further, on information and belief, Defendant is subject to the Court's personal jurisdiction at least due to its sale of products and/or services within the Eastern District of Texas. Defendant has committed such purposeful acts and/or transactions in the Eastern District of Texas and Texas such that it reasonably should know and expect that it could be haled into this Court as a consequence of such activity.

6. Venue is proper in this district under 28 U.S.C. § 1400(b). On information and belief, Defendant has a place of business at 2601 Preston Rd. #1176, Frisco, Texas 75034, in Collin County, Texas. On information and belief, from and within this District Defendant has committed at least a portion of the infringements at issue in this case.

7. For these reasons, personal jurisdiction exists and venue is proper in this Court under 28 U.S.C. § 1400(b).

III. COUNT I
(PATENT INFRINGEMENT OF UNITED STATES PATENT NO. 6,289,434)

8. Plaintiff incorporates the above paragraphs herein by reference.

9. On September 11, 2001, United States Patent No. 6,289,434 (“the ‘434 Patent”) was duly and legally issued by the United States Patent and Trademark Office. The application leading to the ‘434 patent was filed on February 27, 1998. (Ex. A at cover).

10. The ‘434 Patent is titled “Apparatus and Method of Implementing Systems on Silicon Using Dynamic-Adaptive Run-Time Reconfigurable Circuits for Processing Multiple, Independent Data and Control Streams of Varying Rates.” A true and correct copy of the ‘434 Patent is attached hereto as Exhibit A and incorporated herein by reference.

11. Plaintiff is the assignee of all right, title and interest in the ‘434 patent, including all rights to enforce and prosecute actions for infringement and to collect damages for all relevant times against infringers of the ‘434 Patent. Accordingly, Plaintiff possesses the exclusive right and standing to prosecute the present action for infringement of the ‘434 Patent by Defendant.

12. The invention in the ‘434 Patent relates to the field of runtime reconfigurable dynamic-adaptive digital circuits which can implement a myriad of digital processing functions related to systems control, digital signal processing, communications, image processing, speech and voice recognition or synthesis, three-dimensional graphics rendering, and video processing. (Ex. A at col. 1:32-38). The object of the invention is to provide a new method and apparatus for implementing systems on silicon or other chip material which will enable the user a means for achieving the performance of fixed-function implementations at a lower cost. (*Id.* at col. 2:64 – col. 3:1).

13. The most common method of implementing various functions on an integrated circuit is by specifically designing the function or functions to be performed by placing on silicon

an interconnected group of digital circuits in a non-modifiable manner (hard-wired or fixed function implementation). (*Id.* at col. 1:42-47). These circuits are designed to provide the fastest possible operation of the circuit in the least amount of silicon area. (*Id.* at col. 1:47-49). In general, these circuits are made up of an interconnection of various amounts of random-access memory and logic circuits. (*Id.* at col. 1:49-51). Complex systems on silicon are broken up into separate blocks and each block is designed separately to only perform the function that it was intended to do. (*Id.* at col. 1:51-54). Each block has to be individually tested and validated, and then the whole system has to be tested to make sure that the constituent parts work together. (*Id.* at col. 1:54-56). This process is becoming increasingly complex as we move into future generations of single-chip system implementations. (*Id.* at col. 1:57-59). Systems implemented in this way generally tend to be the highest performing systems since each block in the system has been individually tuned to provide the expected level of performance. (*Id.* at col. 1:59-62). This method of implementation may be the smallest (cheapest in terms of silicon area) method when compared to three other distinct ways of implementing such systems. (*Id.* at col. 1:62-65). Each of the other three have their problems and generally do not tend to be the most cost-effective solution. (*Id.* at col. 1:65-67).

14. The first way is implemented in software using a microprocessor and associated computing system, which can be used to functionally implement any system. (*Id.* at col. 2:1-2). However, such systems would not be able to deliver real-time performance in a cost-effective manner for the class of applications that was described above. (*Id.* at col. 2:3-5). Their use is best for modeling the subsequent hard-wired/fixed-function system before considerable design effort is put into the system design. (*Id.* at col. 2:5-8).

15. The second way of implementing such systems is by using an ordinary digital signal processor (DSP). (*Id.* at col. 2:9-10). This class of computing machines is useful for real-time processing of certain speech, audio, video and image processing problems and in certain control functions. (*Id.* at col. 2:10-13). However, they are not cost-effective when it comes to performing certain real time tasks which do not have a high degree of parallelism in them or tasks that require multiple parallel threads of operation such as three-dimensional graphics. (*Id.* at col. 2:13-17).

16. The third way of implementing such systems is by using field programmable gate arrays (FPGA). (*Id.* at col. 2:18-19). These devices are made up of a two-dimensional array of fine grained logic and storage elements which can be connected together in the field by downloading a configuration stream which essentially routes signals between these elements. (*Id.* at col. 2:19-23). This routing of the data is performed by pass-transistor logic. (*Id.* at col. 2:24-25). FPGAs are by far the most flexible of the three methods mentioned. (*Id.* at col. 2:25-26). The problem with trying to implement complex real-time systems with FPGAs is that although there is a greater flexibility for optimizing the silicon usage in such devices, the designer has to trade it off for increase in cost and decrease in performance. (*Id.* at col. 2:26-30). The performance may (in some cases) be increased considerably at a significant cost, but still would not match the performance of hard-wired fixed function devices. (*Id.* at col. 2:30-33).

17. These three ways do not reduce the cost or increase the performance over fixed-function systems. (*Id.* at col. 2:35-37). In terms of performance, fixed-function systems still outperform the three ways for the same cost. (*Id.* at col. 2:37-39).

18. The three systems can theoretically reduce cost by removing redundancy from the system. (*Id.* at col. 2:40-41). Redundancy is removed by re-using computational blocks and memory. (*Id.* at col. 2:41-42). The only problem is that these systems themselves are increasingly

complex, and therefore, their computational density when compared with fixed-function devices is very high. (*Id.* at col. 2:42-45).

19. Most systems on silicon are built up of complex blocks of functions that have varying data bandwidth and computational requirements. (*Id.* at col. 2:46-48). As data and control information moves through the system, the processing bandwidth varies enormously. (*Id.* at col. 2:48-50). Regardless of the fact that the bandwidth varies, fixed-function systems have logic blocks that exhibit a “temporal redundancy” that can be exploited to drastically reduce the cost of the system. (*Id.* at col. 2:50-53). This is true, because in fixed function implementations all possible functional requirements of the necessary data processing must be implemented on the silicon regardless of the final application of the device or the nature of the data to be processed. (*Id.* at col. 2:53-57). Therefore, if a fixed function device must adaptively process data, then it must commit silicon resources to process all possible flavors of the data. (*Id.* at col. 2:58-60). Furthermore, state-variable storage in all fixed function systems are implemented using area inefficient storage elements such as latches and flip-flops. (*Id.* at col. 2:60-63).

20. The inventors therefore sought to provide a new apparatus for implementing systems on a chip that will enable the user to achieve performance of fixed-function implementation at a lower cost. (*Id.* at col. 2:64 – col. 3:1). The lower cost is achieved by removing redundancy from the system. (*Id.* at col. 3:1-2). The redundancy is removed by re-using groups of computational and storage elements in different configurations. (*Id.* at col. 3:2-4). The cost is further reduced by employing only static or dynamic ram as a means for holding the state of the system. (*Id.* at col. 3:4-6). This invention provides a way for effectively adapting the configuration of the circuit to varying input data and processing requirements. (*Id.* at col. 3:6-8).

