

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

DAEDALUS PRIME LLC,

Plaintiff,

v.

MEDIATEK INC.

Defendant.

Civil Action No. 2:24cv235

JURY TRIAL DEMANDED

**COMPLAINT FOR PATENT INFRINGEMENT AND DAMAGES
AND DEMAND FOR JURY TRIAL**

Plaintiff Daedalus Prime LLC (“Daedalus” or “Plaintiff”) files this Complaint for Patent Infringement and Damages against MediaTek Inc. (“MediaTek” or “Defendant”) and alleges as follows:

INTRODUCTION

1. The novel inventions disclosed in U.S. Patent Nos. 8,769,316 (the “’316 Patent”); 10,372,197 (the “’197 Patent”); 10,740,281 (the “’281 Patent”); 8,984,228 (the “’228 Patent”); 11,507,167 (the “’167 Patent”); 9,887,838 (the “’838 Patent”); 10,705,960 (the “’960 Patent”) and 10,725,919 (the “’919 Patent”) (collectively, the “Asserted Patents”) in this matter were invented by Intel Corporation (“Intel”). Intel pioneered the field of microprocessor and semiconductor chip technology. This technology provides capabilities that are crucial to electronic devices such as personal computers and smartphones. Every year, Intel spends billions of dollars on research and development to invent, market, and sell new technology, and Intel obtains patents on many of the novel inventions that come out of that work, including the Asserted Patents.

THE PARTIES

2. Plaintiff is the current owner and assignee of the Asserted Patents.
3. Plaintiff is a Delaware limited liability company with its principal place of business located at 51 Pondfield Road, Suite 3, Bronxville, New York 10708.
4. On information and belief, Defendant MediaTek Inc. is a corporation organized and existing under the laws of Taiwan, and located at No. 1, Dusing Road 1, Hsinchu Science Park, Hsinchu City 30078, Taiwan.
5. On information and belief, Defendant directly and/or indirectly develops, designs, manufactures, distributes, markets, offers to sell and/or sells infringing products and services in the United States, including in the Eastern District of Texas, and otherwise direct infringing activities to this District in connection with their products and services as set forth in this Complaint.

JURISDICTION

6. This civil action arises under the Patent Laws of the United States, 35 U.S.C. § 1 *et seq.*, including without limitation 35 U.S.C. §§ 271, 281, 283, 284, and 285. Accordingly, this Court has subject matter jurisdiction under, *inter alia*, 28 U.S.C. §§ 1331 and 1338(a).
7. MediaTek sells semiconductors and/or processors that are used in mobile phones and other consumer products in the United States. For example, upon information and belief, MediaTek sells Dimensity 9300, Dimensity 9200 and Dimensity 9000 smartphone chips, as well as the MediaTek Kompanio 1380 Chromebook chips and the Dimensity Auto Cockpit CX-1 automotive chips in the U.S.¹:

¹ <https://www.mediatek.com/products/smartphones-2/mediatek-dimensity-9300>;
<https://www.mediatek.com/products/smartphones-2/mediatek-dimensity-9200>;
<https://www.mediatek.com/products/smartphones-2/mediatek-dimensity-9000>;

Powering 2 Billion+ Devices a year

MediaTek Makes **Great Technology** Available To Everyone

Chances are you already have a MediaTek powered device in your life. MediaTek chips power more than **2 billion devices** every year.

New MediaTek Powered Devices Each Year

2,013,258,324+

Powered by
MediaTek Kompanio

As the No. 1 chipmaker for Arm-based Chromebooks, MediaTek has played a pivotal role in keeping students and teachers connected for in-person and virtual learning. MediaTek's Kompanio series of chipsets powers the Chromebooks that students and teachers love.

Powering the brands you love



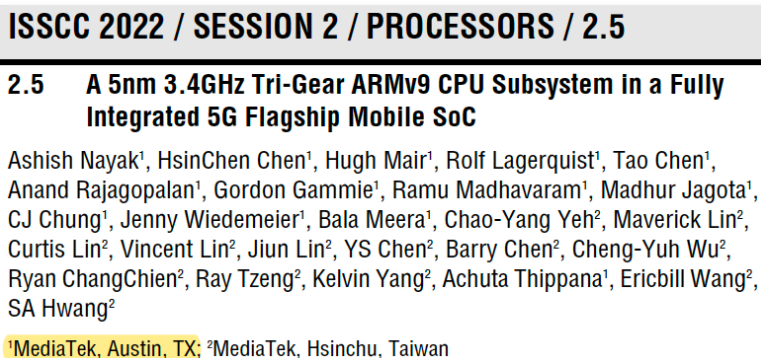


8. This Court has personal jurisdiction over MediaTek Inc. at least because MediaTek Inc. sells, offers for sale, uses, makes and/or imports products that are and have been used, offered for sale, sold, and purchased in the Eastern District of Texas, and MediaTek Inc. has committed, and continues to commit, acts of infringement in the Eastern District of Texas, has conducted business in the Eastern District of Texas, and/or has engaged in continuous and systematic activities in the Eastern District of Texas.

9. Under 28 U.S.C. §§ 1391(b)-(d) and 1400(b), venue is proper in this judicial district as to MediaTek Inc. at least because MediaTek Inc. is a foreign corporation subject to personal jurisdiction in this judicial district and has committed acts of infringement within this judicial district giving rise to this action.

10. On information and belief, Defendant or Defendant's subsidiaries have physical facilities and employees in Texas, including an office at 2435 North Central Expressway, Suite

750, Richardson, Texas 75080. On information and belief, Defendant or Defendant's subsidiaries maintain multiple offices in Texas and have numerous employees in Texas²:



11. Defendant has not contested proper venue and exercise of personal jurisdiction in this District for patent infringement actions. *See, e.g.*, Answer, ¶¶ 12-20, *Mosaid Techs. Inc. v. Mediatek Inc. et al.*, No. 2:23-cv-00129, ECF 24 (E.D. Tex. July 18, 2023); Answer, ¶¶ 11-13, *Am. Patents, LLC v. Mediatek Inc., et al.*, No. 4:22-cv-487, ECF 18 (E.D. Tex. Sept. 6, 2022).

THE ASSERTED PATENTS

12. The Intel inventions contained in the Asserted Patents in this case relate to groundbreaking improvements to microprocessor circuitry and mobile wireless, and have particular application in consumer electronics such as smartphones, tablets, and personal computers.

U.S PATENT NO. 8,769,316

13. On July 1, 2014, the United States Patent Office duly and legally issued the '316 Patent, entitled "Dynamically Allocating a Power Budget Over Multiple Domains of a Processor."³

² Ashish Nayak et al., *A 5nm 3.4GHz Tri-Gear ARMv9 CPU Subsystem in a Fully Integrated 5G Flagship Mobile SoC*, ISSCC, 2022.

³ <https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/8769316>.

14. Daedalus is the owner and assignee of all right, title, and interest in and to the '316 Patent, including the right to assert all causes of action arising under the '316 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

15. The '316 Patent describes, among other things, a method for determining a power budget for a multi-domain processor for a current time interval, determining a portion of the power budget to be allocated to first and second domains of the processor, and controlling a frequency of the domains based on the allocated portions. '316 Patent, Abstract. As the '316 Patent explains, one issue with multicore processors was that “the different circuitry can consume differing amounts of power based on their workloads” but “suitable mechanisms to ensure that these different units have sufficient power do not presently exist.” *Id.* at 1:18-22.

16. The '316 Patent seeks to solve the problem with multicore processors. The novel inventions of the '316 Patent are recited in the claims. For example, claim 8 of the '316 Patent recites:

8. A method comprising:

determining, in a power controller of a multi-domain processor, a power budget for the multi-domain processor for a current time interval, the multi-domain processor including at least a first domain and a second domain;

determining, in the power controller, a portion of the power budget to be allocated to the first and second domains, including allocating a minimum reservation value to the first domain and a minimum reservation value to the second domain, and sharing a remaining portion of the power budget according to a first sharing policy value for the first domain and a second sharing policy value for the second domain; and

controlling a frequency of the first domain and a frequency of the second domain based on the allocated portions.

'316 Patent, Cl. 8.

17. Figure 6 of the '316 Patent, reproduced below, shows a block diagram of a portion of a system of one embodiment of the claimed invention. “As shown in FIG. 6, processor **300** may

be a multicore processor including a plurality of cores [310-310n,] ... [where] each [] core may be of an independent power domain ... configured to operate at an independent voltage and/or frequency”. ’316 Patent, 8:21-25. The cores may be “coupled via an interconnect **315**” with a shared cache 330. *Id.* at 8:26-27. “[P]ower control unit **355** may include a power sharing logic **359** ... [that] perform[s] dynamic control and re-allocation of an available power budget between multiple independent domains of the processor.” *Id.* at 8:33-36.

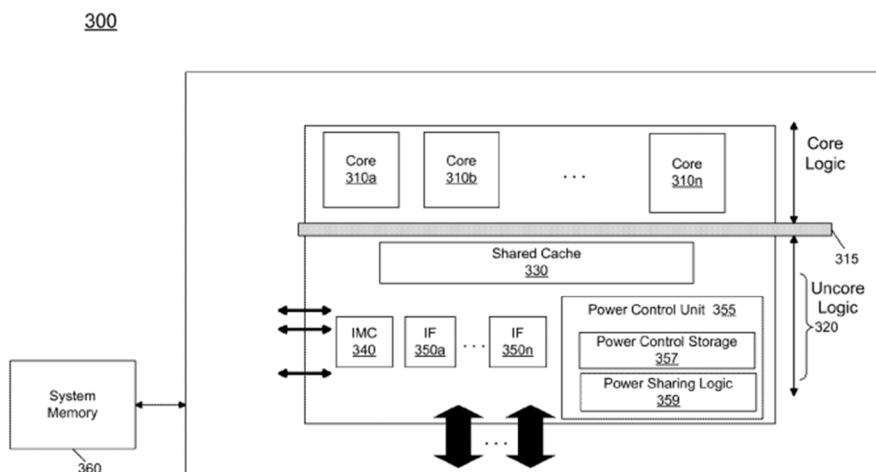


FIG. 6

’316 Patent, Fig. 6.

U.S. PATENT NO. 10,372,197

18. On August 6, 2019, the United States Patent Office duly and legally issued the ’197 Patent, entitled “User Level Control of Power Management Policies.”⁴

19. Daedalus is the owner and assignee of all right, title, and interest in and to the ’197 Patent, including the right to assert all causes of action arising under the ’197 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

⁴ <https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/10372197>.

20. The '197 Patent describes, among other things, a multicore processor comprising a power controller that receives a workload configuration input and a plurality of energy performance bias values, determines a global energy performance bias value to update one or more power settings of one or more management features. '197 Patent, Fig. 4, 6:61-7:14; 7:55-8:4.

21. The novel features of the invention are recited in the claims. For example, claim 1 of the '197 Patent recites:

1. A processor comprising:

a plurality of cores;

a cache memory;

an interconnect to couple the plurality of cores and the cache memory; and

a power controller to control a plurality of power management features of the processor, wherein the power controller includes a tuning circuit to receive a workload configuration input regarding a workload, receive a plurality of energy performance bias (EPB) values and determine a global EPB value based thereon, and update at least one setting of at least one of the plurality of power management features based on the workload configuration input and the global EPB value.

Id. at Cl. 1.

22. Figure 4 of the '197 Patent, reproduced below, is a block diagram of a processor in accordance with an embodiment of the inventions disclosed in the '197 Patent. As shown in Figure 4, processor 300 may be a multicore processor including a plurality of cores 310a-31. The various cores may be coupled via an interconnect 315 to a system agent or uncore 320 that includes various components. The uncore 320 may include a shared cache 330 which may be a last level cache. In addition, the uncore may include a power control unit 355.

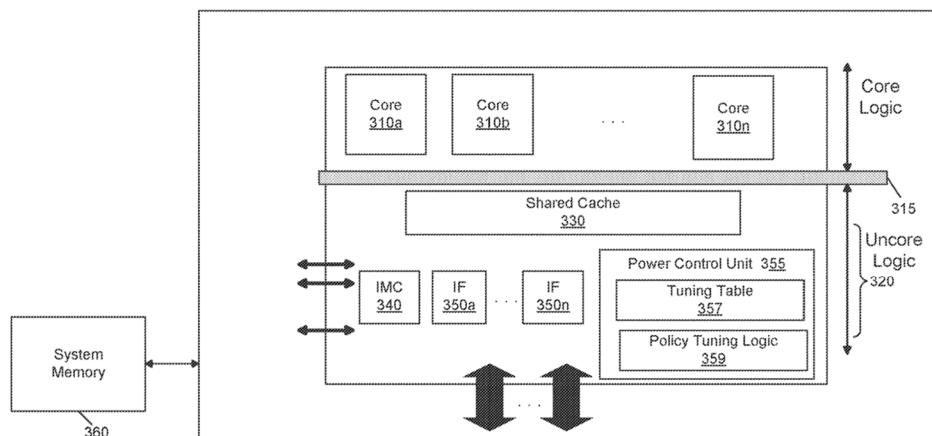


FIG. 4

'197 Patent, Fig. 4.

U.S PATENT NO. 10,740,281

23. On August. 11, 2020, the United States Patent Office duly and legally issued the '281 Patent, entitled "Asymmetric Performance Multicore Architecture with Same Instruction Set Architecture."⁵

24. Daedalus is the owner and assignee of all right, title, and interest in and to the '281 Patent, including the right to assert all causes of action arising under the '281 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

25. The '281 Patent describes, among other things, a method of operating enabled cores of a multi-core processor such that both cores support respective software routines with a same instruction set, with a first core being higher performance and consuming more power than a second core under a same set of applied supply voltage and operating frequency. '281 Patent, Abstract. The '281 Patent describes "a new approach in which at least one of the cores **401** is designed to be lower performance and therefore consume less power than other cores **402** in the

⁵ <https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/10740281>.

processor. However, the lower power core(s) **401** has a same logic design as the higher power core(s) **402** and therefore supports the same instruction set **403** as the high power core(s) **402**. The low power core(s) **401** achieve a lower power design point by having narrower drive transistor widths than the higher power core(s) and/or having other power consumption related design features”. ’281 Patent, 3:58-67.

26. The novel features of the invention are recited in the claims. For example, claim 8 of the ’281 Patent recites:

8. A method comprising:

monitoring a demand for a multi-core processor by an operating system executing on the multi-core processor, wherein the multi-core processor comprises a first plurality of cores and a second plurality of cores that support a same instruction set, the first plurality of cores are higher performance and consume more power than the second plurality of cores, each of the second plurality of cores have a maximum operating frequency that is less than a maximum operating frequency of each of the first plurality of cores, and a caching layer shared by the first plurality of cores and the second plurality of cores; and

controlling a core mix of the first plurality of cores and the second plurality of cores based on the demand with power management hardware of the multi-core processor.

’281 Patent, Cl. 8.

27. Figure 4 of the ’281 Patent, reproduced below, shows a block diagram of a portion of a system of one embodiment of the claimed invention. As shown in Figure 4, cores 401 are “lower performance and therefore consume less power than [the higher power] cores **402** in the processor.” The “lower power core[s] **401** has a same logic design as the higher power core[s] **402** and therefore support[] the same instruction set **403** as the high power core[s] **402**. *Id.* at 3:58-65.




FIG. 4

'281 Patent, Fig. 4.