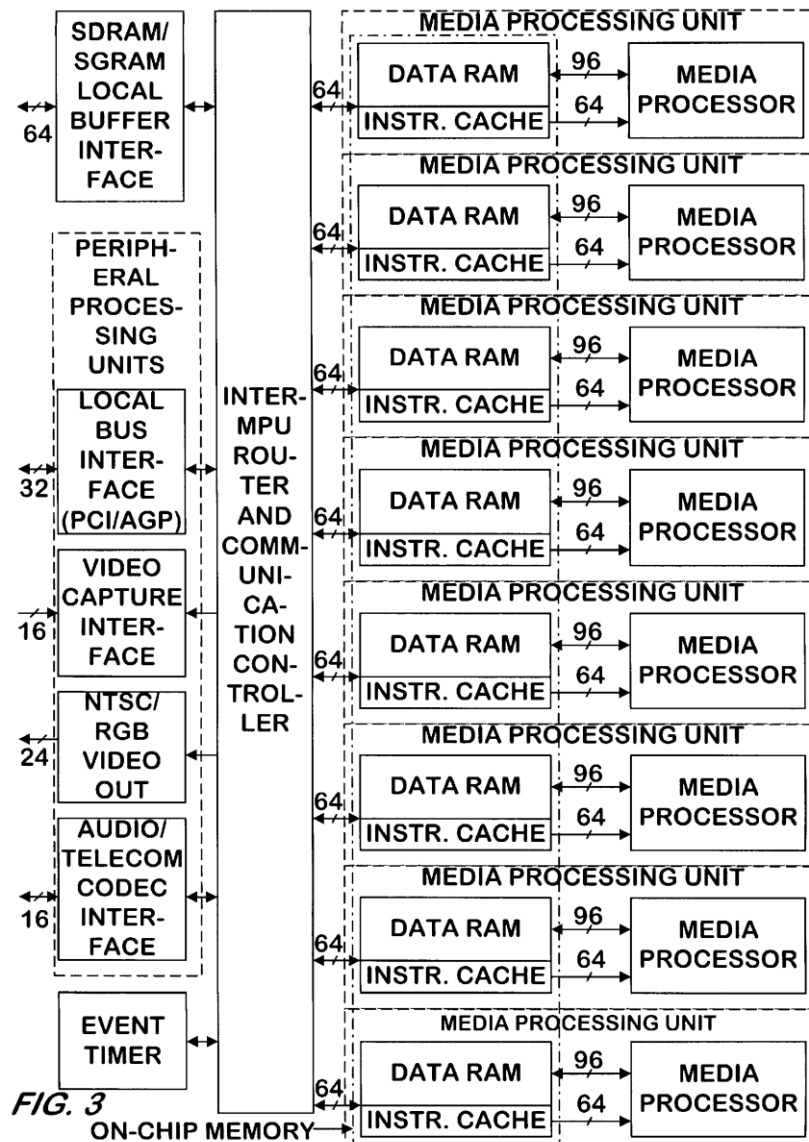
All of this reconfiguration can take place dynamically in run-time without any degradation of performance over fixed-function implementations. (*Id.* at col. 3:8-11).

21. The present invention is therefore an apparatus for adaptively dynamically reconfiguring groups of computations and storage elements in run-time to process multiple separate streams of data and control at varying rates. (*Id.* at col. 3:14-18). The '434 patent refers to the aggregate of the dynamically reconfigurable computational and storage elements as a "media processing unit."

22. The claimed apparatus has addressable memory for storing data and a plurality of instructions that can be provided through a plurality of inputs/outputs that is couple to the input/output of a plurality of media processing units. (*Id.* at col. 55:21-30). The media processing unit comprises a multiplier, an arithmetic unit, and arithmetic logic unit and a bit manipulation unit. (*Id.* at col. 55:31 – col. 56:20). The '434 patent provides examples to explain each of the parts of the media processing unit. (*Id.* at col. 16:27-61 (multiplier and adder); *id.* at col. 16:62 – col. 17:1-9 (arithmetic logic unit); and *id.* at col. 17:10 – col. 17:43 (bit manipulation unit)). Each of the parts has a data input coupled to the media processing unit input/output, an instruction input coupled to the mediate processing unit input/output, and a data output coupled to the mediate processing unit input/output. (*Id.* at col. 55:31 – col. 56:20). Furthermore, the arithmetic logic unit must be capable of operating concurrently with either the multiplier and arithmetic unit. (*Id.* at col. 56:6-12). And the bit manipulation unit must be capable of operating concurrently with the arithmetic logic unit and at least either the multiplier or the arithmetic unit. (*Id.* at col. 56:13-20). Each of the plurality of media processing units must be capable of performing an operating simultaneously with the performance of other operations by other media processing units. (*Id.* at col. 56:21-24). An operation comprises the media processing unit receiving an instruction and

data from memory, processing the data responsive to the instruction to produce a result, and providing the result to the media processor input/output. (*Id.* at col. 56:26-33).

23. An exemplary block diagram of the claimed systems is shown in Figure 3 of the '434 patent:



(*Id.* at Fig. 3). Exemplary architecture and coding for the apparatus is disclosed in the '599 patent. (*E.g., id.* at col. 16:15 – col. 52:20; Figs. 9 – 106).

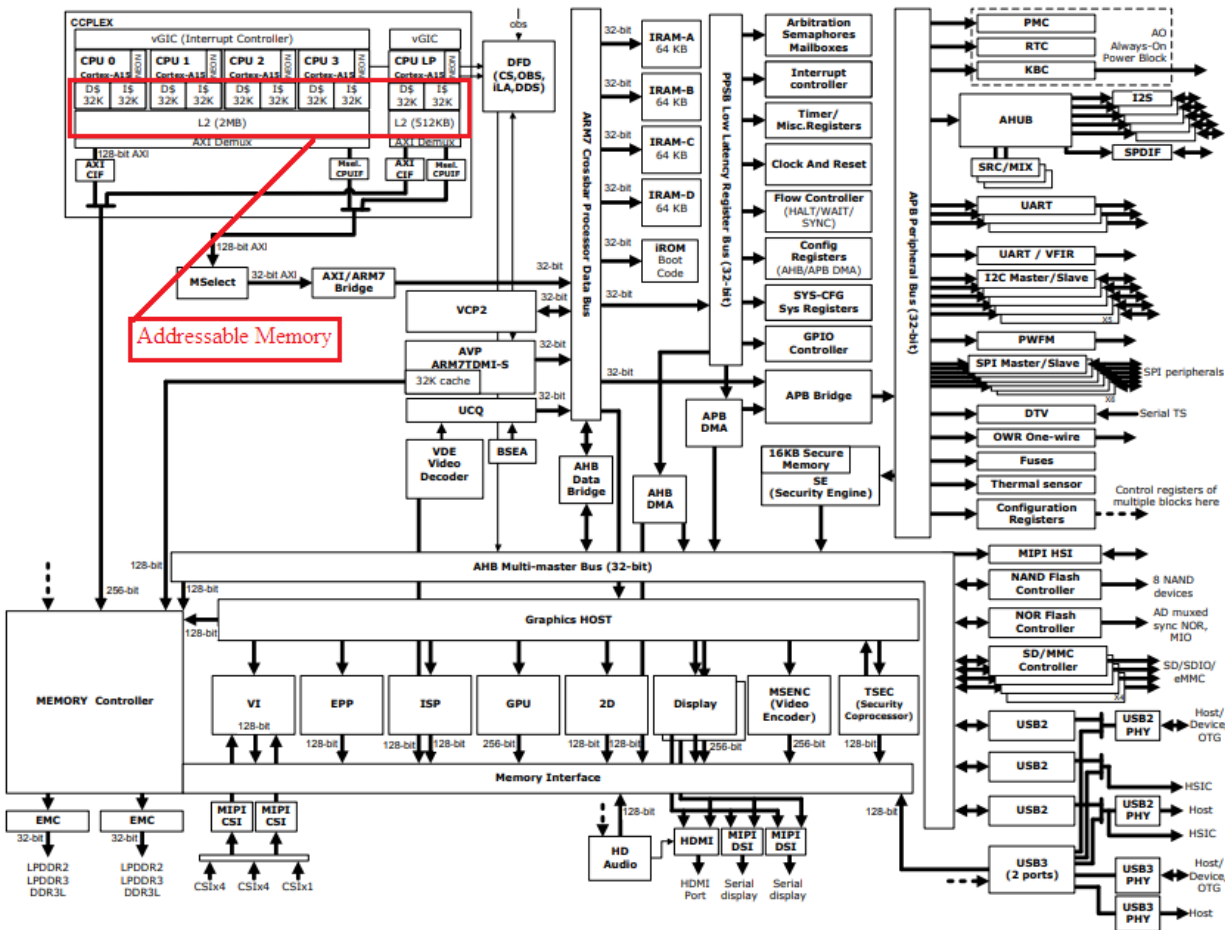
24. As further demonstrated by the prosecution history of the '434 patent, the claimed invention in the '434 patent was unconventional. Claim 1 of the '434 patent was an originally filed claim that issued without any amendment. There was no rejection in the prosecution history contending that claim 1 was anticipated by any prior art.