U.S PATENT NO. 8,984,228

28. On March 17, 2015, the United States Patent Office duly and legally issued the '282 Patent, entitled "Providing Common Caching Agent for Core and Integrated Input/Output (IO) Module."⁶

29. Daedalus is the owner and assignee of all right, title, and interest in and to the '228 Patent, including the right to assert all causes of action arising under the '228 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

30. The '228 Patent describes, among other things, a multicore processor having a plurality of cores, a shared cache memory, an integrated input/output module to interface between

⁶ <https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/8984228>.

the multicore processor and at least one IO device coupled to the multicore processor, and a caching agent to perform cache coherency operations for the plurality of cores and the integrated input/output module. '228 Patent, Abstract. As the '228 Patent explains, “problems arise once an IO component is integrated on the same chip with a multiprocessor. Traditional IO integration treats the IO component as a separate caching agent, meaning that dedicated logic is associated with the IO component to handle cache coherency operations. When an IO agent is performing read/write operations to main memory, it has to snoop the CPU side cache to maintain cache coherency.” *Id.* at 1:19-25. The integrated input/output module, however, “reduces the amount of snoop traffic needed since a reduced number of caching agents per system can be realized.” *Id.* at 3:8-9.

31. The novel features of the invention are recited in the claims. For example, claim 1 of the '228 Patent recites:

1. An apparatus comprising:

a multicore processor including a plurality of cores, a shared cache memory, an integrated input/output (IIO) module to interface between the multicore processor and at least one IO device coupled to the multicore processor, and a caching agent to perform cache coherency operations for the plurality of cores and the IIO module, the caching agent a single caching agent for the multicore processor and including a plurality of distributed portions each associated with a corresponding one of the plurality of cores.

'228 Patent, Cl. 1.

32. Figure 6 of the '228 Patent, reproduced below, shows a block diagram of a portion of a system of one embodiment of the claimed invention. As shown in Figure 6, “processor **700** includes a distributed configuration having partitions or slices each including a core **710** and a partition of a caching agent **715** and a LLC **720**. Note that while distributed caching agents are shown, understand that these distributed portions form a single caching agent, and which is

configured to handle cache coherency operations both for the cores as well as an IIO module 750.”
Id. at 6:23-29.




FIG. 6

'228 Patent, Fig. 6.

U.S. PATENT NO. 11,507,167

33. On November 22, 2022, the United States Patent Office duly and legally issued the '167 Patent, entitled “Controlling Operating Voltage of a Processor.”⁷

34. Daedalus is the owner and assignee of all right, title, and interest in and to the '167 Patent, including the right to assert all causes of action arising under the '167 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

35. The '167 Patent describes, among other things, a processor that includes a core domain with a plurality of cores, and a power controller capable of instructing a voltage regulator

⁷ <https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/11507167>.

to increase the operating voltage. '167 Patent, Abstract. As the '167 Patent explains, “there is a vital need for energy efficiency and conservation associated with integrated circuits.” *Id.* at 1:44-45. The '167 Patent explains that “voltage transitions within a processor may be segmented into two or more segments. In an embodiment, a dispatcher or other control logic of the processor may controllably cause such multi-phase voltage ramps. In operation, a first segment is a transition to an interim or safe voltage level, which is at a sufficient voltage level to cover all active agents (and at least one additional agent) running at a lower frequency in a particular transition. Any additional voltage increase to enable a pending frequency increase requested for one or more of the agents is handled in a second segment of the transition, which can occur after a low power state exit of the additional agent. [...] In this way, a reduced latency for allowing an agent to exit a low power state may be realized.” *Id.* at 2:28-46.

36. The novel features of the invention are recited in the claims. For example, claim 1 of the '167 Patent recites:

1. A multicore processor comprising:

a plurality of cores, wherein each core comprises a processor configured to operate at an independent voltage and frequency level;

wherein at least one core is coupled to a plurality of levels of cache memory;

a power control unit configured to cause an operating voltage to be updated for one or more of the cores in response to receiving a request to alter an operating state of the one or more of the cores;

wherein the power control unit is further configured to:

receive a first request to alter an operating state of a first core to a modified operating state, the modified operating state operating at a third voltage level;

responsive to the first request, cause a voltage regulator to increase an operating voltage of the first core from a first voltage level to a second voltage level lower than the third voltage level;

enable a second core to exit an inactive state and enter an active state while the operating voltage of the first core is at the second voltage level;

increase the operating voltage of the first core from the second voltage level to the third voltage level after the second core enters the active state.

'167 Patent, Cl. 1.

37. Figure 1 of the '167 Patent, reproduced below, shows a block diagram of a portion of a system of one embodiment of the claimed invention. "As shown in FIG. 1, system 100 may include various components, including a processor 110 which as shown is a multicore processor. Processor 110 may be coupled to a power supply 150 via an external voltage regulator 160, which may perform a first voltage conversion to provide a primary regulated voltage to processor 110." *Id.* at 2:49-54.




FIG. 1

'167 Patent, Fig. 1.

U.S PATENT NO. 9,887,838

38. On February 6, 2018, the United States Patent Office duly and legally issued the '838 Patent, entitled "Method and Device for Secure Communications Over a Network Using a Hardware Security Engine."⁸

39. Daedalus is the owner and assignee of all right, title, and interest in and to the '838 Patent, including the right to assert all causes of action arising under the '838 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

40. The '838 Patent describes, among other things, establishing a secure communication session with a server including initiating a request for a secure communication session with a server using a nonce value generated in a security engine of a system-on-a-chip (SOC) of a client device. '838 Patent, Abstract. As the '838 Patent explains, "there is a vital need for energy efficiency and conservation associated with integrated circuits." *Id.* at 1:44-45. The '838 Patent explains that a "security engine 110 may be embodied as a security co-processor or processing circuitry separate from the processor core 118. The security engine 110 includes the security key 150 and the secure memory 114, which is accessible only by the security engine 110. The security engine 110 stores the security key 150, and other cryptographic keys as discussed below, in the secure memory 114." *Id.* at 4:17-23.

41. The novel features of the invention are recited in the claims. For example, claim 9 of the '838 Patent recites:

9. A method comprising:

generating a random nonce in a security engine that is separate from a processor core of a system-on-a-chip of a client device;
initiating, using the client device, a request for a secure communication session with a remote server over a network, the request including the random nonce;

⁸ <https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/9887838>.

performing a cryptographic key exchange, using the security engine of the system-on-a-chip, with the remote server;

generate a symmetric session key to encrypt messages sent to the remote server and decrypt messages received from the remote server during the secure communication session;

encrypting the session key based on a security key that was encoded in a secure memory of the security engine during a manufacturing process of the system-on-a-chip;

storing the encrypted session key in the secure memory of the security engine of the system-on-a-chip; and

establishing, using the client device, the secure communication session with the remote server using the session key.

'838 Patent, Cl. 9.

42. Figure 1 of the '838 Patent, reproduced below, shows a block diagram of a portion of a system of one embodiment of the claimed invention. "In the illustrative embodiment of FIG. 1, the SOC **112** includes the security engine **110**, a memory controller **116**, a processor core **118**, and a plurality of hardware peripherals **130**, which are communicatively coupled to each other via a link **120**." *Id.* at 1:63-67. As shown in Figure 1, "a system **100** [establishes] a secure communication session include[ing] a client device **102**, a server **104**, and a network **106**. In operation, the client device **102** initiates a request for a secure communication session with the server **104** over the network **106**. *Id.* at 3:24-28.

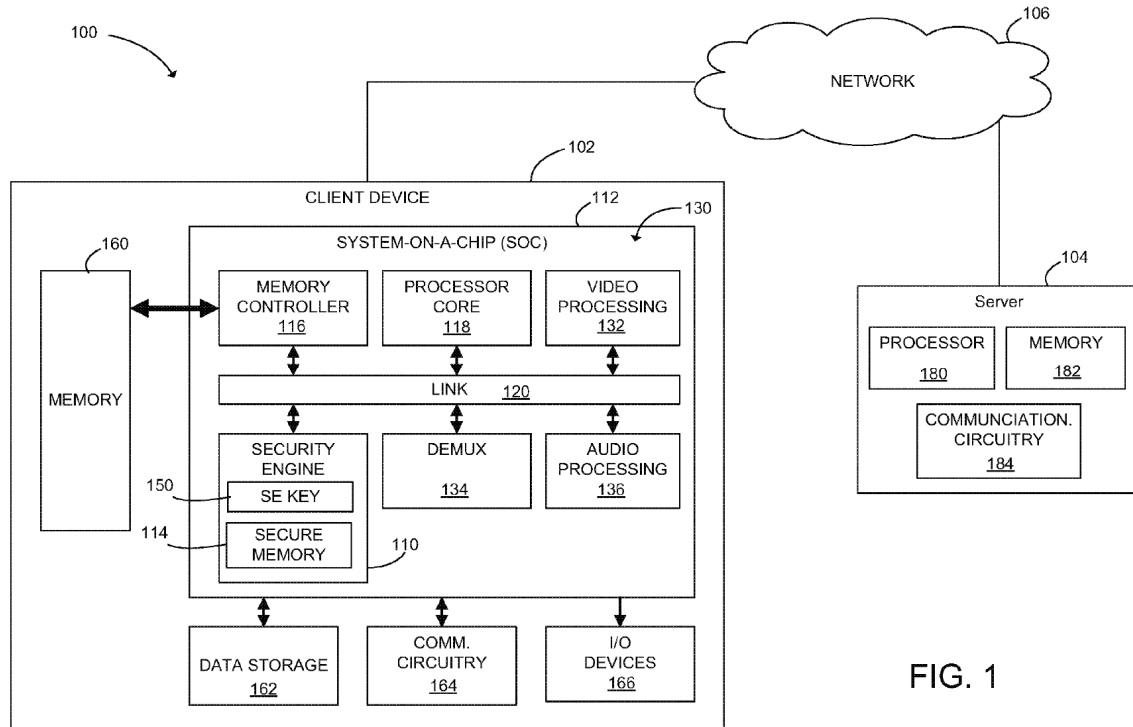


FIG. 1

'838 Patent, Fig. 1.

U.S. PATENT NOS. 10,705,960 and 10,725,919

43. On July 7, 2020, the United States Patent Office duly and legally issued the '960 Patent, entitled "Processors Having Virtually Clustered Cores and Cache Slices."⁹

44. On July 28, 2020, the United States Patent Office duly and legally issued the '919 Patent, entitled "Processors Having Virtually Clustered Cores and Cache Slices."¹⁰

45. Daedalus is the owner and assignee of all right, title, and interest in and to the '960 Patent and the '919 Patent, including the right to assert all causes of action arising under the '960 Patent and '919 Patent, and the right to sue and obtain any remedies for past, present, or future infringement.

46. The '960 Patent and the '919 Patent describe, among other things, a system comprising a plurality of processors each having one or more corresponding lower-level caches,

⁹ <https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/10705960>.

¹⁰ <https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/10725919>.

and a shared higher-level cache, which includes a plurality of distributed cache slices. '960 Patent, Abstract. The claimed processors include logic to direct an access that misses in one or more lower-level caches of a corresponding logical processor to a subset of the distributed cache slides in a virtual cluster that corresponds to the logical processor. *Id.* As the '960 Patent explains, “many processors now have multiple to many cores that are monolithically integrated on a single integrated circuit or die[,]” which “generally help to allow multiple threads or other workloads to be performed concurrently, which generally helps to increase execution throughput.” *Id.* at 1:26-31. “However, the multiple cores may have a downside in terms of longer hit and/or miss latencies to a shared cache. [...] In addition, the multiple or many cores also tend to increase the memory address entropy at memory controllers, which may tend to result in lower effective memory bandwidth.” *Id.* at 1:32-47. The inventions described and claimed in the '960 Patent overcome these challenges by providing novel processors with virtually clustered cores and cache slices, which has the effect of resulting in higher effective memory bandwidth.

47. The novel features of the invention are recited in the claims. For example, claim 15 of the '960 Patent recites:

15. A method comprising:

executing instructions and processing data with a plurality of cores, the plurality of cores comprising symmetric multi-threaded cores;

storing the instructions and the data in a cache subsystem, the cache subsystem comprising a plurality of first level caches and at least one higher level distributed cache comprising a plurality of distributed cache portions that are physically distributed across a die, each first level cache integral to one of the plurality of cores and each distributed cache portion accessible to each of the plurality of cores;

sharing the plurality of distributed cache portions among the plurality of cores;

providing coherent, non-uniform access to the plurality of distributed cache portions by the plurality of cores;

enabling a first frequency to be set for a first cluster of the plurality of cores which are physically proximate to one another and a second frequency to be set for a second cluster of the plurality of cores which are physically proximate to one another, wherein an average distance between cores in the first cluster is less than an average distance between the plurality of cores;

selectively gating power to the first cluster of the plurality of cores and distributed cache portions that correspond to the first cluster and/or the second cluster of the plurality of cores and distributed cache portions that correspond to the second cluster;

controlling access by the symmetric multi-threaded cores to a first system memory with a first integrated memory controller; and

controlling access by the symmetric multi-threaded cores to a second system memory with a second integrated memory controller.

'960 Patent, Cl. 15.

48. Further, claim 16 of the '919 Patent recites:

16. A method comprising:

executing instructions and processing data with a plurality of cores, the plurality of cores comprising symmetric multi-threaded cores;

storing the instructions and the data in a cache subsystem, the cache subsystem comprising a plurality of first-level caches and at least one higher-level distributed cache comprising a plurality of distributed cache portions that are physically distributed across a die, each first-level cache integral to one of the plurality of cores and each distributed cache portion accessible to each of the plurality of cores;

sharing the plurality of distributed cache portions among the plurality of cores;


providing coherent, non-uniform access to the plurality of distributed cache portions by the plurality of cores;

enabling a first frequency to be set for a first cluster of the plurality of cores which are physically proximate to one another and a second frequency to be set for a second cluster of the plurality of cores which are physically proximate to one another, wherein an average distance between cores in the first cluster is less than an average distance between all of the cores; and

selectively gating power to the first cluster of the plurality of cores and distributed cache portions of the at least one higher-level distributed cache that correspond to the first cluster and/or the second cluster of the plurality of cores and distributed cache portions of the at least one higher-level distributed cache that correspond to the second cluster.

'919 Patent, Cl. 16.

49. Figure 2 of the '960 Patent, reproduced below, shows a block diagram of an embodiment of a processor 201 having a first virtual cluster 215-1 and a second virtual cluster 215-2. The processor includes eighteen cores and eighteen corresponding cache slices. The cores/slices are coupled with first and second ring interconnects, which are coupled by a first inter-ring connection logic and a second inter-ring connection logic.



'960 Patent, Fig. 2.

MEDIATEK'S USE OF THE PATENTED TECHNOLOGY

50. According to its website, MediaTek is the world's 5th largest global fabless semiconductor company.¹¹ MediaTek powers more than 2 billion devices a year, which are in 20

¹¹ <https://www.mediatek.com/who-we-are>.

percent of homes and nearly 1 of every 3 mobile phones globally.¹² Upon information and belief, MediaTek's revenue in 2023 is approximately \$13 billion USD¹³.

51. On information and belief, MediaTek makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States various semiconductor chips which infringe the Asserted Patents. For example, MediaTek makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States the MediaTek Dimensity SoCs. As described in the counts below, these and other MediaTek products that include processors based on the ARMv8.2 architecture, as well as subsequent revisions to the ARM architecture such as the ARMv9 architecture, include power management, multiprocessor, cache and security technology that infringe the Asserted Patents.

FIRST COUNT

(Infringement of U.S Patent No. 8,769,316)

52. Daedalus incorporates by reference the allegations set forth in Paragraphs 1-51 of the Complaint as though fully set forth herein.

53. The claims of the '316 Patent are valid and enforceable.

54. On information and belief, in violation of 35 U.S.C. § 271(a), MediaTek has directly infringed and continues to directly infringe one or more claims of the '316 Patent, including at least Claim 8 of the '316 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '316 Patent, including but not limited to its electronic devices containing SoCs or microprocessors based on or derived from ARMv8.2 architecture, as well as subsequent revisions

¹² *Id.*; <https://www.poweredbymediatek.com/>.

¹³ <https://corp.mediatek.com/investor-relations/investor-relation-news/2023-q4-financial-results>.

to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9300 SoCs, and all reasonably similar products (the “’316 Patent Accused Products”).