25. A key element behind the invention is one of reconfigurability and reusability. (*Id.* at col. 13:26-27). Each apparatus is therefore made up of very high-speed core elements that on a pipelined basis can be configured to form a more complex function. (*Id.* at col. 13:27-30). This leads to a lower gate count, thereby giving a smaller die size and ultimately a lower cost. (*Id.* at col. 13:30-31). Since the apparatuses are virtually identical to each other, writing software becomes very easy. (*Id.* at col. 13:32-33). The RISC-like nature of each of the media processing units also allows for a consistent hardware platform for simple operating system and driver development. (*Id.* at col. 13:33-36). Any one of the media processing units can take on a supervisory role and act as a central controller if necessary. (*Id.* at col. 13:36-37). This can be very useful in set top applications where a controlling CPU may not be necessary, further reducing system cost. (*Id.* at col. 13:37-40). The claimed apparatus is therefore an unconventional way of implementing processors that can achieve the performance of fixed-function implementations at a lower cost. (*Id.* at col. 2:64 – col. 3:11).

26. **Direct Infringement.** Upon information and belief, Defendant has been directly infringing claims of the '434 patent in the Eastern District of Texas and Texas, and elsewhere in the United States, by making, using, selling, and offering for sale an apparatus for processing data for media processing that satisfies each and every limitation of at least claim 1, including without limitation the Microsoft Surface 2 (“Accused Instrumentality”). (*E.g.*, <https://www.amazon.com/Microsoft-Surface-2-32-GB/dp/B00FF6J532>).

27. The Accused Instrumentality comprises a processing unit (*e.g.*, Nvidia Tegra 4) which has multiple media processing units (*e.g.*, ARM Quad core Cortex-A15). (*E.g.*, <http://www.nvidia.com/object/tegra-4-processor.html>;
https://www.nvidia.com/docs/IO/116757/NVIDIA_Quad_a15_whitepaper_FINALv2.pdf;
https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxIV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTCld2Tuc31pIRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17). The Accused Instrumentality comprises an addressable memory (*e.g.*, memory system of the Accused Instrumentality) for storing the data, and a plurality of instructions, and having a plurality of input/outputs, each said input/output for providing and receiving at least one selected from the data and the instructions. As shown below, the Accused Instrumentality comprises a memory system which is coupled to multicore ARM processors through multiple internal inputs/outputs. The memory system provides instructions and stored data for processing and receives processed data.

Figure 1: Tegra 4 Series Processor Block Diagram



https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxIV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTCId2Tuc31pIRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17

Caches are used on CPUs to reduce the number of off-chip accesses to system memory. Caches store the most frequently used data on-chip enabling the CPU to access the data faster and improving the performance and efficiency of the CPU. Each core of the quad core ARM Cortex-A15 CPU complex on NVIDIA Tegra 4 has its own 32KB Instruction cache and 32KB of Data cache. All four cores share a common large 2MB L2 cache, which is 16-way set associative. The large 128 entry deep Out-of-order buffer allows the L2 cache latency to be largely hidden. Along with the 32KB L1 Caches, the 2MB L2 cache works to minimize off-chip fetches to system memory, both increasing performance and reducing power as DRAM fetches are more power intensive than on-chip SRAM fetches.

https://www.nvidia.com/docs/IO/116757/NVIDIA_Quad_a15_whitepaper_FINALv2.pdf).

The NVIDIA® Tegra® 4 series processor is a complete applications and digital media system built around several powerful hardware elements:

- CPU Complex: Quad Cortex™-A15 Symmetric Multi-Processing ARM® Cores in a 4-PLUS-1™ configuration with a quad-core fast CPU complex and a fifth Battery Saver Core. The Cortex-A15 core features triple instruction issue and both out-of-order and speculative execution. It has full cache coherency support for the quad symmetric processors. All processors have 32 KB Instruction and 32 KB Data Level 1 caches; and there is a 2 MB shared Level 2 cache for the quad-core complex and a 512 KB Level 2 cache for the fifth core. The NVIDIA 4-PLUS-1 architecture uses the fifth Battery Saver Core, which operates exclusively with the main CPU complex, for very low-power, low-leakage operation at the light CPU loads common to multimedia and lightly loaded use situations.
- Memory Controller: dual-channel (2x 32-bit) DRAM interface providing more than twice the available bandwidth of Tegra 3 devices. LP-DDR2, LP-DDR3 and DDR3 DRAM types are all supported.

https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxlV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTClD2Tuc31plRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17).

28. The Accused Instrumentality comprises a plurality of media processing units (*e.g.*, ARM cortex-A15 multicore processors), each media processing unit having an input/output coupled to at least one of the addressable memory input/outputs. As shown below, the Accused Instrumentality comprises ARM cortex-A15 multicore processors, each processor comprises a NEON media coprocessor and acts as a media processing unit. The ARM processors are coupled to the memory system. The processors receive instructions and data from the memory system by multiple internal inputs and provides processed data to the memory system by multiple internal outputs.

INTRODUCING NVIDIA® TEGRA® 4, THE WORLD'S FASTEST MOBILE PROCESSOR

Take your mobile experience further than ever before with Tegra 4's perfect balance of power and efficiency. Explore new camera capabilities, lightning-fast web browsing, stunning visuals for mobile gaming, and blazing LTE networking through an optional chipset—all with incredible battery life.

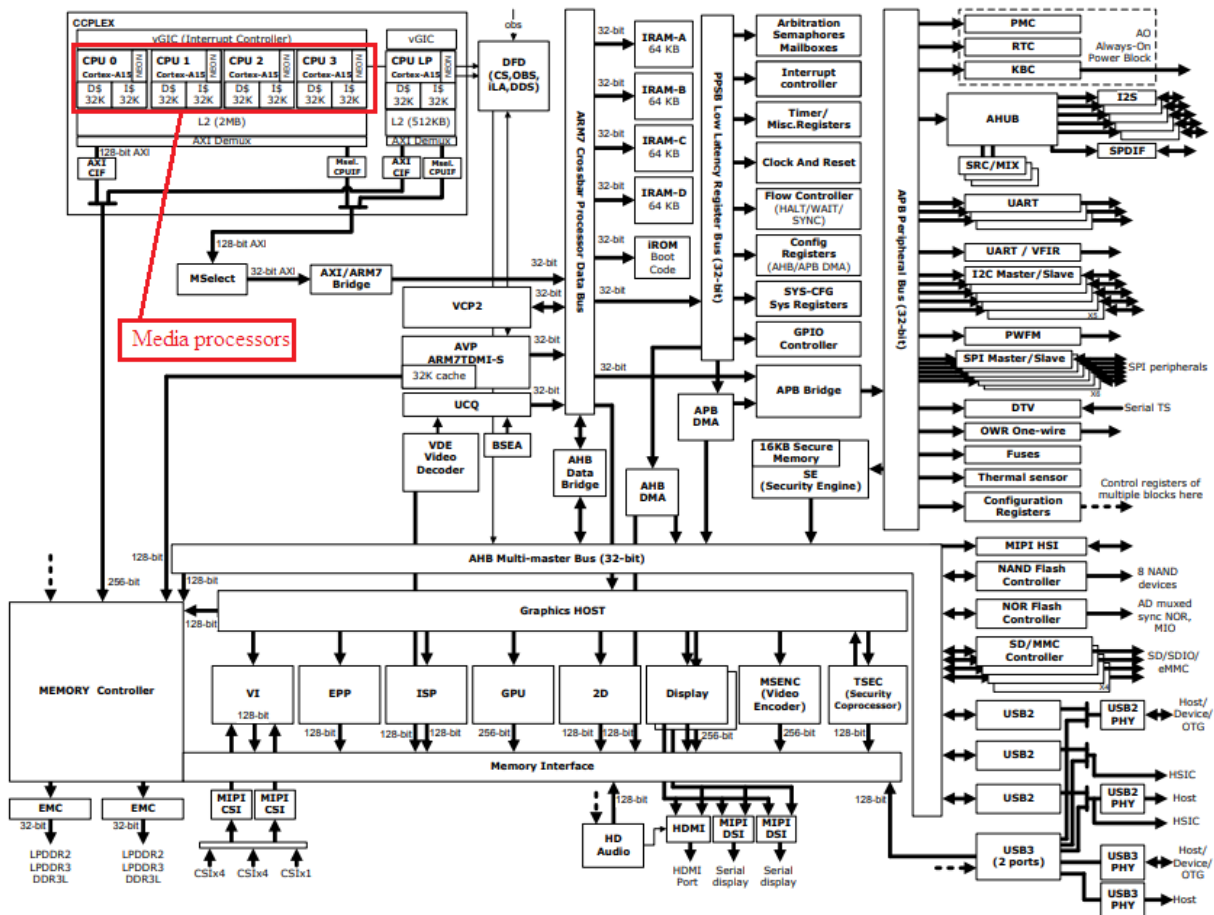


KEY INNOVATIONS

- > NVIDIA® GPU with up to 72 custom cores - Enjoy unique mobile device innovations in photography, media, gaming, and web—including High Dynamic Range (HDR) imaging, WebGL, and HTML5.
- > Quad-Core ARM Processor - NVIDIA Tegra 4 processor harnesses ARM's most advanced CPU cores ever, plus a second-generation battery-saver core, to deliver record levels of performance and battery life. The ARM Cortex-A15 CPU is the engine behind Tegra 4, while Tegra 4i is powered by the new ARM Cortex-A9 r4 CPU—which was defined by ARM with help from NVIDIA - and the most efficient CPU core in its class.

(<http://www.nvidia.com/object/tegra-4-processor.html>).

Figure 1: Tegra 4 Series Processor Block Diagram



https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxlV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTCld2Tuc31plRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17).

18.0 CPU

The NVIDIA® Tegra® 4 series processor CPU complex contains quad ARM® Cortex™-A15 CPUs in a 4-PLUS-1 configuration with a fifth architecturally identical power-saving Cortex-A15 Companion Core.

18.1 Cortex-A15 CPU

Cortex-A15 is an advanced processor design with many features for high instruction throughput. It integrates the L2 cache controller into the CPU complex unlike Cortex-A9. All of the CPUs include the NEON Media Processing Engine. Further details of the Cortex-A15 itself are available from ARM.

These two documents are the key references on Cortex-A15, and both are available from ARM's website:

- Cortex-A15
Revision: r2p1
Technical Reference Manual

Published by ARM Limited, document number ARM DDI 0438D.
- ARM Architecture Reference Manual
ARM v7-A and ARM v7-R edition

Published by ARM Limited, document number ARM DDI 0406C.

(*Id.*).

Caches are used on CPUs to reduce the number of off-chip accesses to system memory. Caches store the most frequently used data on-chip enabling the CPU to access the data faster and improving the performance and efficiency of the CPU. Each core of the quad core ARM Cortex-A15 CPU complex on NVIDIA Tegra 4 has its own 32KB Instruction cache and 32KB of Data cache. All four cores share a common large 2MB L2 cache, which is 16-way set associative. The large 128 entry deep Out-of-order buffer allows the L2 cache latency to be largely hidden. Along with the 32KB L1 Caches, the 2MB L2 cache works to minimize off-chip fetches to system memory, both increasing performance and reducing power as DRAM fetches are more power intensive than on-chip SRAM fetches.

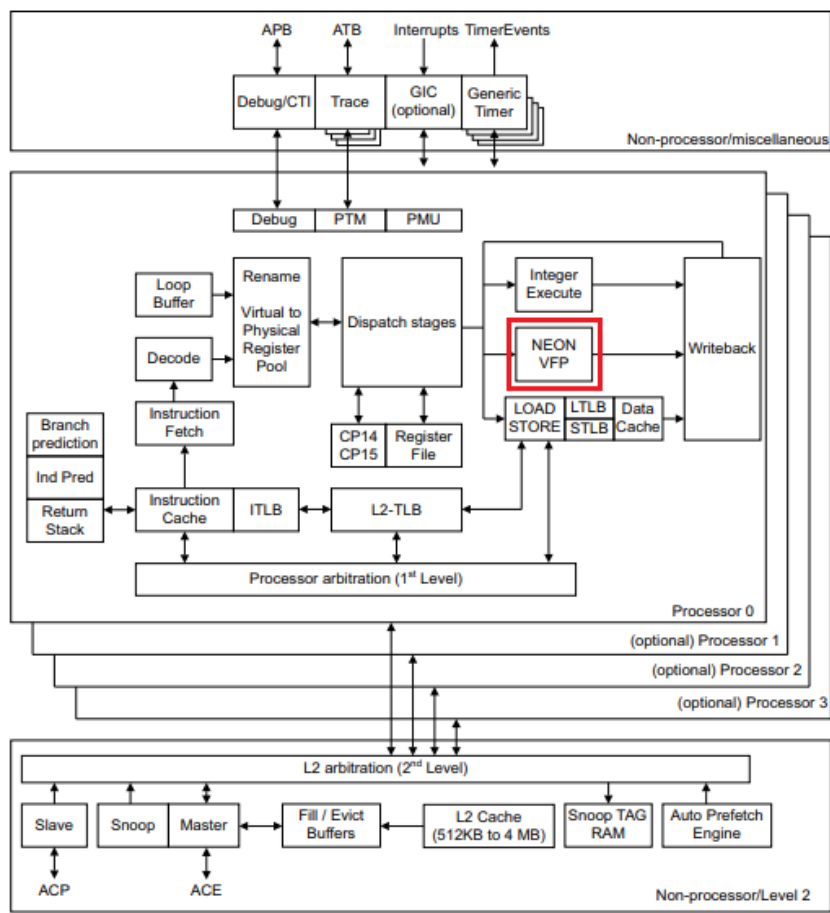
https://www.nvidia.com/docs/IO/116757/NVIDIA_Quad_a15_whitepaper_FINALv2.pdf).

The NVIDIA® Tegra® 4 series processor is a complete applications and digital media system built around several powerful hardware elements:

- CPU Complex: Quad Cortex™-A15 Symmetric Multi-Processing ARM® Cores in a 4-PLUS-1™ configuration with a quad-core fast CPU complex and a fifth Battery Saver Core. The Cortex-A15 core features triple instruction issue and both out-of-order and speculative execution. It has full cache coherency support for the quad symmetric processors. All processors have 32 KB Instruction and 32 KB Data Level 1 caches; and there is a 2 MB shared Level 2 cache for the quad-core complex and a 512 KB Level 2 cache for the fifth core. The NVIDIA 4-PLUS-1 architecture uses the fifth Battery Saver Core, which operates exclusively with the main CPU complex, for very low-power, low-leakage operation at the light CPU loads common to multimedia and lightly loaded use situations.
- Memory Controller: dual-channel (2x 32-bit) DRAM interface providing more than twice the available bandwidth of Tegra 3 devices. LP-DDR2, LP-DDR3 and DDR3 DRAM types are all supported.

https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxIV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTCld2Tuc31pIRJdQe bODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17

Figure 2-1 shows a block diagram of the Cortex-A15 processor.



(http://infocenter.arm.com/help/topic/com.arm.doc.ddi0438c/DDI0438C_cortex_a15_r2p0_trm.pdf).

29. The Accused Instrumentality comprises media processors with each processor comprising a multiplier (*e.g.*, an Integer MUL or FP MUL) having a data input coupled to the media processing unit input/output, an instruction input coupled to the media processing unit input/output, and a data output coupled to the media processing unit input/output. As shown below, the Accused Instrumentality comprises multiple ARM cortex-A15 multicore processor, each processor comprises a NEON media coprocessor and acts as a media processing unit. NEON media coprocessor comprises a multiplier which is coupled to the inputs/outputs of the processor. Upon information and belief, the multiplier comprises a data input, an instruction input, and a data output coupled to the input/output of the processor.

18.0 CPU

The NVIDIA® Tegra® 4 series processor CPU complex contains quad ARM® Cortex™ -A15 CPUs in a 4-PLUS-1 configuration with a fifth architecturally identical power-saving Cortex-A15 Companion Core.

18.1 Cortex-A15 CPU

Cortex-A15 is an advanced processor design with many features for high instruction throughput. It integrates the L2 cache controller into the CPU complex unlike Cortex-A9. All of the CPUs include the NEON Media Processing Engine. Further details of the Cortex-A15 itself are available from ARM.