55. Each of the ’316 Patent Accused Products implements a method comprising determining, in a power controller of a multi-domain processor, a power budget for the multi-domain processor for a current time interval, the multi-domain processor including at least a first domain and a second domain.

56. For example, SoCs or microprocessors derived from the ARMv8.2 architecture, as well as subsequent revisions to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9300, include logic such as ARM Power Policy Units that are configured by systems such as an ARM System Control Processor.¹⁴

57. Upon information and belief, Dimensity 9300 includes a Power Controller that determines a power budget for the multi-domain processor for a current time interval.¹⁵ The multi-domain processor includes a first domain comprising of the CPU and a second domain comprising of the GPU.¹⁶

58. Each of the ’316 Patent Accused Products implements a method comprising determining, in the power controller, a portion of the power budget to be allocated to the first and second domains, including allocating a minimum reservation value to the first domain and a minimum reservation value to the second domain, and sharing a remaining portion of the power budget according to a first sharing policy value for the first domain and a second sharing policy value for the second domain.

¹⁴ “Arm DynamIQ Shared Unit-110”, page 78, *available at* <https://documentation-service.arm.com/static/62bb28beb334256d9ea8cc32>; *id.*, page 51; *id.*, page 80; *id.*, page 77.

¹⁵ ARM, *High-level Considerations for Power Management of a big.LITTLE™ System*, Application Note 424 (2016), p. 15, <https://developer.arm.com/documentation/dai0424/a/I1007542>.

¹⁶ Xin Wang, *Intelligent Power Allocation, Maximize Performance in the Thermal Envelope*, ARM White Paper (March 2017), at pp. 11-14, <https://developer.arm.com/Tools%20and%20Software/Intelligent%20Power%20Allocation>.

59. For example, the Dimensity 9300 uses intellectual power allocation (IPA) to dynamically allocate power budget between two domains pursuant to a power allocation policy.¹⁷ Under the policy, “[e]ach cooling device is allocated with a share of the power budget, depending on the proportion of the device’s requested power in the total requested power;” and the extra power is allocated between the devices based on the weight for each device.¹⁸




Figure 8 IPA thermal management approach

To keep the system within the thermal envelope, IPA uses the PID controller to dynamically allocate power budget, and allows a short term boosting by exploiting thermal headroom.

IPA manages performance requests from different cooling devices. Each cooling device can request different performance levels, and has a power model to estimate its power consumption and the impact of the performance request for the cooling device.

The Power Arbiter provides guaranteed minimum performance. You can configure the Power Arbiter with policies to allocate power among different cooling devices.

¹⁷ *Id.*, *supra* note 8, pp. 12, 14-15.

¹⁸ *Id.* at pp. 14-15.

60. Further, on information and belief, MediaTek has actively induced and/or contributed to infringement of at least Claim 8 of the '316 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

61. Users of the '316 Patent Accused Products directly infringe at least Claim 8 of the '316 Patent when they use the '316 Patent Accused Products in the ordinary, customary, and intended way. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the '316 Patent Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the '316 Patent Accused Products to consumers within the United States and instructing and encouraging such customers to use the '316 Patent Accused Products in the ordinary, customary, and intended way, which MediaTek knew infringes at least Claim 8 of the '316 Patent, or, alternatively, was willfully blind to the infringement.

62. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the '316 Patent Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the '316 Patent Accused Products in the United States, which MediaTek knew infringes at least Claim 8 of the '316 Patent, or, alternatively, was willfully blind to the infringement.

63. On information and belief, in violation of 35 U.S.C. § 271(c), MediaTek's contributory infringement further includes offering to sell or selling within the United States, or

importing into the United States, components of the patented invention of at least Claim 8 of the '316 Patent, constituting a material part of the invention. On information and belief, MediaTek knows and has known the same to be especially made or especially adapted for use in an infringement of the '316 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

64. On information and belief, in violation of 35 U.S.C. § 271(f)(1), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 8 of the '316 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

65. On information and belief, in violation of 35 U.S.C. § 271(f)(2), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 8 of the '316 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

66. MediaTek is not licensed or otherwise authorized to practice the claims of the '316 Patent.

67. Thus, by its acts, MediaTek has injured Daedalus and is liable to Daedalus for directly and/or indirectly infringing one or more claims of the '316 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 8.

68. On information and belief, MediaTek has known about the '316 Patent at least since August 23, 2022.¹⁹ At a minimum, MediaTek has knowledge of the '316 Patent at least as of the filing of this Complaint. Accordingly, MediaTek's infringement of the '316 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

69. As a result of MediaTek's infringement of the '316 Patent, Daedalus has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for MediaTek's infringement, but in no event less than a reasonable royalty with interest and costs.

70. On information and belief, MediaTek will continue to infringe the '316 Patent unless enjoined by this Court. MediaTek's infringement of Daedalus' rights under the '316 Patent will continue to damage Daedalus, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

SECOND COUNT

(Infringement of U.S Patent No. 10,372,197)

71. Daedalus incorporates by reference the allegations set forth in Paragraphs 1-70 of the Complaint as though fully set forth herein.

72. The claims of the '197 Patent are valid and enforceable.

¹⁹ *Daedalus Prime LLC v. Mazda Motor Corp., et al.*, No. 22-cv-01108 (D. Del. Aug. 23, 2022).

73. On information and belief, in violation of 35 U.S.C. § 271(a), MediaTek has directly infringed and continues to directly infringe one or more claims of the '197 Patent, including at least Claim 1 of the '197 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '197 Patent, including but not limited to its electronic devices containing SoCs or microprocessors based on or derived from ARMv8.2 architecture, as well as subsequent revisions to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9300 SoCs, and all reasonably similar products (the "'197 Patent Accused Products").

74. Each of the '197 Patent Accused Products comprises a processor. For example, the Dimensity 9300 contains one or more microprocessors based on or derived from the ARM Cortex-X4 architecture and the ARM Cortex-A720 architecture.

75. Each of the '197 Patent Accused Products comprises a plurality of cores.

76. Specifically, the '197 Patent Accused Products include one or more clusters comprising a plurality of cores. For example, Dimensity 9300 SoCs comprise four Cortex-X4 and 4 Cortex A720 cores²⁰:

²⁰ <https://mediatek-marketing.files.svcdn.com/production/documents/Infographics/MediaTek-Dimensity-9300-Infographic.pdf?dm=1698856450>.



MediaTek Dimensity 9300

All Big Core CPU

World's first flagship smartphone chip to use all big cores for extreme performance.

- 4x Cortex-X4 CPU up to 3.25GHz
- 4x Cortex-A720 CPU up to 2.0GHz
- 15% increase in single-core performance
- 40% increase in multi-core performance

5G

MediaTek

Dimensity 9300

High-Performance 12-core GPU

Fast raytracing and console-level technologies on the Immortalis-G720 ensure incredibly smooth gameplay.

- 46% peak performance boost
- 40% increase in power savings
- 40% memory BW savings in geometry-heavy games
- 46% improvement in HW raytracing performance
- Console-level global illumination effects

Advantages in Power Efficiency

Precise CPU management for superior power efficiency.

- Up to 33% multi-core power saving vs previous gen CPU
- 3rd gen TSMC 4nm chip production
- 2nd gen thermally optimized IC design and package

Next-Gen Cinematic Video Capture

Our industry-leading video capture technologies just got even better.

- AI Semantic Analysis Video Engine
- 4K cinematic mode with depth of field
- 100% pixel-level autofocus + dual-lossless zoom
- 3-Microphone HDR audio recording

Generative AI Engine with Private, Personalized AI

New 7th Gen APU brings hardware-accelerated Generative AI into smartphones.

- 8x faster transformer-based generative AI
- 2x faster integer and floating-point compute improvement
- 45% more power efficient
- Up to 33 billion parameters
- Exclusive hardware-accelerated memory compression technology
- First to support on-device LoRA Fusion

Superior Security

Introducing a user privacy-focused security design and secure smartphone ecosystem.

- Secure Processor + HWRoT
- New Arm MTE Technology

AI-Display: Flagship Smart TV tech meets Smartphones

We are bringing many of our flagship technologies from the big screen to the small screen.

- World's first AI-depth engine display
- Support Google Ultra HDR
- 10% power savings with display low power technology

i.mediatek.com/dimensity-9300
Copyright 2023 © MediaTek, Inc. All rights reserved.

MEDIATEK

77. Each of the '197 Patent Accused Products comprises a cache memory.

78. For example, Dimensity 9300 SoCs comprise L1 and L2 cache memories²¹.

²¹ Arm® Cortex-X4 Core Technical Reference Manual, p. 41; Arm® Cortex-A720 Core Technical Reference Manual, p. 37.

Figure 3-1: Cortex-X4 core components

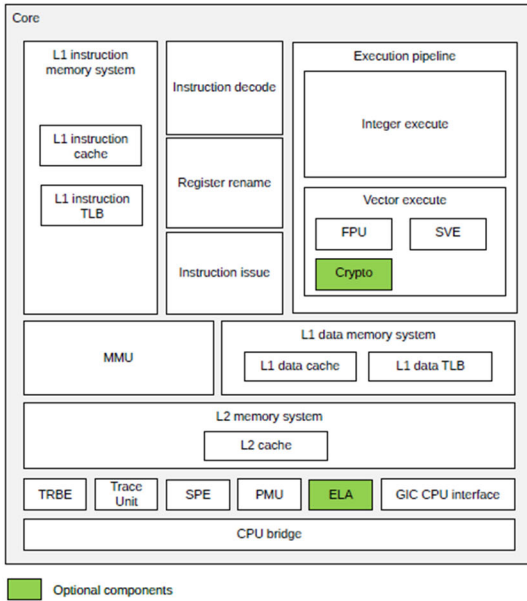
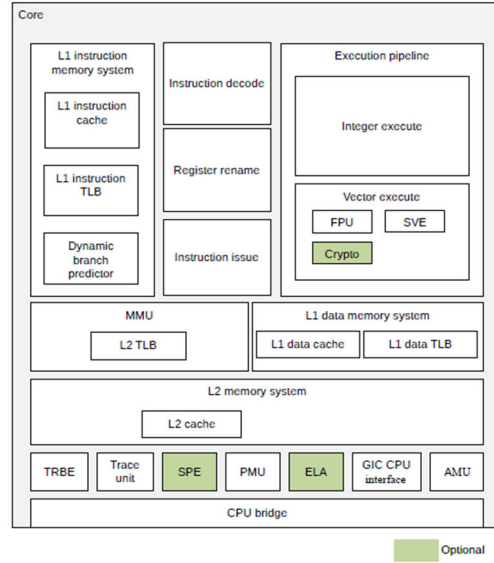


Figure 3-1: Cortex-A720 core components



79. Each of the '197 Patent Accused Products comprises an interconnect to couple the plurality of cores and the cache memory.


80. For example, the Dimensity 9300 SoCs include a DynamIQ Shared Unit (DSU). The DSU couples the plurality of cores to the L3 cache memory²².

²² Stefan Rosinger & Saurabh Pradhan, *Dimensity 9000 – A Flagship Smartphone SoC*, at p. 8, <https://hc34.hotchips.org/assets/program/conference/day2/Mobile%20and%20Edge/HC2022.Mediatek.EricbillWang.v08.pptx.pdf>.

CPU Highlights vs. Dimensity 1200

- Arm Cortex-X2 +40% integer performance over Arm Cortex-A78
- Arm Cortex-A510 +35% integer performance over Arm Cortex-A55
- -50% CPU power @ iso-performance
- Geekbenchv5 single-thread 1278 (+36%) , multi-thread 4400 (+33%)

Cortex-X2 1MB L2	Cortex-A7 10 512KB L2	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10
Cortex-A7 10 512KB L2	Cortex-A7 10 512KB L2	Shared 512KB L2\$ Shared VPU		Shared 512KB L2\$ Shared VPU	
DynamIQ Shared Unit equipped 8MB L3\$					


8

81. Each of the '197 Patent Accused Products comprises a power controller to control a plurality of power management features of the processor, wherein the power controller includes a tuning circuit to receive a workload configuration input regarding a workload, receive a plurality of energy performance bias (EPB) values and determine a global EPB value based thereon, and update at least one setting of at least one of the plurality of power management features based on the workload configuration input and the global EPB value.

82. For example, SoCs or microprocessors derived from the ARMv8.2 architecture, as well as subsequent revisions to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9300, include logic such as ARM Power Policy Units that are configured by systems such as an ARM System Control Processor.²³

83. On information and belief, the '197 Patent Accused Products use ARM's Intelligent Power Allocation technology in conjunction with the Power Policy units and a System Control Processor or Resource and Power Manger to receive a workload configuration input regarding a

²³ "Arm DynamIQ Shared Unit-110", page 78, available at <https://documentation-service.arm.com/static/62bb28beb334256d9ea8cc32>; *id.*, page 51; *id.*, page 80; *id.*, page 77.

workload, receive a plurality of energy performance bias (EPB) values and determine a global EPB value based thereon, and update at least one setting of at least one of the plurality of power management features based on the workload configuration input and the global EPB value.

84. For example, on information and belief, in the Dimensity 9300 SoCs, the Intelligent Power Allocation logic receives real-time CPU and GPU performance requests and based on the requested workload configuration and power models to cause settings of Power Policy Units to be updated to maximize requested performance without exceeding the Thermal Design Power for the SoC²⁴:

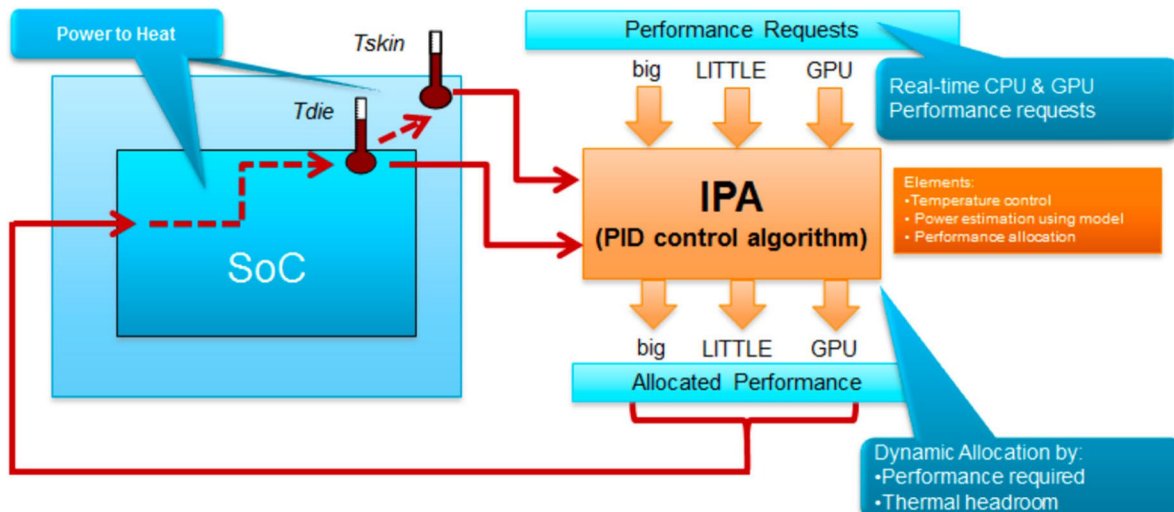


Figure 7 ARM Intelligent Power Allocation

85. Further, on information and belief, MediaTek has actively induced and/or contributed to infringement of at least Claim 1 of the '197 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

86. Users of the '197 Patent Accused Products directly infringe at least Claim 1 of the '197 Patent when they use the '197 Patent Accused Products in the ordinary, customary, and

²⁴ Wang, *supra* note 8, pp. 11-14.

intended way. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the '197 Patent Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the '197 Patent Accused Products to consumers within the United States and instructing and encouraging such customers to use the '197 Patent Accused Products in the ordinary, customary, and intended way, which MediaTek knew infringes at least Claim 1 of the '197 Patent, or, alternatively, was willfully blind to the infringement.

87. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the '197 Patent Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the '197 Patent Accused Products in the United States, which MediaTek knew infringes at least Claim 1 of the '197 Patent, or, alternatively, was willfully blind to the infringement.

88. On information and belief, in violation of 35 U.S.C. § 271(c), MediaTek's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of at least Claim 1 of the '197 Patent, constituting a material part of the invention. On information and belief, MediaTek knows and has known the same to be especially made or especially adapted for use in an infringement of the '197 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

89. On information and belief, in violation of 35 U.S.C. § 271(f)(1), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 1 of the '197 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

90. On information and belief, in violation of 35 U.S.C. § 271(f)(2), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 1 of the '197 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

91. MediaTek is not licensed or otherwise authorized to practice the claims of the '197 Patent.

92. Thus, by its acts, MediaTek has injured Daedalus and is liable to Daedalus for directly and/or indirectly infringing one or more claims of the '197 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 1.

93. On information and belief, MediaTek has known about the '197 Patent at least since August 23, 2022.²⁵ At a minimum, MediaTek has knowledge of the '197 Patent at least as of the filing of this Complaint. Accordingly, MediaTek's infringement of the '197 Patent has been and

²⁵ *Daedalus Prime LLC v. Mazda Motor Corp., et al.*, No. 22-cv-01108 (D. Del. Aug. 23, 2022).

continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

94. As a result of MediaTek's infringement of the '197 Patent, Daedalus has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for MediaTek's infringement, but in no event less than a reasonable royalty with interest and costs.

95. On information and belief, MediaTek will continue to infringe the '197 Patent unless enjoined by this Court. MediaTek's infringement of Daedalus' rights under the '197 Patent will continue to damage Daedalus, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

THIRD COUNT

(Infringement of U.S Patent No. 10,740,281)

96. Daedalus incorporates by reference the allegations set forth in Paragraphs 1-95 of the Complaint as though fully set forth herein.

97. The claims of the '281 Patent are valid and enforceable.

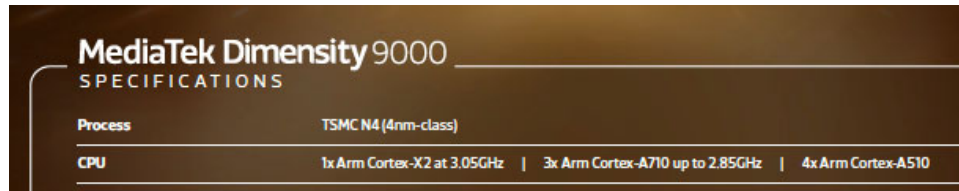
98. On information and belief, in violation of 35 U.S.C. § 271(a), MediaTek has directly infringed and continues to directly infringe one or more claims of the '281 Patent, including at least Claim 8 of the '281 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '281 Patent, including but not limited to its electronic devices containing SoCs or microprocessors based on or derived from ARMv8.2 architecture, as well as subsequent revisions

to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9000 SoCs, and all reasonably similar products (the “’281 Patent Accused Products”).

99. Each of the ’281 Patent Accused Products comprises a multi-core processor. For example, the Dimensity 9000 contains one or more microprocessors based on or derived from the ARM Cortex-X2 architecture, the ARM Cortex-A710 architecture and the ARM Cortex-A510 architecture.

100. Each of the ’281 Patent Accused Products comprises a first plurality of cores and a second plurality of cores.

101. Specifically, the ’281 Patent Accused Products include one or more clusters comprising a plurality of cores. For example, Dimensity 9000 SoCs comprise three Cortex-A710 cores and four Cortex-A510 cores²⁶:



The infographic displays the specifications for the MediaTek Dimensity 9000. It features a dark background with white text. The title 'MediaTek Dimensity 9000' is prominently displayed at the top, followed by 'SPECIFICATIONS' in a smaller font. Below this, there are two rows of specifications, each with a label on the left and the corresponding details on the right. The first row is labeled 'Process' and specifies 'TSMC N4 (4nm-class)'. The second row is labeled 'CPU' and specifies '1x Arm Cortex-X2 at 3.05GHz | 3x Arm Cortex-A710 up to 2.85GHz | 4x Arm Cortex-A510'.

MediaTek Dimensity 9000	
SPECIFICATIONS	
Process	TSMC N4 (4nm-class)
CPU	1x Arm Cortex-X2 at 3.05GHz 3x Arm Cortex-A710 up to 2.85GHz 4x Arm Cortex-A510


²⁶ <https://mediatek-marketing.files.svcdn.com/production/documents/Dimensity-9000-Infographic.pdf>; <https://i.mediatek.com/dimensity-9000>; Nayak, et al., *supra* note 2.

The ARM logo is displayed in a large, light grey, lowercase font.

WORLD'S 1ST CORTEX-X2 IN A SMARTPHONE CHIP

The Dimensity 9000 uses new Armv9 architecture CPUs and GPU to deliver unparalleled performance. Its octa-core CPU includes an Arm Cortex-X2 that bursts to epic 3GHz, while new LPDDR5X memory makes data immediately available, eliminating the wait to give immediate responsiveness in any app, whatever you're doing.

- Ultra-Core - 1x Arm Cortex-X2 at 3.05GHz
- Super-Cores - 3x Arm Cortex-A710 up to 2.85GHz
- Efficiency Cores - 4x Arm Cortex-A510
- World's first Arm Mali-G710 MC10 graphics processor
- Big caches - 8MB L3 cache + 6MB system cache
- LPDDR5X 7500Mbps support – 20% more power efficient than LPDDR5



102. Each of the '281 Patent Accused Products comprises a first plurality of cores and a second plurality of cores that support a same instruction set.

103. For example, Dimensity 9000 SoCs comprise clusters of Cortex-A710 cores and Cortex-A510 cores, all supporting the ARMv9 instruction set²⁷:

2 The Cortex®-A510 core

The Cortex®-A510 core is a high-efficiency, low-power product that implements the Arm®v9.0-A architecture. The Arm®v9.0-A architecture extends the architecture defined in the Arm®v8-A architectures up to Arm®v8.5-A.

²⁷ Arm® Cortex®-A510 Core Technical Reference Manual, p. 22-23; Arm® Cortex®-A710 Core Technical Reference Manual, p. 22-23; Nayak, et al., *supra* note 2.

2.1 Cortex®-A510 core features

The Cortex®-A510 core might be used in standalone DynamIQ™ configurations where a homogenous DSU-110 DynamIQ™ cluster includes one to eight Cortex®-A510 cores. The Cortex®-A510 core might also be used as a high efficiency core or a high-performance core in a heterogenous DSU-110 DynamIQ™ cluster.

However, regardless of the cluster configuration, the Cortex®-A510 core always has the same features.

Core features

- Implementation of the Arm®v9.0-A A64 instruction set

2 The Cortex®-A710 core

The Cortex®-A710 core is a high-performance, low-power, and constrained area product that implements the Arm®v9.0-A architecture. The Arm®v9.0-A architecture extends the architecture defined in the Armv8-A architectures up to Arm®v8.5-A. The Cortex®-A710 core targets clamshell and premium high-end smartphone applications.

2.1 Cortex®-A710 core features

The Cortex®-A710 core might be used in standalone DynamIQ™ configurations, that is in a homogenous cluster of one to four Cortex®-A710 cores. It might also be used either as the high-performance or balanced-performance core in a heterogenous cluster.

However, regardless of the cluster configuration, the Cortex®-A710 core always has the same features.

Core features

- Implementation of the Armv9-A A32, T32, and A64 instruction sets

The heterogeneous CPU complex, shown in Fig. 2.5.2, is organized into 3 gears. The 1st gear is a single HP core which utilizes the ARMv9 Cortex-X2 microarchitecture with 64KB L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.

All processor cores in the CPU subsystem incorporate the ARMv9 instruction set with key architectural advances. Memory Tagging Extension (MTE) enables greater security by locking data in the memory using a tag which can only be accessed by the correct key held by the pointer accessing the memory location, as shown in Fig. 2.5.1. Further, a Scalable Vector Extension 2 (SVE2) allows a scalable vector length in multiples of 128b, up to 2048b, enabling increased DSP and ML vector-processing capabilities, as shown in Fig. 2.5.1.

104. Each of the '281 Patent Accused Products comprises a first plurality of cores that are higher performance and consume more power than a second plurality of cores.

105. For example, the Dimensity 9000 SoCs include a plurality of “Balanced Performance (BP)” Cortex-A710 cores and a plurality of “High Efficiency (HE)” Cortex-A510 cores., wherein the Cortex-A710 cores are higher performance and consume more power than the Cortex-A510 cores²⁸:

The heterogeneous CPU complex, shown in Fig. 2.5.2, is organized into 3 gears. The 1st gear is a single HP core which utilizes the ARMv9 Cortex-X2 microarchitecture with 64KB L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.

The HP core runs up to 3.4GHz clock speed to meet high-speed compute demands, while the HE cores are optimized to operate efficiently at ultra-low voltage. The BP cores provide a balance of power and performance for average workloads. Depending on the dynamic computing needs, workloads can be seamlessly switched and assigned across different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.

²⁸ Nayak, et al., *supra* note 2; Aditya Bedi, *The Foundation of Total Compute: First Armv9 Cortex CPUs* (May 25, 2021), <https://community.arm.com/arm-community-blogs/b/architectures-and-processors-blog/posts/first-armv9-cpu-cores>.



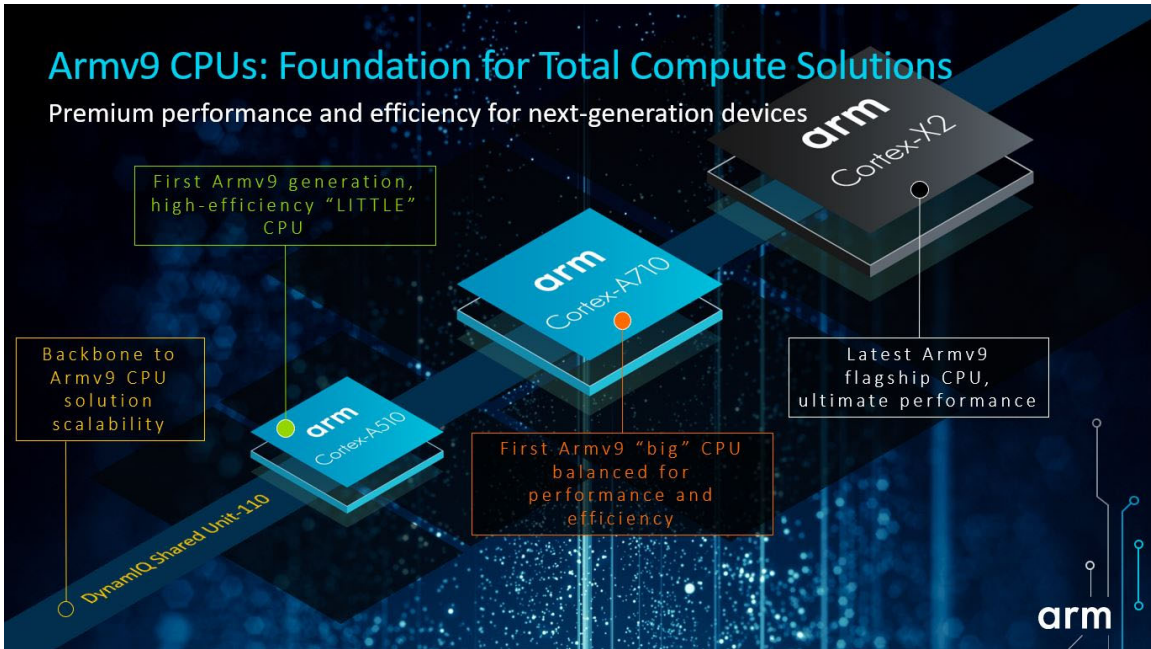


Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.



Core	Focus	Microarchitectural Highlights
Cortex-X2	Peak Performance	Branch prediction : Decoupled from fetch, improved accuracy and prediction SVE 128b vector length implementation bfloat16, int8 matmul data types support for ML apps Reduced a pipeline stage at dispatch Larger Out of order window size, load-store window/structures. d-TLB Data prefetch enhancements for accuracy and coverage
A710	Sustained Performance	Branch prediction : improved accuracy, capacity doubled for BTB, GHB Increased capacity of L1 instruction TLB Improved efficiency with reduction of mid core width Reduced a pipeline stage at dispatch Reduced DSU access, DRAM refills Data prefetch enhancements for accuracy and coverage
A510	Power Efficiency	Two cores grouped in a complex, multiple complexes per cluster L2 cache, L2 TLB, vector paths data path shared across a complex Fine-grained scheduling between the cores in the complex 3-wide decode and issue, 3 integer ALU pipeline state of art multi-stage branch prediction Advance data prefetchers for accuracy and coverage

Figure 2.5.2: ARMv9 CPU cluster.



Cortex Application Processors: Design Philosophies

The right CPU for the right workload

Cortex-X series	Cortex-A700 series	Cortex-A500 series
<ul style="list-style-type: none"> Maximize Performance on Single-Thread / Burst Workloads Pipeline structured and provisioned to push IPC 	<ul style="list-style-type: none"> Prioritize Sustained Multi-Core Workloads Best Balance of Efficiency and Performance 	<ul style="list-style-type: none"> Efficiency first on background and lightweight Workloads In-order design with "big-core" inspired prefetch & prediction

Total Compute solutions | Next-Gen Architecture Support
Focus on use-case driven cluster-level workloads and configurations

15 © 2021 Arm Limited

We are also announcing the [Arm Cortex-A710](#). This is our first Armv9 generation "big" CPU, with it providing the best balance of performance and efficiency. Accompanying the "big" Cortex-A710 is the first Armv9 high efficiency "LITTLE" CPU, the [Arm Cortex-A510](#), which is the successor to the highly popular [Arm Cortex-A55 CPU](#).

106. Each of the '281 Patent Accused Products comprises a second plurality of cores that have a maximum operating frequency that is less than the maximum operating frequency of each of the first plurality of cores.

107. For example, the Dimensity 9000 SoCs include a plurality of Cortex-A510 cores, which have a maximum operating frequency that is less than the maximum operating frequency of the plurality of Cortex-A710 cores.²⁹


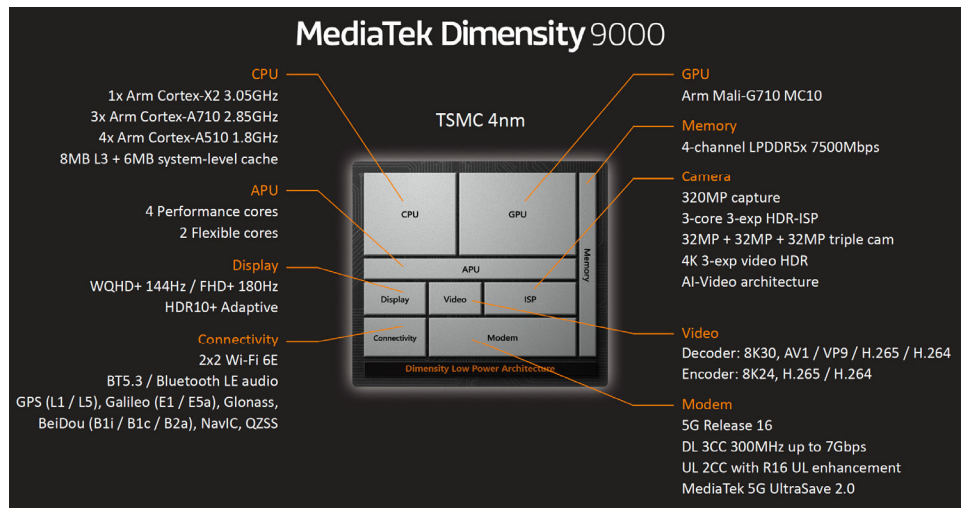



Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

²⁹ Rosinger & Pradhan, *supra* note 14, p 3; Nayak, et al., *supra* note 2; <https://i.mediatek.com/dimensity-9000;> <https://mediatek-marketing.files.svdcdn.com/production/documents/Dimensity-9000-Infographic.pdf>.