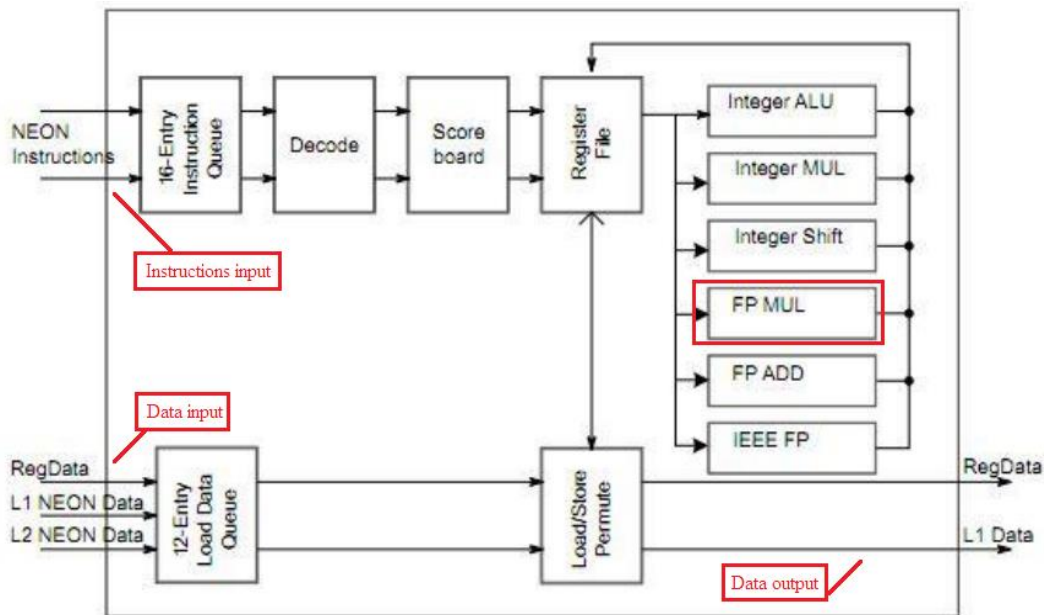
These two documents are the key references on Cortex-A15, and both are available from ARM's website:

- Cortex-A15
Revision: r2p1
Technical Reference Manual

Published by ARM Limited, document number ARM DDI 0438D.
- ARM Architecture Reference Manual
ARM v7-A and ARM v7-R edition

Published by ARM Limited, document number ARM DDI 0406C.

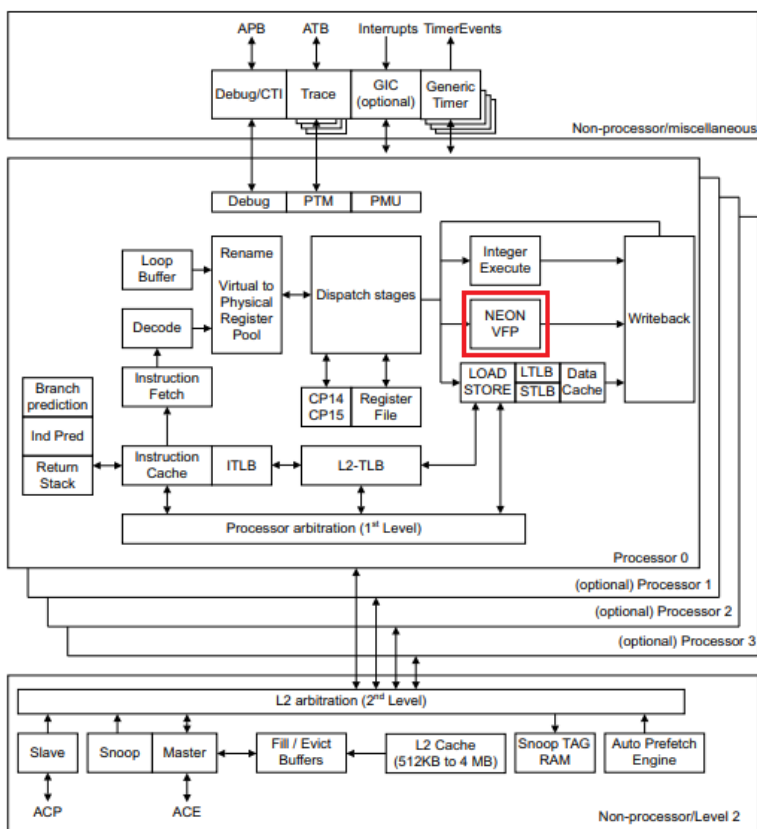
(*Id.*).



(E.g., <http://www.add.ece.ufl.edu/4924/docs/arm/ARM%20NEON%20Development.pdf>).

30. The Accused Instrumentality comprises media processors with each processor comprising an arithmetic unit (e.g., an FP ADD) having a data input coupled to the media processing unit input/output, an instruction input coupled to the media processing unit input/output, and a data output coupled to the media processing unit input/output. As shown below, the Accused Instrumentality comprises multiple ARM cortex-A15 multicore processor, each processor comprises a NEON media coprocessor and acts as a media processing unit. NEON media coprocessor comprises an arithmetic unit which is coupled to the inputs/outputs of the processor. Upon information and belief, the arithmetic unit comprises a data input, an instruction input, and a data output coupled to the input/output of the processor.

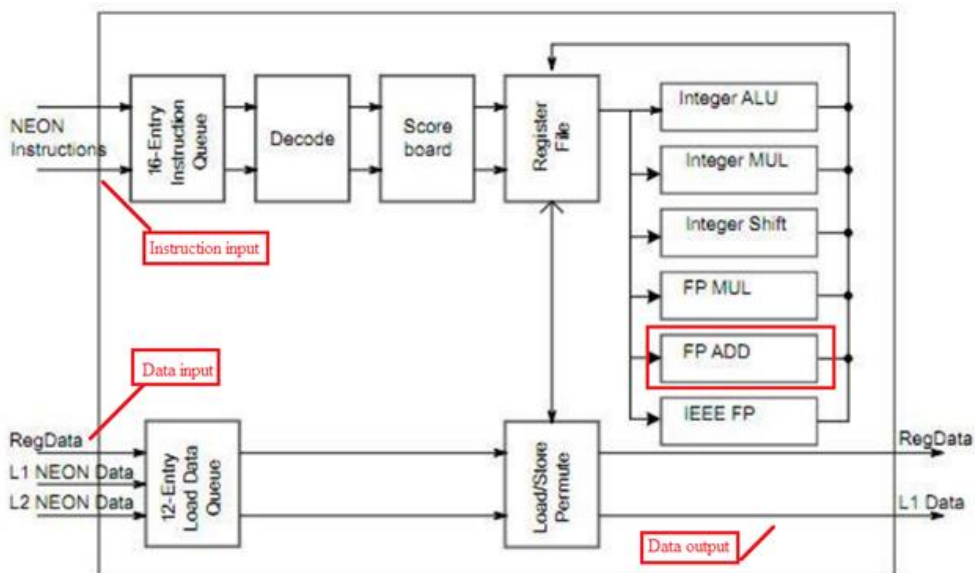
Figure 2-1 shows a block diagram of the Cortex-A15 processor.



(E.g.,

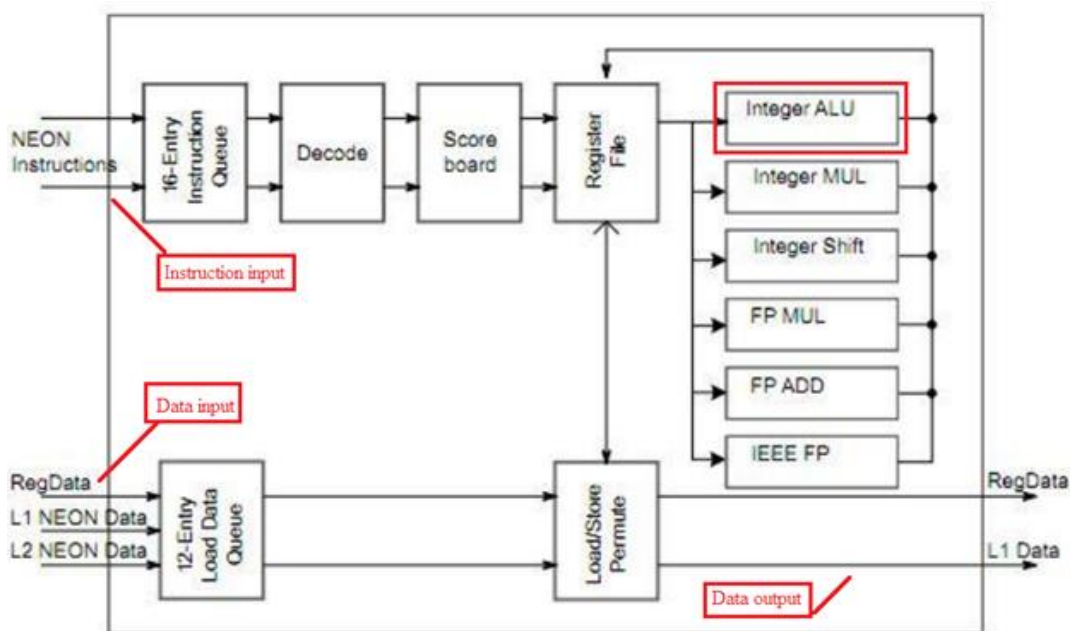
http://infocenter.arm.com/help/topic/com.arm.doc.ddi0438c/DDI0438C_cortex_a15_r2p0_trm.p

df).