MediaTek Dimensity 9000	
SPECIFICATIONS	
Process	TSMC N4 (4nm-class)
CPU	1x Arm Cortex-X2 at 3.05GHz 3x Arm Cortex-A710 up to 2.85GHz 4x Arm Cortex-A510



WORLD'S 1ST CORTEX-X2 IN A SMARTPHONE CHIP

The Dimensity 9000 uses new Armv9 architecture CPUs and GPU to deliver unparalleled performance. Its octa-core CPU includes an Arm Cortex-X2 that bursts to epic 3GHz, while new LPDDR5X memory makes data immediately available, eliminating the wait to give immediate responsiveness in any app, whatever you're doing.

- Ultra-Core - 1x Arm Cortex-X2 at 3.05GHz
- Super-Cores - 3x Arm Cortex-A710 up to 2.85GHz
- Efficiency Cores - 4x Arm Cortex-A510
- World's first Arm Mali-G710 MC10 graphics processor
- Big caches - 8MB L3 cache + 6MB system cache
- LPDDR5X 7500Mbps support – 20% more power efficient than LPDDR5

108. Each of the '281 Patent Accused Products comprises a caching layer shared by the first plurality of cores and the second plurality of cores.

109. For example, the Dimensity 9000 SoCs include a DynamIQ Shared Unit (DSU). The DSU couples the plurality of cores to the shared L3 cache memory³⁰.

³⁰ Rosinger & Pradhan, *supra* note 14, pp. 3, 8, 11; Nayak, et al., *supra* note 2.

CPU

1x Arm Cortex-X2 3.05GHz
 3x Arm Cortex-A710 2.85GHz
 4x Arm Cortex-A510 1.8GHz
8MB L3 + 6MB system-level cache

CPU Highlights vs. Dimensity 1200

- Arm Cortex-X2 +40% integer performance over Arm Cortex-A78
- Arm Cortex-A510 +35% integer performance over Arm Cortex-A55
- -50% CPU power @ iso-performance

- Geekbenchv5 single-thread 1278 (+36%) , multi-thread 4400 (+33%)



Cortex-X2 1MB L2	Cortex-A7 10 512KB L2	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10
Cortex-A7 10 512KB L2	Cortex-A7 10 512KB L2	Shared 512KB L2\$ Shared VPU		Shared 512KB L2\$ Shared VPU	
DynamIQ Shared Unit equipped 8MB L3\$					



8

Arm DynamIQ Shared Unit-110

- DSU-110 for Cache Coherency & Shared L3\$
- Optimized Ring Transport Network
 - Bi-directional, Dual-ring
- 2X Bandwidth and 25% Lower Leakage
 - Support Partial SRAM & Logic Shutdown
- Cache Partition for QoS

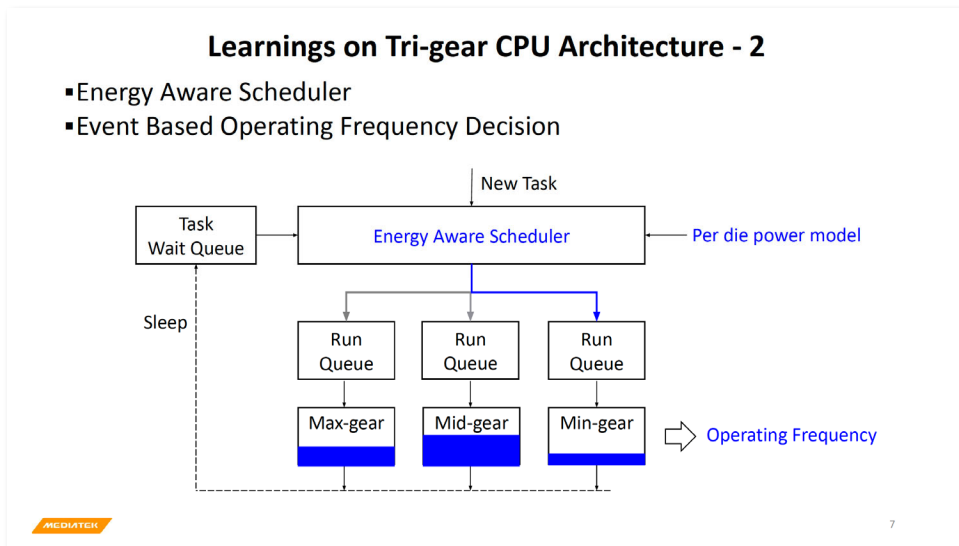


11

The heterogeneous CPU complex, shown in Fig. 2.5.2, is organized into 3 gears. The 1st gear is a single HP core which utilizes the ARMv9 Cortex-X2 microarchitecture with 64KB L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.


110. Each of the '281 Patent Accused Products comprises power management hardware to enable and disable the first plurality of cores and the second plurality of cores, wherein an operating system that executes on the multi-core processor is to monitor a demand for the multi-core processor and control a core mix of the first plurality of cores and the second plurality of cores based on the demand with the power management hardware.

111. For example, the Dimensity 9000 SoCs include an Energy Aware Scheduler and other logic such as ARM Power Policy Units that are configured by systems such as an ARM System Control Processor.³¹



³¹ Rosinger & Pradhan, *supra* note 14, p. 7; Nayak, et al., *supra* note 2.

▪ Usable Tri-gear Power Efficiency Curve



MEDIATEK

The HP core runs up to 3.4GHz clock speed to meet high-speed compute demands, while the HE cores are optimized to operate efficiently at ultra-low voltage. The BP cores provide a balance of power and performance for average workloads. Depending on the dynamic computing needs, workloads can be seamlessly switched and assigned across different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.




Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

In summary, a tri-gear ARMv9 CPU subsystem for a flagship 5G smartphone SoC is introduced with a high-performance core achieving 3.4GHz at robust yield and delivering up to 27% higher peak performance through microarchitectural and implementation advancements. Further, circuit innovation to enable continuous monitoring of on-die power supply is shown. Finally, a new variation aware AVS technology is introduced to further improve CPU power efficiency.

112. On information and belief, the '281 Patent Accused Products use ARM's Intelligent Power Allocation technology in conjunction with the Power Policy Units and a System Control Processor or Resource and Power Manager to monitor demand and enable and disable the first plurality of cores and the second plurality of cores.³²

³² "Arm DynamIQ Shared Unit-110" at 36-37.

Power management and Power Policy Units

The DynamIQ™ cluster shared logic integrates several *Power management and Power Policy Units* (PPUs) to control power modes and resets. The PPU can be programmed to directly select a

Copyright © 2019–2022 Arm Limited (or its affiliates). All rights reserved.
Non-Confidential

Page 36 of 862

Arm® DynamIQ™ Shared Unit-110 Technical Reference
Manual

Document ID: 101381_0400_11_en
Issue: 11
Technical overview

specific power mode or can be programmed to autonomously switch between power modes within a specified range, based on the requirements of the cluster. The PPU can be programmed from your *System Control Processor (SCP)* using the utility bus to access them.

113. For example, on information and belief, in the Dimensity 9000 SoCs, the Power Policy Units are utilized to enable and disable the first plurality of cores and the second plurality of cores³³:

5.1 Power management in the DSU-110

The *DynamIQ™ Shared Unit-110* (DSU-110) provides various mechanisms to control both dynamic and static power dissipation. These mechanisms are associated with a set of power domains, power modes, and operational modes. Some of these mechanisms are brought under software control using *Power Policy Units* (PPUs).

The power management techniques employed by the DSU-110 and cores in the cluster include:

- Internal core clock gating where different internal parts of the core are clock idle
- Per-core *Dynamic Voltage and Frequency Scaling* (DVFS)
- Powerdown of components of the cluster which can include:
 - Cores
 - All the L3 cache or parts of the L3 cache. See [5.4.1 L3 cache RAM powerdown](#) on page 57 and [5.4.2 L3 cache slice powerdown](#) on page 61.
 - Retention which is a low-power mode that retains the register and RAM state. Retention can be applied to the following components of the cluster:
 - Cache RAMs in the cores
 - All of the L3 cache or parts of the L3 cache

³³ *Id.* at 50.

114. Further, on information and belief, MediaTek has actively induced and/or contributed to infringement of at least Claim 8 of the '281 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

115. Users of the '281 Patent Accused Products directly infringe at least Claim 8 of the '281 Patent when they use the '281 Patent Accused Products in the ordinary, customary, and intended way. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the '281 Patent Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the '281 Patent Accused Products to consumers within the United States and instructing and encouraging such customers to use the '281 Patent Accused Products in the ordinary, customary, and intended way, which MediaTek knew infringes at least Claim 8 of the '281 Patent, or, alternatively, was willfully blind to the infringement.

116. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the '281 Patent Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the '281 Patent Accused Products in the United States, which MediaTek knew infringes at least Claim 8 of the '281 Patent, or, alternatively, was willfully blind to the infringement.

117. On information and belief, in violation of 35 U.S.C. § 271(c), MediaTek's contributory infringement further includes offering to sell or selling within the United States, or

importing into the United States, components of the patented invention of at least Claim 8 of the '281 Patent, constituting a material part of the invention. On information and belief, MediaTek knows and has known the same to be especially made or especially adapted for use in an infringement of the '281 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

118. On information and belief, in violation of 35 U.S.C. § 271(f)(1), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 8 of the '281 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

119. On information and belief, in violation of 35 U.S.C. § 271(f)(2), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 8 of the '281 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

120. MediaTek is not licensed or otherwise authorized to practice the claims of the '281 Patent.

121. Thus, by its acts, MediaTek has injured Daedalus and is liable to Daedalus for directly and/or indirectly infringing one or more claims of the '281 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 8.

122. On information and belief, MediaTek has known about the '281 Patent at least since August 23, 2022.³⁴ At a minimum, MediaTek has knowledge of the '281 Patent at least as of the filing of this Complaint. Accordingly, MediaTek's infringement of the '281 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

123. As a result of MediaTek's infringement of the '281 Patent, Daedalus has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for MediaTek's infringement, but in no event less than a reasonable royalty with interest and costs.

124. On information and belief, MediaTek will continue to infringe the '281 Patent unless enjoined by this Court. MediaTek's infringement of Daedalus' rights under the '281 Patent will continue to damage Daedalus, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

FOURTH COUNT

(Infringement of U.S Patent No. 8,984,228)

125. Daedalus incorporates by reference the allegations set forth in Paragraphs 1-124 of the Complaint as though fully set forth herein.

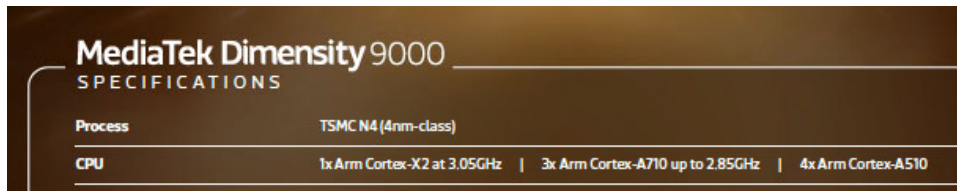
126. The claims of the '228 Patent are valid and enforceable.

³⁴ *Daedalus Prime LLC v. Mazda Motor Corp., et al.*, No. 22-cv-01108 (D. Del. Aug. 23, 2022).

127. On information and belief, in violation of 35 U.S.C. § 271(a), MediaTek has directly infringed and continues to directly infringe one or more claims of the '228 Patent, including at least Claims 1 and 11 of the '228 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '228 Patent, including but not limited to its electronic devices containing SoCs or microprocessors based on or derived from ARMv8.2 architecture, as well as subsequent revisions to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9000 SoCs, and all reasonably similar products (the "'228 Patent Accused Products").

128. Each of the '228 Patent Accused Products comprises a multi-core processor including a plurality of cores. For example, the Dimensity 9000 contains one or more microprocessors based on or derived from the ARM Cortex-X2 architecture, the ARM Cortex-A710 architecture and the ARM Cortex-A510 architecture.

129. Specifically, the '228 Patent Accused Products include one or more clusters comprising a plurality of cores. For example, Dimensity 9000 SoCs comprise a Cortex-X2 core, three Cortex-A710 cores and four Cortex-A510 cores³⁵:



The infographic displays the specifications for the MediaTek Dimensity 9000 processor. It features a dark background with white text. The title 'MediaTek Dimensity 9000' is prominently displayed at the top, followed by 'SPECIFICATIONS'. Below this, two rows of specifications are listed: 'Process' (TSMC N4 (4nm-class)) and 'CPU' (1x Arm Cortex-X2 at 3.05GHz | 3x Arm Cortex-A710 up to 2.85GHz | 4x Arm Cortex-A510).

MediaTek Dimensity 9000 SPECIFICATIONS	
Process	TSMC N4 (4nm-class)
CPU	1x Arm Cortex-X2 at 3.05GHz 3x Arm Cortex-A710 up to 2.85GHz 4x Arm Cortex-A510

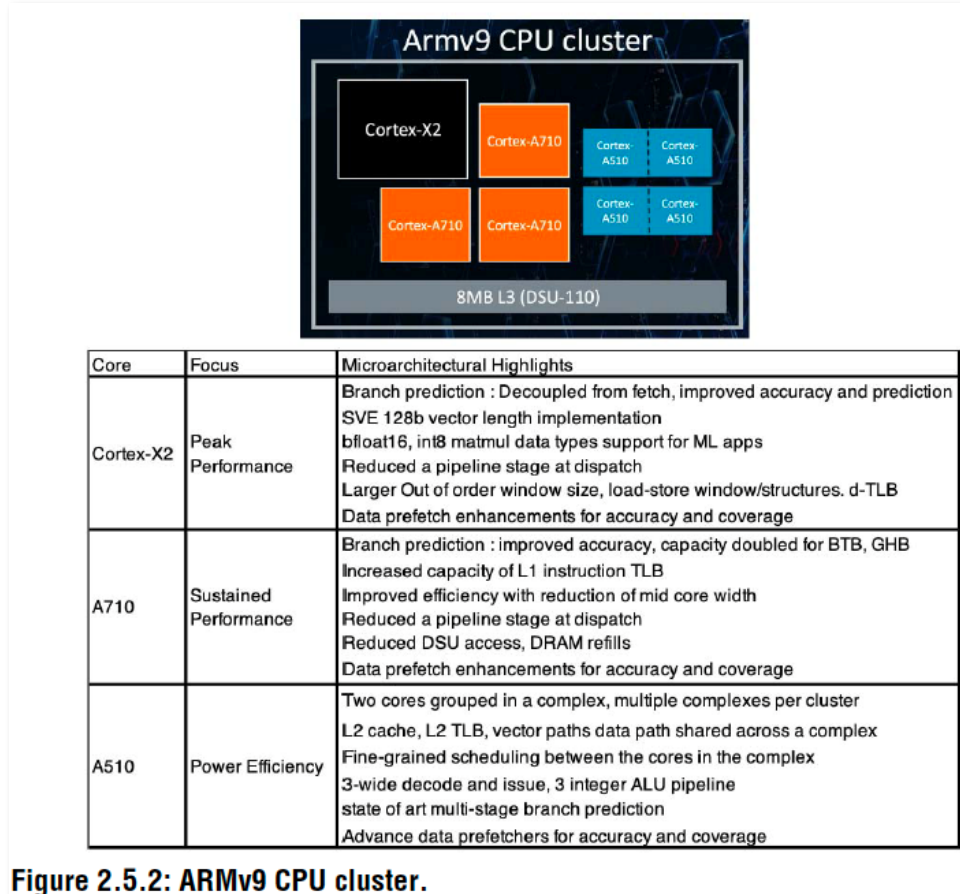
³⁵ <https://mediatek-marketing.files.svcdn.com/production/documents/Dimensity-9000-Infographic.pdf>; <https://i.mediatek.com/dimensity-9000>; Nayak, et al., *supra* note 2.