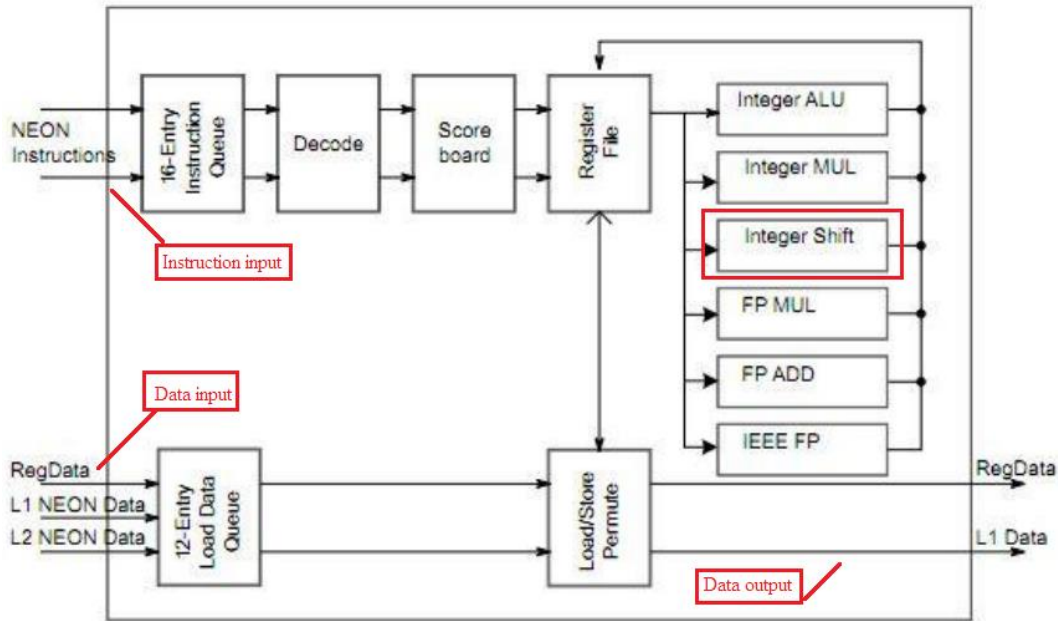
(E.g., <http://www.add.ece.ufl.edu/4924/docs/arm/ARM%20NEON%20Development.pdf>).

31. The Accused Instrumentality comprises media processors with each processor comprising an arithmetic logic unit (e.g., an ALU) having a data input coupled to the media processing unit input/output, an instruction input coupled to the media processing unit input/output, and a data output coupled to the media processing unit input/output, capable of operating concurrently with at least one selected from the multiplier (e.g., an Integer MUL or FP MUL) and arithmetic unit (e.g., a FP ADD). As shown below, the Accused Instrumentality comprises multiple ARM cortex-A15 multicore processor, each processor comprises a NEON media coprocessor and acts as a media processing unit. NEON media coprocessor comprises an arithmetic logical unit which is coupled to the inputs/outputs of the processor. Upon information and belief, the arithmetic logical unit comprises a data input, an instruction input, and a data output coupled to the input/output of the processor. Upon information and belief, the arithmetic logical unit (e.g., the Integer ALU) is capable of operating concurrently with at least one selected from the multiplier (e.g., the Integer MUL or FP MUL) and arithmetic unit (e.g., the FP ADD).



(E.g., <http://www.add.ece.ufl.edu/4924/docs/arm/ARM%20NEON%20Development.pdf>).

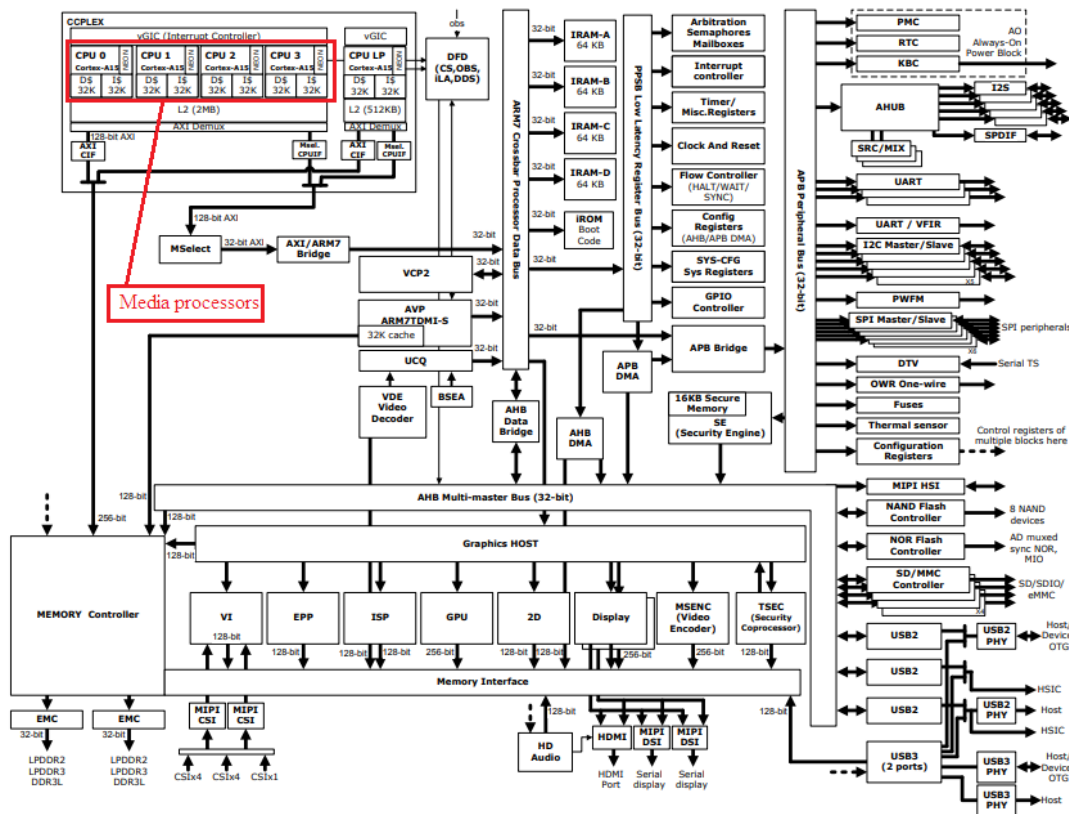
32. The Accused Instrumentality comprises media processors with each processor comprising a bit manipulation unit (e.g., an Integer Shift unit) having a data input coupled to the media processing unit input/output, an instruction input coupled to the media processing unit input/output, and a data output coupled to the media processing unit input/output, capable of operating concurrently with the arithmetic logic unit (e.g., an Integer ALU) and at least one selected from the multiplier (e.g., an Integer MUL or FP MUL) and arithmetic unit (e.g., a FP ADD). As shown below, the Accused Instrumentality comprises multiple ARM cortex- A15 multicore processors, each processor comprising a NEON media coprocessor that acts as a media processing unit. The NEON media coprocessor comprises an integer shift unit (i.e., bit manipulation unit) which is coupled to the inputs/outputs of the processor. Upon information and belief, the integer shift unit (i.e., bit manipulation unit) comprises a data input, an instruction input, and a data output coupled to the input/output of the processor. Upon information and belief, the integer shift unit (i.e., bit manipulation unit) is capable of operating concurrently with the arithmetic logic unit (e.g., the Integer ALU) and at least one selected from the multiplier (e.g., the Integer MUL or FP MUL) and arithmetic unit (e.g., the FP ADD).



(E.g., <http://www.add.ece.ufl.edu/4924/docs/arm/ARM%20NEON%20Development.pdf>).

33. The Accused Instrumentality comprises a plurality of media processors (e.g., ARM cortex-A15 multicore processors) for performing at least one operation, simultaneously with the performance of other operations by other media processing units (e.g., other ARM cortex-A15 multicore processors on the same chip).

Figure 1: Tegra 4 Series Processor Block Diagram



(E.g.,

[https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06_Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxIV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTCId2Tuc31plRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17\).](https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06_Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxIV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTCId2Tuc31plRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17).)

18.0 CPU

The NVIDIA® Tegra® 4 series processor CPU complex contains quad ARM® Cortex™-A15 CPUs in a 4-PLUS-1 configuration with a fifth architecturally identical power-saving Cortex-A15 Companion Core.

18.1 Cortex-A15 CPU

Cortex-A15 is an advanced processor design with many features for high instruction throughput. It integrates the L2 cache controller into the CPU complex unlike Cortex-A9. All of the CPUs include the NEON Media Processing Engine. Further details of the Cortex-A15 itself are available from ARM.

These two documents are the key references on Cortex-A15, and both are available from ARM's website:

- Cortex-A15
Revision: r2p1
Technical Reference Manual

Published by ARM Limited, document number ARM DDI 0438D.
- ARM Architecture Reference Manual
ARM v7-A and ARM v7-R edition

Published by ARM Limited, document number ARM DDI 0406C.

https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxIV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTCld2Tuc31pIRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17

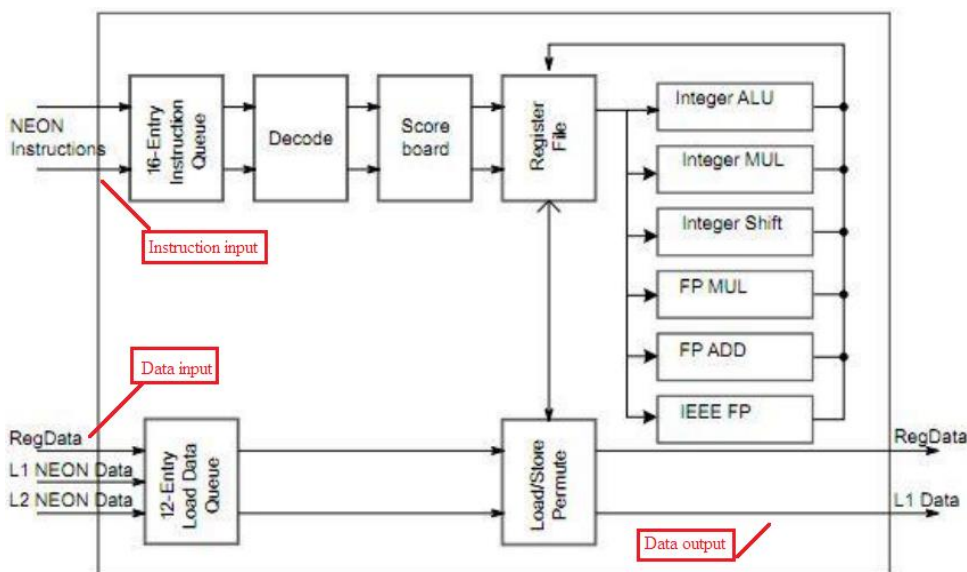
The NVIDIA® Tegra® 4 series processor is a complete applications and digital media system built around several powerful hardware elements:

- CPU Complex: Quad Cortex™-A15 Symmetric Multi-Processing ARM® Cores in a 4-PLUS-1™ configuration with a quad-core fast CPU complex and a fifth Battery Saver Core. The Cortex-A15 core features triple instruction issue and both out-of-order and speculative execution. It has full cache coherency support for the quad symmetric processors. All processors have 32 KB Instruction and 32 KB Data Level 1 caches; and there is a 2 MB shared Level 2 cache for the quad-core complex and a 512 KB Level 2 cache for the fifth core. The NVIDIA 4-PLUS-1 architecture uses the fifth Battery Saver Core, which operates exclusively with the main CPU complex, for very low-power, low-leakage operation at the light CPU loads common to multimedia and lightly loaded use situations.
- Memory Controller: dual-channel (2x 32-bit) DRAM interface providing more than twice the available bandwidth of Tegra 3 devices. LP-DDR2, LP-DDR3 and DDR3 DRAM types are all supported.

(*Id.*).

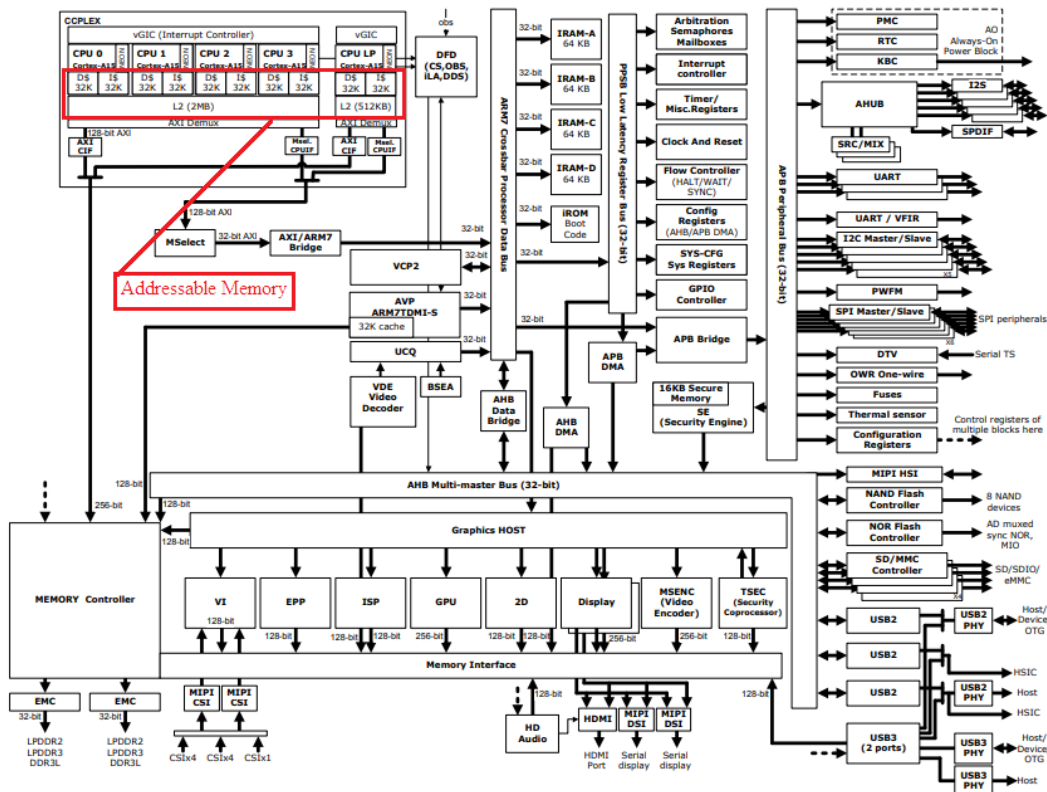
34. The Accused Instrumentality comprises a plurality of media processors (*e.g.*, ARM cortex-A15 multicore processors), each processor receiving at the media processor input/output an instruction and data from the memory, and processing the data responsive to the instruction received to produce at least one result. As previously shown, each ARM cortex-A15 multicore

media processor comprises a NEON media coprocessor which receives instructions and data from memory and processes the data responsive to the instruction received in order to produce a result.