WORLD'S 1ST CORTEX-X2 IN A SMARTPHONE CHIP

The Dimensity 9000 uses new Armv9 architecture CPUs and GPU to deliver unparalleled performance. Its octa-core CPU includes an Arm Cortex-X2 that bursts to epic 3GHz, while new LPDDR5X memory makes data immediately available, eliminating the wait to give immediate responsiveness in any app, whatever you're doing.

- Ultra-Core - 1x Arm Cortex-X2 at 3.05GHz
- Super-Cores - 3x Arm Cortex-A710 up to 2.85GHz
- Efficiency Cores - 4x Arm Cortex-A510
- World's first Arm Mali-G710 MC10 graphics processor
- Big caches - 8MB L3 cache + 6MB system cache
- LPDDR5X 7500Mbps support – 20% more power efficient than LPDDR5



130. Each of the '228 Patent Accused Products comprises a shared cache memory.

131. For example, the Dimensity 9000 SoCs include a DynamIQ Shared Unit (DSU).

The DSU couples the plurality of cores to the shared cache memory³⁶.

CPU

1x Arm Cortex-X2 3.05GHz
 3x Arm Cortex-A710 2.85GHz
 4x Arm Cortex-A510 1.8GHz
8MB L3 + 6MB system-level cache

³⁶ Rosinger & Pradhan, *supra* note 14, pp. 3, 8, 11; Nayak, et al., *supra* note 2.

CPU Highlights vs. Dimensity 1200

- Arm Cortex-X2 +40% integer performance over Arm Cortex-A78
- Arm Cortex-A510 +35% integer performance over Arm Cortex-A55
- -50% CPU power @ iso-performance
- Geekbenchv5 single-thread 1278 (+36%) , multi-thread 4400 (+33%)



Cortex-X2 1MB L2	Cortex-A7 10 512KB L2	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10
Cortex-A7 10 512KB L2	Cortex-A7 10 512KB L2	Shared 512KB L2\$ Shared VPU		Shared 512KB L2\$ Shared VPU	
DynamIQ Shared Unit equipped 8MB L3\$					



8

Arm DynamIQ Shared Unit-110

- DSU-110 for Cache Coherency & Shared L3\$
- Optimized Ring Transport Network
 - Bi-directional, Dual-ring
- 2X Bandwidth and 25% Lower Leakage
 - Support Partial SRAM & Logic Shutdown
- Cache Partition for QoS



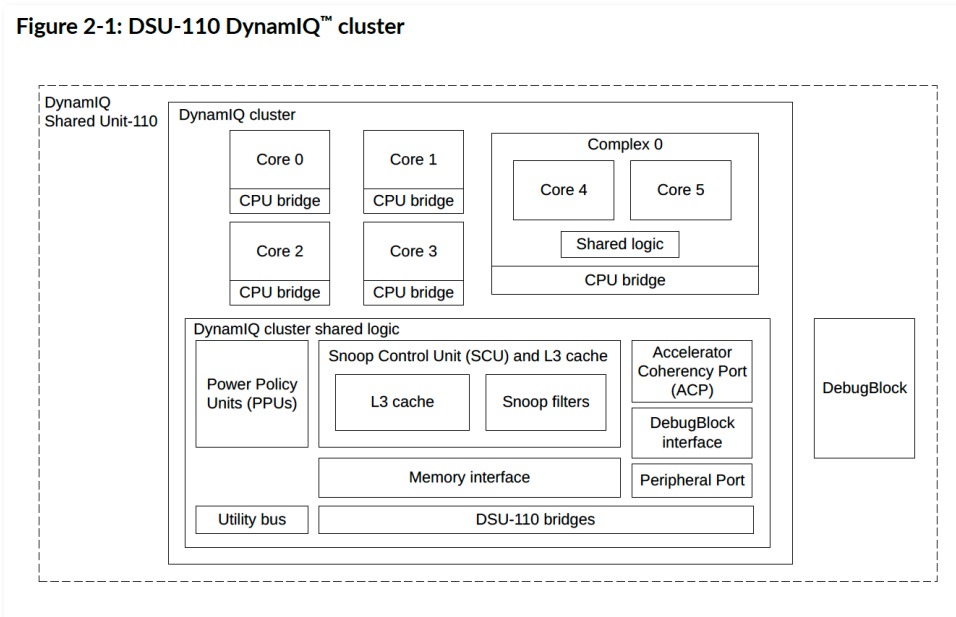
11

The heterogeneous CPU complex, shown in Fig. 2.5.2, is organized into 3 gears. The 1st gear is a single HP core which utilizes the ARMv9 Cortex-X2 microarchitecture with 64KB L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.

132. Each of the '228 Patent Accused Products comprises an integrated input/output (IIO) module to interface between the multicore processor and at least one IO device coupled to the multicore processor.

133. For example, the Dimensity 9000 SoCs include a DynamIQ™ Shared Unit-110 (DSU-110) that provides a shared L3 memory system, snoop control and filtering, and other control logic to support a cluster of A-class architecture cores and that contains external interfaces including those to the cores.³⁷

Figure 2-1: DSU-110 DynamIQ™ cluster



³⁷ “Arm DynamIQ Shared Unit-110” at 17-18; Nayak, et al., *supra* note 2.

All cores in the DSU-110 DynamIQ™ cluster, including those in complexes, are coherently connected to an L3 memory system that includes an L3 cache and a *Snoop Control Unit* (SCU). The SCU maintains coherency between caches in the cores and the L3 cache, and includes a snoop filter to optimize coherency maintenance operations. The shared L3 cache simplifies process migration between the cores.

The DSU-110 DynamIQ™ cluster can be implemented with various power domains to target power performance levels. These power domains are managed through the *Power Policy Units* (PPUs). The DSU-110 DynamIQ™ cluster supports many mechanisms to reduce static and dynamic power dissipation. For example, placing the cores and L3 cache into retention and powering down parts of the L3 cache.

All the external interfaces including those to the cores are provided through the DSU-110 to the *System on Chip* (SoC). Main system transactions are supported through the memory interface which can be implemented as a coherent or non-coherent interface. A peripheral port is provided to support low latency access to external system components but also can be used as a non-coherent master interface. The *Accelerator Coherency Port* (ACP) provides coherent access for non-cached masters that need I/O coherency with the cluster. The utility bus is a memory-mapped port that provides a programming interface to the PPU's and some of the other system components.

This paper presents a tri-gear ARMv9 CPU subsystem incorporated in a 5G flagship mobile SoC. Implemented in a 5nm technology node, a 3.4GHz High-Performance (HP) core is introduced along with circuit and implementation techniques to achieve CPU PPA targets. A die photograph is shown in Fig. 2.5.7. The SoC integrates a 5G modem supporting NR sub-6GHz with downlink and uplink speed up to 7.01Gb/s and 2.5Gb/s, respectively, an ARMv9 CPU subsystem, an ARM Mali G710 GPU for 3D graphics, an in-house Vision Processing Unit (VPU), and a Deep-Learning Accelerator (DLA) for high-performance and power-efficient AI processing. The integrated display engine can provide portrait panel resolution up to QHD+ 21:9 (1600×3360) and frame rates up to 144Hz. Multimedia and imaging subsystems decode 8K video at 30fps, while encoding 4K video at 60fps; camera resolutions up to 320MPixels are supported. LPDDR5-6400/LPDDR5X-7500 memory interfaces facilitate up to 24GB of external SDRAM over four 16b channels for a peak transfer rate of 0.46Tb/s.

134. On information and belief, the '228 Patent Accused Products contain ARM's DSU-110 that provides an input/output (IIO) module to interface between the multicore processor and at least one IO device.³⁸

³⁸ Rosinger & Pradhan, *supra* note 14, p. 8.

CPU Highlights vs. Dimensity 1200

- Arm Cortex-X2 +40% integer performance over Arm Cortex-A78
- Arm Cortex-A510 +35% integer performance over Arm Cortex-A55
- -50% CPU power @ iso-performance
- Geekbenchv5 single-thread 1278 (+36%) , multi-thread 4400 (+33%)

Cortex-X2 1MB L2	Cortex-A7 10 512KB L2	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10
Cortex-A7 10 512KB L2	Cortex-A7 10 512KB L2	Shared 512KB L2\$ Shared VPU		Shared 512KB L2\$ Shared VPU	
DynamIQ Shared Unit equipped 8MB L3\$					

MEDIATEK

8

135. Each of the '228 Patent Accused Products comprises a caching agent to perform cache coherency operations for the plurality of cores and the IIO module, the caching agent a single caching agent for the multicore processor and including a plurality of distributed portions each associated with a corresponding one of the plurality of cores.

136. For example, on information and belief, in the Dimensity 9000 SoCs, the DSU-110 provides coherency features, including a snoop control unit, and coherent bus interfaces³⁹:

Coherency and snoop control

The DSU-110 has the following coherency and snoop control features:

- *Snoop Control Unit* (SCU) maintains coherency and consistency in the memory system internal to the cluster, and (optionally) external to the cluster.
- SCU includes a set of snoop filters, automatically sized, one for each cache slice.

Interface features

The DSU-110 has the following interface features:

- Optional AMBA 5 CHI Issue E 256-bit coherent master bus interface, supports up to four *Coherent Hub Interface* (CHI) bus master ports.
- Optional AMBA AXI5 Issue H 256-bit non-coherent master interface, supports up to four *Advanced eXtensible Interface* (AXI) bus master ports.
- Optional 128-bit or 256-bit wide I/O-coherent *Accelerator Coherency Port* (ACP) based on AMBA ACE5-Lite.

³⁹ “Arm DynamIQ Shared Unit-110” at 19.

137. For example, on information and belief, the Dimensity 9000 SoCs provide a single caching agent for the multicore processor, and a plurality of distributed portions each associated with a corresponding one of the plurality of cores⁴⁰:

- L3 cache slice support, for improved bandwidth and cache RAM layout, up to eight slices supported
- L3 cache powerdown based either on cache slices or cache ways
- Cache partitioning support, compliant with *Memory System Resource Partitioning and Monitoring* (MPAM) architecture

- The DSU-110 has an internal transport mechanism that is responsible for all communication between components in the design. The topology of the transport is defined by the number of cores and number of L3 cache slices.

Transport configuration

The topology of the transport mechanism is automatically determined, dependent on the number of cores and L3 cache slices in your cluster. However, you can set transport data path width. For information on the DSU-110 transport, see *RTL configuration process* in *Arm® DynamIQ™ Shared Unit-110 Configuration and Integration Manual*.

138. Further, on information and belief, MediaTek has actively induced and/or contributed to infringement of at least Claims 1 and 11 of the '228 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

139. Users of the '228 Patent Accused Products directly infringe at least Claims 1 and 11 of the '228 Patent when they use the '228 Patent Accused Products in the ordinary, customary, and intended way. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the '228 Patent Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the '228 Patent Accused Products to consumers within the United States and instructing and encouraging such customers to use the '228 Patent Accused Products in the ordinary, customary, and intended

⁴⁰ *Id.* at 19, 21.

way, which MediaTek knew infringes at least Claims 1 and 11 of the '228 Patent, or, alternatively, was willfully blind to the infringement.

140. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the '228 Patent Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the '228 Patent Accused Products in the United States, which MediaTek knew infringes at least Claims 1 and 11 of the '228 Patent, or, alternatively, was willfully blind to the infringement.

141. On information and belief, in violation of 35 U.S.C. § 271(c), MediaTek's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of at least Claims 1 and 11 of the '228 Patent, constituting a material part of the invention. On information and belief, MediaTek knows and has known the same to be especially made or especially adapted for use in an infringement of the '228 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

142. On information and belief, in violation of 35 U.S.C. § 271(f)(1), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claims 1 and 11 of the '228 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United

States in a manner that would infringe the patent if such combination occurred within the United States.

143. On information and belief, in violation of 35 U.S.C. § 271(f)(2), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claims 1 and 11 of the '228 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

144. MediaTek is not licensed or otherwise authorized to practice the claims of the '228 Patent.

145. Thus, by its acts, MediaTek has injured Daedalus and is liable to Daedalus for directly and/or indirectly infringing one or more claims of the '228 Patent, whether literally or under the doctrine of equivalents, including without limitation Claims 1 and 11.

146. On information and belief, MediaTek has known about the '228 Patent at least since August 23, 2022.⁴¹ At a minimum, MediaTek has knowledge of the '228 Patent at least as of the filing of this Complaint. Accordingly, MediaTek's infringement of the '228 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

⁴¹ *Daedalus Prime LLC v. Mazda Motor Corp., et al.*, No. 22-cv-01108 (D. Del. Aug. 23, 2022).

147. As a result of MediaTek's infringement of the '228 Patent, Daedalus has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for MediaTek's infringement, but in no event less than a reasonable royalty with interest and costs.

148. On information and belief, MediaTek will continue to infringe the '228 Patent unless enjoined by this Court. MediaTek's infringement of Daedalus' rights under the '228 Patent will continue to damage Daedalus, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

FIFTH COUNT

(Infringement of U.S Patent No. 11,507,167)

149. Daedalus incorporates by reference the allegations set forth in Paragraphs 1-148 of the Complaint as though fully set forth herein.

150. The claims of the '167 Patent are valid and enforceable.

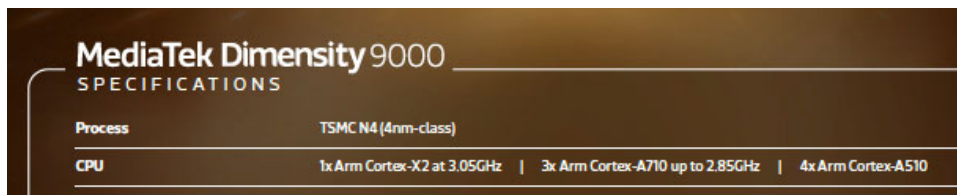
151. On information and belief, in violation of 35 U.S.C. § 271(a), MediaTek has directly infringed and continues to directly infringe one or more claims of the '167 Patent, including at least Claim 1 of the '167 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '167 Patent, including but not limited to its electronic devices containing SoCs or microprocessors based on or derived from ARMv8.2 architecture, as well as subsequent revisions to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9000 SoCs, and all reasonably similar products (the "'167 Patent Accused Products").

152. Each of the '167 Patent Accused Products comprises a multi-core processor. For example, the Dimensity 9000 contains one or more microprocessors based on or derived from the

ARM Cortex-X2 architecture, the ARM Cortex-A710 architecture and the ARM Cortex-A510 architecture.

153. Each of the '167 Patent Accused Products comprises a plurality of cores, wherein each core comprises a processor configured to operate at an independent voltage and frequency level.

154. Specifically, the '167 Patent Accused Products include one or more independent cores or clusters of cores. For example, Dimensity 9000 SoCs comprise one Cortex-X2 core, three Cortex-A710 cores and two clusters of Cortex-A510 cores, configured to operate at an independent voltage and frequency level⁴²:



The infographic displays the specifications for the MediaTek Dimensity 9000 processor. It features a dark background with white text. The title 'MediaTek Dimensity 9000' is prominently displayed at the top, followed by 'SPECIFICATIONS' in a smaller font. Below this, there are two rows of specifications, each with a label on the left and the corresponding details on the right. The first row is labeled 'Process' and specifies 'TSMC N4 (4nm-class)'. The second row is labeled 'CPU' and lists the core configurations: '1x Arm Cortex-X2 at 3.05GHz | 3x Arm Cortex-A710 up to 2.85GHz | 4x Arm Cortex-A510'.