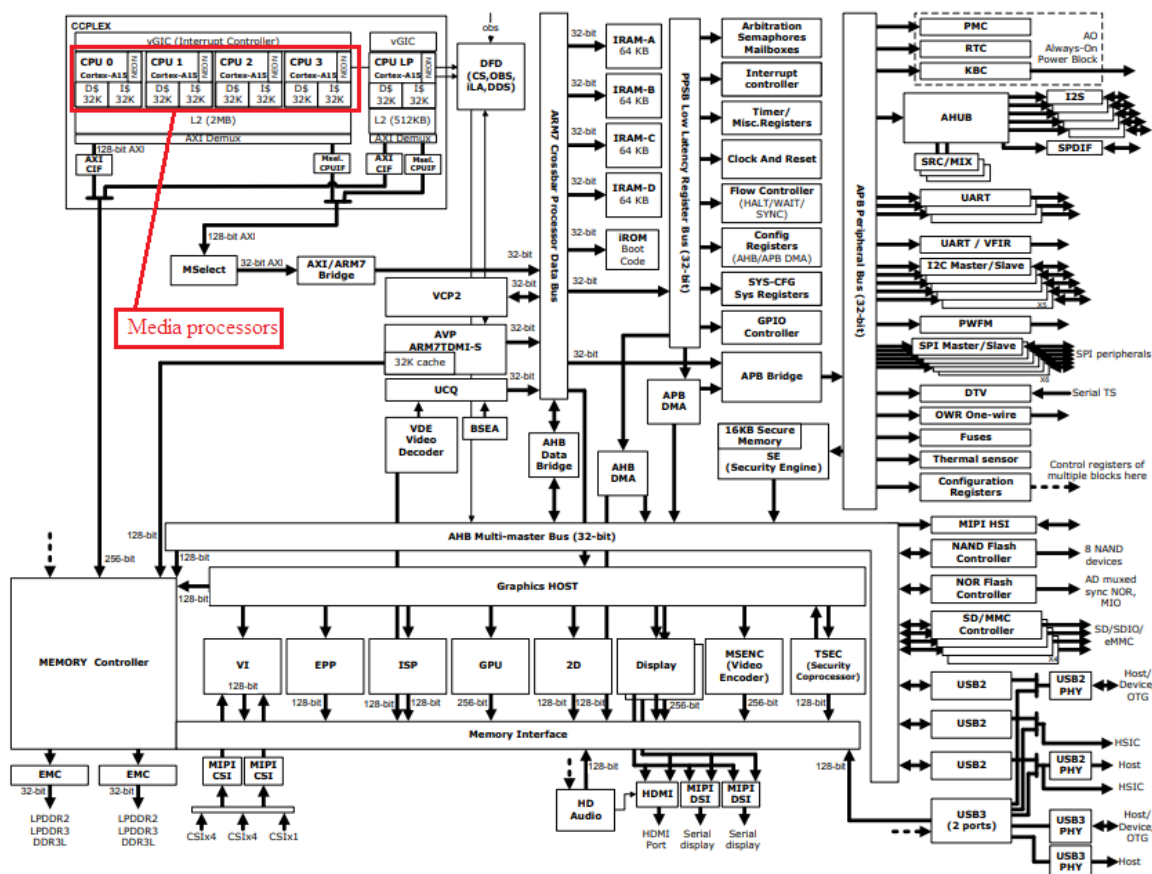
(E.g., <http://www.add.ece.ufl.edu/4924/docs/arm/ARM%20NEON%20Development.pdf>).

Figure 1: Tegra 4 Series Processor Block Diagram



https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxIV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTCld2Tuc31plRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17

Figure 1: Tegra 4 Series Processor Block Diagram



https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxIV83ImQd4YznqZ9sCWxNmXHWwCh5oyc6cKMvTCld2Tuc31plRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17

35. The Accused Instrumentality comprises a plurality of media processors (e.g., ARM cortex-A15 multicore processors), each processor providing at least one of the at least one result at the media processor input/output. (*Supra* ¶34).

18.0 CPU

The NVIDIA® Tegra® 4 series processor CPU complex contains quad ARM® Cortex™-A15 CPUs in a 4-PLUS-1 configuration with a fifth architecturally identical power-saving Cortex-A15 Companion Core.

18.1 Cortex-A15 CPU

Cortex-A15 is an advanced processor design with many features for high instruction throughput. It integrates the L2 cache controller into the CPU complex unlike Cortex-A9. All of the CPUs include the NEON Media Processing Engine. Further details of the Cortex-A15 itself are available from ARM.

These two documents are the key references on Cortex-A15, and both are available from ARM's website:

- Cortex-A15
Revision: r2p1
Technical Reference Manual

Published by ARM Limited, document number ARM DDI 0438D.
- ARM Architecture Reference Manual
ARM v7-A and ARM v7-R edition

Published by ARM Limited, document number ARM DDI 0406C.

(E.g.,

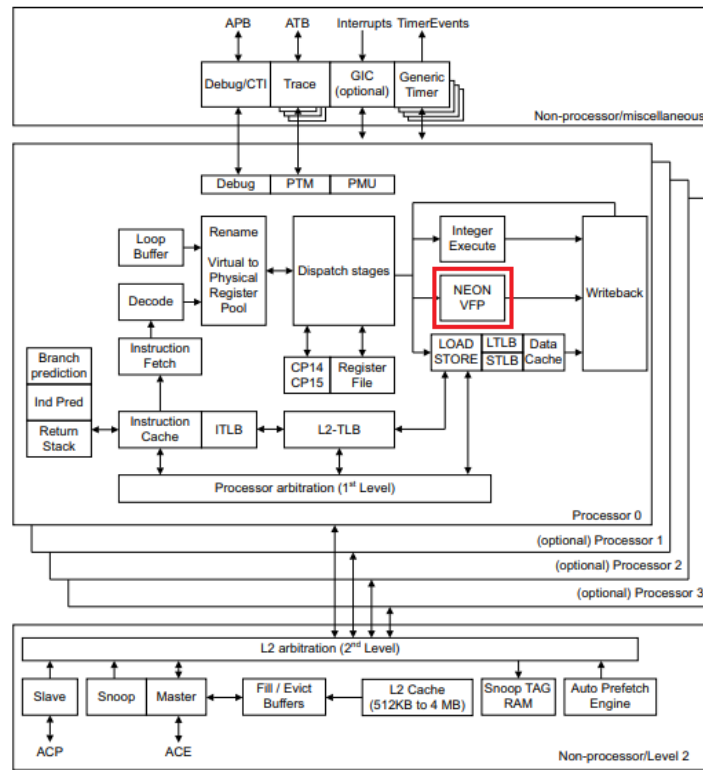
https://developer.download.nvidia.com/assets/embedded/secure/docs/Tegra4_publicTRMv01_06-Sep.pdf?IXpAikYEePajX58WX2jYTQn7IzzqBT0u-C-C4_tz2TUwzkmYYxIV83ImQd4YzmqZ9sCWxNmXHWwCh5oyc6cKMvTClD2Tuc31plRJdQebODwoFIVT2aqQeHnP2IVCLcl4p6Hjx7WotqWgJPbSbjKhkUouB5BEQF17).

The NVIDIA® Tegra® 4 series processor is a complete applications and digital media system built around several powerful hardware elements:

- CPU Complex: Quad Cortex™-A15 Symmetric Multi-Processing ARM® Cores in a 4-PLUS-1™ configuration with a quad-core fast CPU complex and a fifth Battery Saver Core. The Cortex-A15 core features triple instruction issue and both out-of-order and speculative execution. It has full cache coherency support for the quad symmetric processors. All processors have 32 KB Instruction and 32 KB Data Level 1 caches; and there is a 2 MB shared Level 2 cache for the quad-core complex and a 512 KB Level 2 cache for the fifth core. The NVIDIA 4-PLUS-1 architecture uses the fifth Battery Saver Core, which operates exclusively with the main CPU complex, for very low-power, low-leakage operation at the light CPU loads common to multimedia and lightly loaded use situations.
- Memory Controller: dual-channel (2x 32-bit) DRAM interface providing more than twice the available bandwidth of Tegra 3 devices. LP-DDR2, LP-DDR3 and DDR3 DRAM types are all supported.

(Id.).

Figure 2-1 shows a block diagram of the Cortex-A15 processor.



(E.g.,

http://infocenter.arm.com/help/topic/com.arm.doc.ddi0438c/DDI0438C_cortex_a15_r2p0_trm.pdf).

36. Plaintiff has been damaged as a result of Defendant's infringing conduct. Defendant is thus liable to Plaintiff for damages in an amount that adequately compensates Plaintiff for such Defendant's infringement of the '434 patent, *i.e.*, in an amount that by law cannot be less than would constitute a reasonable royalty for the use of the patented technology, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

37. On information and belief, Defendant has had at least constructive notice of the '434 patent by operation of law, and there are no marking requirements that have not been complied with.

IV. JURY DEMAND

Plaintiff, under Rule 38 of the Federal Rules of Civil Procedure, requests a trial by jury of any issues so triable by right.

V. PRAYER FOR RELIEF

WHEREFORE, Plaintiff respectfully requests that the Court find in its favor and against Defendant, and that the Court grant Plaintiff the following relief:

- a. Judgment that one or more claims of United States Patent No. 6,289,434 have been infringed, either literally and/or under the doctrine of equivalents, by Defendant;
- b. Judgment that Defendant account for and pay to Plaintiff all damages to and costs incurred by Plaintiff because of Defendant's infringing activities and other conduct complained of herein;
- c. That Plaintiff be granted pre-judgment and post-judgment interest on the damages caused by Defendant's infringing activities and other conduct complained of herein; and
- d. That Plaintiff be granted such other and further relief as the Court may deem just and proper under the circumstances.

Dated: July 30, 2018

Respectfully submitted,

/s/ David R. Bennett

By: David R. Bennett
DIRECTION IP LAW
P.O. Box 14184
Chicago, IL 60614-0184
Telephone: (312) 291-1667
e-mail: dbennett@directionip.com

**ATTORNEY FOR PLAINTIFF
ALTAIR LOGIX LLC**