MediaTek Dimensity 9000 SPECIFICATIONS	
Process	TSMC N4 (4nm-class)
CPU	1x Arm Cortex-X2 at 3.05GHz 3x Arm Cortex-A710 up to 2.85GHz 4x Arm Cortex-A510


⁴² <https://mediatek-marketing.files.svcdn.com/production/documents/Dimensity-9000-Infographic.pdf>; <https://i.mediatek.com/dimensity-9000>; Nayak, et al., *supra* note 2; Arm® Cortex®-A510 Core Technical Reference Manual, p. 39-40; Arm® Cortex®-A710 Core Technical Reference Manual, p. 37.

arm

WORLD'S 1ST CORTEX-X2 IN A SMARTPHONE CHIP

The Dimensity 9000 uses new Armv9 architecture CPUs and GPU to deliver unparalleled performance. Its octa-core CPU includes an Arm Cortex-X2 that bursts to epic 3GHz, while new LPDDR5X memory makes data immediately available, eliminating the wait to give immediate responsiveness in any app, whatever you're doing.

- Ultra-Core - 1x Arm Cortex-X2 at 3.05GHz
- Super-Cores - 3x Arm Cortex-A710 up to 2.85GHz
- Efficiency Cores - 4x Arm Cortex-A510
- World's first Arm Mali-G710 MC10 graphics processor
- Big caches - 8MB L3 cache + 6MB system cache
- LPDDR5X 7500Mbps support – 20% more power efficient than LPDDR5



Core	Focus	Microarchitectural Highlights
Cortex-X2	Peak Performance	<ul style="list-style-type: none"> Branch prediction : Decoupled from fetch, improved accuracy and prediction SVE 128b vector length implementation bfloat16, int8 matmul data types support for ML apps Reduced a pipeline stage at dispatch Larger Out of order window size, load-store window/structures. d-TLB Data prefetch enhancements for accuracy and coverage
A710	Sustained Performance	<ul style="list-style-type: none"> Branch prediction : improved accuracy, capacity doubled for BTB, GHB Increased capacity of L1 instruction TLB Improved efficiency with reduction of mid core width Reduced a pipeline stage at dispatch Reduced DSU access, DRAM refills Data prefetch enhancements for accuracy and coverage
A510	Power Efficiency	<ul style="list-style-type: none"> Two cores grouped in a complex, multiple complexes per cluster L2 cache, L2 TLB, vector paths data path shared across a complex Fine-grained scheduling between the cores in the complex 3-wide decode and issue, 3 integer ALU pipeline state of art multi-stage branch prediction Advance data prefetchers for accuracy and coverage

Figure 2.5.2: ARMv9 CPU cluster.

The heterogeneous CPU complex, shown in Fig. 2.5.2, is organized into 3 gears. The 1st gear is a single HP core which utilizes the ARMv9 Cortex-X2 microarchitecture with 64KB L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.

The HP core runs up to 3.4GHz clock speed to meet high-speed compute demands, while the HE cores are optimized to operate efficiently at ultra-low voltage. The BP cores provide a balance of power and performance for average workloads. Depending on the dynamic computing needs, workloads can be seamlessly switched and assigned across different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.




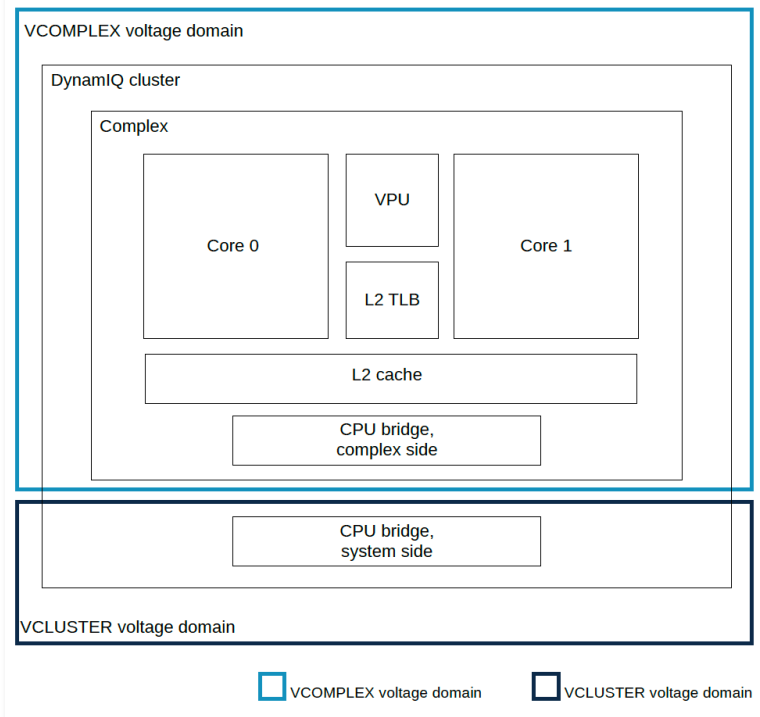
Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

5.1 Voltage and power domains

The *DynamiQ™ Shared Unit-110 (DSU-110) Power Policy Units (PPUs)* control power management for the Cortex®-A510 core. A Cortex®-A510 complex supports separate gated power domains for the complex, for each core inside the complex, and for the *Vector Processing Unit (VPU)*. It also supports a dedicated voltage domain for each complex, and a voltage domain for the DSU-110 DynamiQ™ cluster.

The following figure shows the voltage domains for a Cortex®-A510 configuration with a dual-core complex:

Figure 5-1: Cortex®-A510 voltage domains, dual core



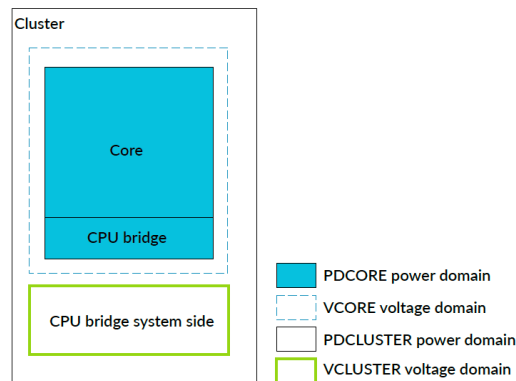
5.1 Voltage and power domains

The *DynamIQ™ Shared Unit-110 (DSU-110) Power Policy Units (PPUs)* control power management for the Cortex®-A710 core. The core supports one power domain, PDCORE, and one system power domain, PDCLUSTER. Similarly, it supports one core voltage domain, VCORE, and one cluster system voltage domain, VCLUSTER. The power domains and voltage domains have the same boundaries.

The PDCORE power domain contains all Cortex®-A710 core logic and part of the core asynchronous bridge that belongs to the VCORE domain. The PDCLUSTER power domain contains the part of the CPU bridge that belongs to the VCLUSTER domain.

The following figure shows the Cortex®-A710 core power domain and voltage domain. It also shows the cluster power domain and voltage domain that cover the system side of the CPU bridge.

Figure 5-1: Cortex®-A710 core voltage domains and power domains



155. Each of the '167 Patent Accused Products comprises at least one core is coupled to a plurality of levels of cache memory.

156. For example, Dimensity 9000 SoCs comprise cores that are coupled to a plurality of levels of cache, including L1, L2 and L3⁴³:

CPU Highlights vs. Dimensity 1200

- Arm Cortex-X2 +40% integer performance over Arm Cortex-A78
- Arm Cortex-A510 +35% integer performance over Arm Cortex-A55
- -50% CPU power @ iso-performance

- Geekbenchv5 single-thread 1278 (+36%) , multi-thread 4400 (+33%)

Cortex-X2 1MB L2	Cortex-A7 10 512KB L2	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10
Cortex-A7 10 512KB L2	Cortex-A7 10 512KB L2	Shared 512KB L2\$ Shared VPU		Shared 512KB L2\$ Shared VPU	
DynamIQ Shared Unit equipped 8MB L3\$					

ARMEDIA
8

The heterogeneous CPU complex, shown in Fig. 2.5.2, is organized into 3 gears. The 1st gear is a single HP core which utilizes the ARMv9 Cortex-X2 microarchitecture with 64KB L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.

⁴³ Rosinger & Pradhan, *supra* note 14, p. 8; Nayak, et al., *supra* note 2; Arm® Cortex®-A510 Core Technical Reference Manual, p. 33; Arm® Cortex®-A710 Core Technical Reference Manual, p. 32; Arm® Cortex®-X2 Core Technical Reference Manual, p. 32.

Figure 3-1: Cortex®-A510 core components




Figure 3-1: Cortex®-A710 core components

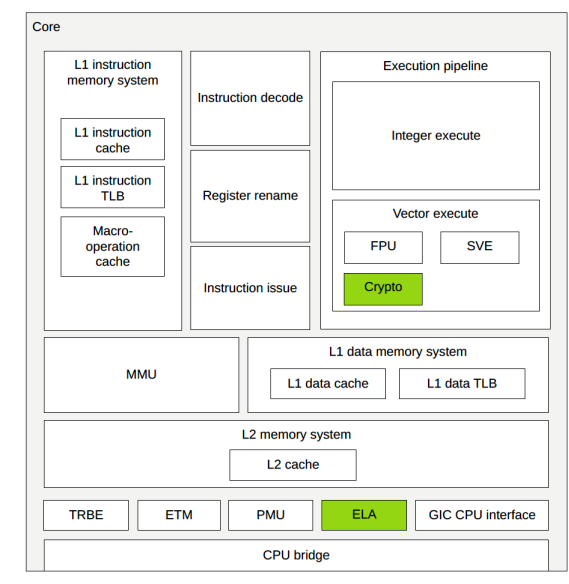
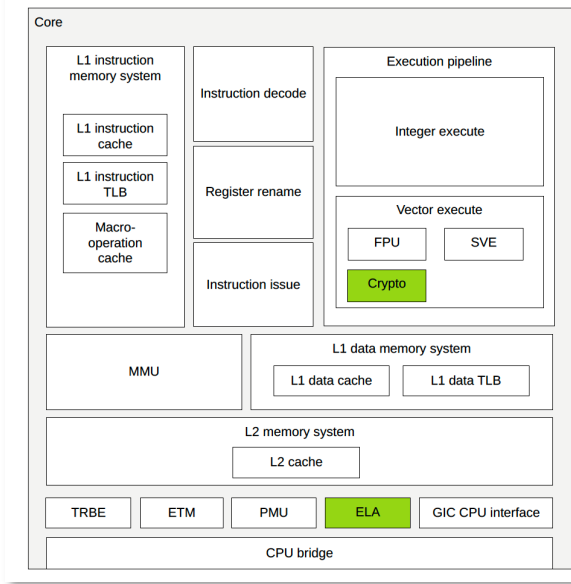


Figure 3-1: Cortex®-X2 core components



157. Each of the '167 Patent Accused Products comprises a power control unit configured to cause an operating voltage to be updated for one or more of the cores in response to receiving a request to alter an operating state of the one or more of the cores.

158. For example, the Dimensity 9000 SoCs include DVFS circuitry capable of controlling the operating voltage for one or more cores, such as an Energy Aware Scheduler and

other logic such as ARM Power Policy Units that are configured by systems such as an ARM System Control Processor⁴⁴:

The HP core runs up to 3.4GHz clock speed to meet high-speed compute demands, while the HE cores are optimized to operate efficiently at ultra-low voltage. The BP cores provide a balance of power and performance for average workloads. Depending on the dynamic computing needs, workloads can be seamlessly switched and assigned across different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.




Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

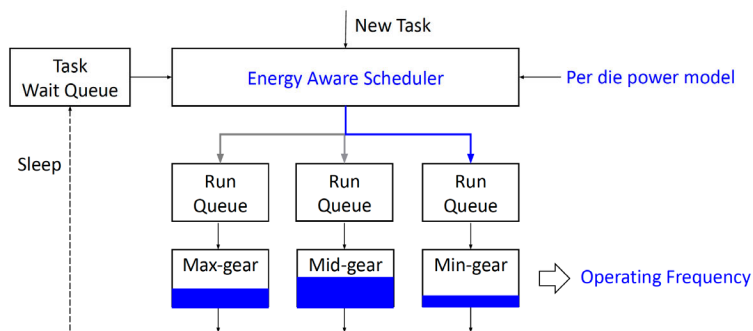
When the operating condition degrades, such as increased IR-drop, the FLL clock frequency will be limited to guarantee safe CPU operation. A voltage increase request, sent to the PMIC, is generated by comparing the FLL output frequency to the PLL input frequency. Conversely, when operating conditions improve and extra voltage margin is no longer needed, the ROSC will oscillate at PLL frequency with a fine code higher than minFC. A voltage decrease request will be sent to the PMIC to reduce the supply voltage until the FLL is frequency locked while using minFC for the ROSC.

In summary, a tri-gear ARMv9 CPU subsystem for a flagship 5G smartphone SoC is introduced with a high-performance core achieving 3.4GHz at robust yield and delivering up to 27% higher peak performance through microarchitectural and implementation advancements. Further, circuit innovation to enable continuous monitoring of on-die power supply is shown. Finally, a new variation aware AVS technology is introduced to further improve CPU power efficiency.

⁴⁴ Nayak, et al., *supra* note 2; Rosinger & Pradhan, *supra* note 14, p. 8; “Arm DynamIQ Shared Unit-110” at 50.

Learnings on Tri-gear CPU Architecture - 2

- Energy Aware Scheduler
- Event Based Operating Frequency Decision



5.1 Power management in the DSU-110

The *DynamiQ™ Shared Unit-110* (DSU-110) provides various mechanisms to control both dynamic and static power dissipation. These mechanisms are associated with a set of power domains, power modes, and operational modes. Some of these mechanisms are brought under software control using *Power Policy Units* (PPUs).

The power management techniques employed by the DSU-110 and cores in the cluster include:

- Internal core clock gating where different internal parts of the core are clock idle
- **Per-core *Dynamic Voltage and Frequency Scaling* (DVFS)**
- Powerdown of components of the cluster which can include:
 - Cores
 - All the L3 cache or parts of the L3 cache. See [5.4.1 L3 cache RAM powerdown](#) on page 57 and [5.4.2 L3 cache slice powerdown](#) on page 61.

159. Each of the '167 Patent Accused Products comprises a power control unit that is further configured to receive a first request to alter an operating state of a first core to a modified operating state at a third voltage level.

160. For example, the Dimensity 9000 SoCs are capable of performing dynamic voltage and frequency scaling (DVFS) and adaptive voltage scaling to adjust operating voltage and frequency.⁴⁵

⁴⁵ Nayak, et al., *supra* note 2.

The HP core runs up to 3.4GHz clock speed to meet high-speed compute demands, while the HE cores are optimized to operate efficiently at ultra-low voltage. The BP cores provide a balance of power and performance for average workloads. Depending on the dynamic computing needs, workloads can be seamlessly switched and assigned across different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.




Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

When the operating condition degrades, such as increased IR-drop, the FLL clock frequency will be limited to guarantee safe CPU operation. A voltage increase request, sent to the PMIC, is generated by comparing the FLL output frequency to the PLL input frequency. Conversely, when operating conditions improve and extra voltage margin is no longer needed, the ROSC will oscillate at PLL frequency with a fine code higher than minFC. A voltage decrease request will be sent to the PMIC to reduce the supply voltage until the FLL is frequency locked while using minFC for the ROSC.

161. Each of the '167 Patent Accused Products comprises a power control unit that is responsive to the first request and can cause a voltage regulator to increase an operating voltage of the first core from a first voltage level to a second voltage level lower than the third voltage level.

162. For example, the Dimensity 9000 SoCs are capable of performing dynamic voltage and frequency scaling (DVFS) and adaptive voltage scaling to adjust operating voltage and

frequency, including at least increasing an operating voltage of the first core from a first voltage level to a second voltage level lower than the third voltage level.⁴⁶

The HP core runs up to 3.4GHz clock speed to meet high-speed compute demands, while the HE cores are optimized to operate efficiently at ultra-low voltage. The BP cores provide a balance of power and performance for average workloads. Depending on the dynamic computing needs, workloads can be seamlessly switched and assigned across different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.




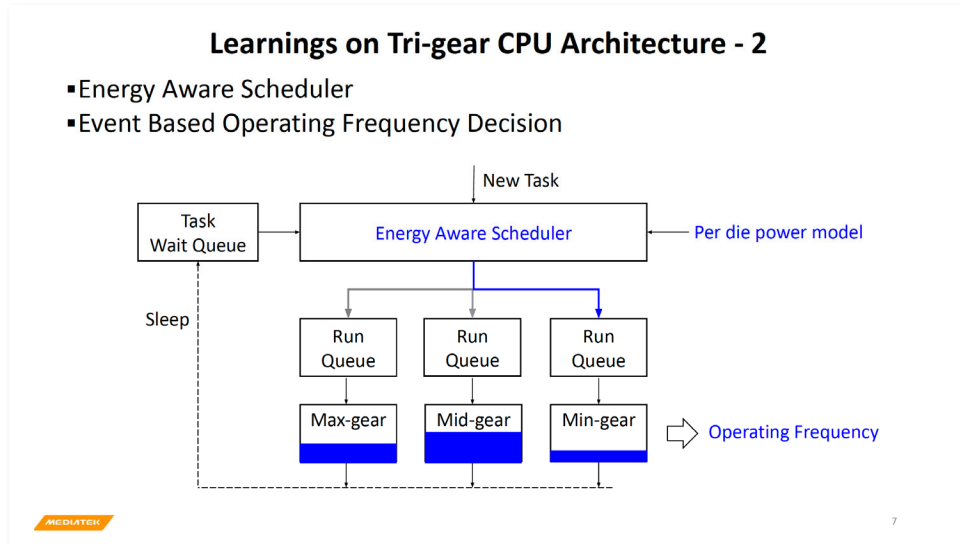
Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

When the operating condition degrades, such as increased IR-drop, the FLL clock frequency will be limited to guarantee safe CPU operation. A voltage increase request, sent to the PMIC, is generated by comparing the FLL output frequency to the PLL input frequency. Conversely, when operating conditions improve and extra voltage margin is no longer needed, the ROSC will oscillate at PLL frequency with a fine code higher than minFC. A voltage decrease request will be sent to the PMIC to reduce the supply voltage until the FLL is frequency locked while using minFC for the ROSC.

163. Each of the '167 Patent Accused Products comprises a power control unit that can enable a second core to exit an inactive state and enter an active state while the operating voltage of the first core is at the second voltage level.

⁴⁶ *Id.*

164. For example, the Dimensity 9000 SoCs are capable of performing dynamic voltage and frequency scaling (DVFS) and adaptive voltage scaling to adjust operating voltage and frequency, and can cause cores to exit an inactive state and enter an active state⁴⁷:



The HP core runs up to 3.4GHz clock speed to meet high-speed compute demands, while the HE cores are optimized to operate efficiently at ultra-low voltage. The BP cores provide a balance of power and performance for average workloads. Depending on the dynamic computing needs, workloads can be seamlessly switched and assigned across different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.

⁴⁷ Rosinger & Pradhan, *supra* note 14, p. 7; Nayak, et al., *supra* note 2; Arm® Cortex®-A510 Core Technical Reference Manual, p. 46; Arm® Cortex®-A710 Core Technical Reference Manual, p. 40-42; Arm® Cortex®-X2 Core Technical Reference Manual, p. 40-42.




Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

5.4 Core power modes

Each core in a Cortex®-A510 complex, as well as the shared logic, has a defined set of power modes and corresponding legal transitions between these power modes. The power mode of each core can be independent of other cores in a complex or DSU-110 DynamIQ™ cluster.

Power modes for a complex are managed at the DynamIQ™ cluster level as *Power Policy Unit (PPU)* modes. See *Power Management* in the *Arm® DynamIQ Shared Unit-110 Technical Reference Manual* for more information.

The following table shows the supported Cortex®-A510 power modes. It describes the meaning of each mode for a Cortex®-A510 core. Although the power mode can affect any logic that is shared between cores in a Cortex®-A510 complex, the table only describes the effect on the core. See 5.5 [Complex power modes](#) on page 50 for more information.

Table 5-1: Cortex®-A510 core power modes

Power mode	Short name	Description
On	ON	The core is powered up and active.
Functional retention	FUNC_RET	The core is fully powered and operational, but the <i>Vector Processing Unit (VPU)</i> is idle.
Full retention	FULL_RET	The core is in retention state. In this mode, only power that is required to retain register and RAM state is available. The core is non-operational. A core must be in <i>Wait for Interrupt (WFI)</i> or <i>Wait for Event (WFE)</i> low-power state before it enters this mode.
Off	OFF	The core is powered down.
Emulated off	OFF_EMU	Emulated off mode permits you to debug the powerup and powerdown cycle without changing the software. In this mode, the core proceeds through all the powerdown steps, except: <ul style="list-style-type: none"> The clock is not gated and power is not removed when the core is powered down. Only the Warm reset is asserted. The debug logic is preserved in the core and remains accessible by the debugger.

5.4.1 On mode

In the On power mode, the Cortex®-A510 core is on and fully operational.

The core can be initialized into the On mode. When a transition to the On mode is completed, all caches are accessible and coherent. Other than the normal architectural steps to enable caches, no additional software configuration is required.

5.4.2 Off mode

In the Off power mode, power is removed completely from the core and no state is retained.

In Off mode, all core logic and RAMs are off. The domain is inoperable and all core state is lost. On transition to Off mode, the L1 and L2 caches are disabled, cleaned, and the core is removed from coherency automatically.

5.4 Core power modes

The Cortex®-A710 core power domain has a defined set of power modes and corresponding legal transitions between these modes. The power mode of each core can be independent of other cores in a cluster.

The *Power Policy Unit* (PPU) of a core manages at the cluster level the transitions between the power modes for that core. See *Power Management* in the *Arm® DynamIQ™ Shared Unit-110 Technical Reference Manual* for more information.

The following table shows the supported Cortex®-A710 core power modes.

Table 5-1: Cortex®-A710 core power modes

Power mode	Short name	Power state
On	ON	The core is powered up and active.
Full retention	FULL_RET	The core is in retention. In this mode, only power that is required to retain register and RAM state is available. The core is not operational. A core must be in <i>Wait for Interrupt</i> (WFI) or <i>Wait for Event</i> (WFE) low-power state before it enters this mode.
Off	OFF	The core is powered down.
Emulated Off	OFF_EMU	Emulated off mode permits you to debug the powerup and powerdown cycle without changing the software. In this mode, the core powerdown is normal, except: <ul style="list-style-type: none"> The clock is not gated and power is not removed when the core is powered down. Only a Warm reset is asserted. The debug logic is preserved in the core and remains accessible by the debugger.
Debug recovery	DBG_RECOV	The RAM and logic are powered up. This mode is for applying a Warm reset to the cluster, while preserving memory and RAS registers for debug purposes. Both cache and RAS state are preserved when transitioning from DBG_RECOV to ON. Caution: This mode must not be used during normal system operation.
Warm reset	WARM_RST	A Warm reset resets all state except for the trace logic and the debug and RAS registers.

5.4.1 On mode

In the On power mode, the Cortex®-A710 core is on and fully operational.

The core can be initialized into the On mode. When a transition to the On mode is completed, all caches are accessible and coherent. Other than the normal architectural steps to enable caches, no additional software configuration is required.

5.4.2 Off mode

In the Off power mode, power is removed completely from the core and no state is retained.

In Off mode, all core logic and RAMs are off. The domain is inoperable and all core state is lost. The L1 and L2 caches are disabled, cleaned and invalidated, and the core is removed from coherency automatically on transition to Off mode.

5.4 Core power modes

The Cortex®-X2 core power domain has a defined set of power modes and corresponding legal transitions between these modes. The power mode of each core can be independent of other cores in a cluster.

The *Power Policy Unit* (PPU) of a core manages at the cluster level the transitions between the power modes for that core. See *Power Management* in the *Arm® DynamIQ™ Shared Unit-110 Technical Reference Manual* for more information.

The following table shows the supported Cortex®-X2 core power modes.

Table 5-1: Cortex®-X2 core power modes

Power mode	Short name	Power state
On	ON	The core is powered up and active.
Full retention	FULL_RET	The core is in retention. In this mode, only power that is required to retain register and RAM state is available. The core is not operational. A core must be in <i>Wait for Interrupt</i> (WFI) or <i>Wait for Event</i> (WFE) low-power state before it enters this mode.
Off	OFF	The core is powered down.
Emulated Off	OFF_EMU	Emulated off mode permits you to debug the powerup and powerdown cycle without changing the software. In this mode, the core powerdown is normal, except: <ul style="list-style-type: none"> The clock is not gated and power is not removed when the core is powered down. Only a Warm reset is asserted. The debug logic is preserved in the core and remains accessible by the debugger.
Debug recovery	DBG_RECOV	The RAM and logic are powered up. This mode is for applying a Warm reset to the cluster, while preserving memory and RAS registers for debug purposes. Both cache and RAS state are preserved when transitioning from DBG_RECOV to ON. Caution: This mode must not be used during normal system operation.
Warm reset	WARM_RST	A Warm reset resets all state except for the trace logic and the debug and RAS registers.

5.4.1 On mode

In the On power mode, the Cortex®-X2 core is on and fully operational.

The core can be initialized into the On mode. When a transition to the On mode is completed, all caches are accessible and coherent. Other than the normal architectural steps to enable caches, no additional software configuration is required.

5.4.2 Off mode

In the Off power mode, power is removed completely from the core and no state is retained.

In Off mode, all core logic and RAMs are off. The domain is inoperable and all core state is lost. The L1 and L2 caches are disabled, cleaned and invalidated, and the core is removed from coherency automatically on transition to Off mode.

165. On information and belief, the 167 Patent Accused Products use ARM’s Intelligent Power Allocation technology in conjunction with the Power Policy Units and a System Control Processor or Resource and Power Manger to enable and disable a second core⁴⁸:

Power management and Power Policy Units

The DynamIQ™ cluster shared logic integrates several *Power management and Power Policy Units* (PPUs) to control power modes and resets. The PPU can be programmed to directly select a

Copyright © 2019–2022 Arm Limited (or its affiliates). All rights reserved.
Non-Confidential

Page 36 of 862

Arm® DynamIQ™ Shared Unit-110 Technical Reference
Manual

Document ID: 101381_0400_11_en
Issue: 11
Technical overview

specific power mode or can be programmed to autonomously switch between power modes within a specified range, based on the requirements of the cluster. The PPU can be programmed from your *System Control Processor* (SCP) using the utility bus to access them.

166. For example, on information and belief, in the Dimensity 9000 SoCs, the Power Policy Units are utilized to enable and disable cores⁴⁹:

⁴⁸ “Arm DynamIQ Shared Unit-110” at 36-37.

⁴⁹ *Id.* at 50.

5.1 Power management in the DSU-110

The *DynamiQ™ Shared Unit-110* (DSU-110) provides various mechanisms to control both dynamic and static power dissipation. These mechanisms are associated with a set of power domains, power modes, and operational modes. Some of these mechanisms are brought under software control using *Power Policy Units* (PPUs).

The power management techniques employed by the DSU-110 and cores in the cluster include:

- Internal core clock gating where different internal parts of the core are clock idle
- Per-core *Dynamic Voltage and Frequency Scaling* (DVFS)
- Powerdown of components of the cluster which can include:
 - Cores
 - All the L3 cache or parts of the L3 cache. See [5.4.1 L3 cache RAM powerdown](#) on page 57 and [5.4.2 L3 cache slice powerdown](#) on page 61.
- Retention which is a low-power mode that retains the register and RAM state. Retention can be applied to the following components of the cluster:
 - Cache RAMs in the cores
 - All of the L3 cache or parts of the L3 cache

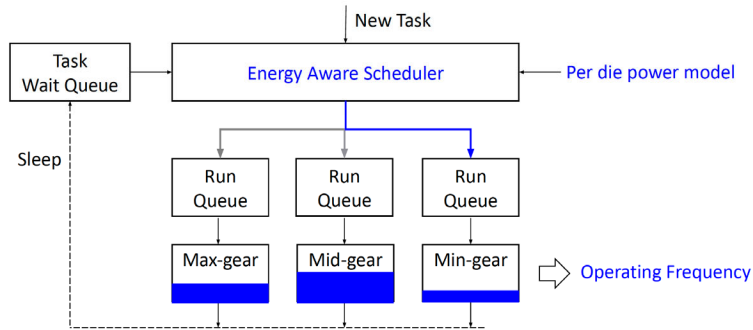
167. Each of the '167 Patent Accused Products comprises a power control unit that can increase the operating voltage of the first core from the second voltage level to the third voltage level after the second core enters the active state.

168. For example, the Dimensity 9000 SoCs are capable of performing dynamic voltage and frequency scaling (DVFS) and adaptive voltage scaling to adjust operating voltage and frequency, and can cause cores to exit an inactive state and enter an active state⁵⁰:

⁵⁰ Rosinger & Pradhan, *supra* note 14, p. 7; Nayak, et al., *supra* note 2; Arm® Cortex®-A510 Core Technical Reference Manual, p. 39; Arm® Cortex®-A710 Core Technical Reference Manual, p. 37; Arm® Cortex®-X2 Core Technical Reference Manual, p. 37.

Learnings on Tri-gear CPU Architecture - 2

- Energy Aware Scheduler
- Event Based Operating Frequency Decision



MEDIATEK

7

The HP core runs up to 3.4GHz clock speed to meet high-speed compute demands, while the HE cores are optimized to operate efficiently at ultra-low voltage. The BP cores provide a balance of power and performance for average workloads. Depending on the dynamic computing needs, workloads can be seamlessly switched and assigned across different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.




Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

5 Power management

The Cortex®-A510 core provides mechanisms to control both dynamic and static power dissipation.

The dynamic power management includes the following features:

- Hierarchical clock gating
- *Per-complex Dynamic Voltage and Frequency Scaling (DVFS)*
- A *Maximum Power Mitigation Mechanism (MPMM)* to control the maximum power

The static power management includes the following features:

- Powerdown
- Per-complex DVFS
- Dynamic retention, a low-power mode that retains the register and RAM state

5 Power management

The Cortex®-A710 core provides mechanisms to control both dynamic and static power dissipation.

The dynamic power management includes the following features:

- Hierarchical clock gating
- *Per-core Dynamic Voltage and Frequency Scaling (DVFS)*

The static power management includes the following features:

- Powerdown
- Dynamic retention, a low-power mode that retains the register and RAM state

5 Power management

The Cortex®-X2 core provides mechanisms to control both dynamic and static power dissipation.

The dynamic power management includes the following features:

- Hierarchical clock gating
- *Per-core Dynamic Voltage and Frequency Scaling (DVFS)*

The static power management includes the following features:

- Powerdown
- Dynamic retention, a low-power mode that retains the register and RAM state

169. Further, on information and belief, MediaTek has actively induced and/or contributed to infringement of at least Claim 1 of the '167 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

170. Users of the '167 Patent Accused Products directly infringe at least Claim 1 of the '167 Patent when they use the '167 Patent Accused Products in the ordinary, customary, and intended way. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the '167 Patent Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the '167 Patent Accused Products to consumers within the United States and instructing and encouraging such customers to use the '167 Patent Accused Products in the ordinary, customary, and intended way, which MediaTek knew infringes at least Claim 1 of the '167 Patent, or, alternatively, was willfully blind to the infringement.

171. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the '167 Patent Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the '167 Patent Accused Products in the United States, which MediaTek knew infringes at least Claim 1 of the '167 Patent, or, alternatively, was willfully blind to the infringement.

172. On information and belief, in violation of 35 U.S.C. § 271(c), MediaTek's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of at least Claim 1 of the '167 Patent, constituting a material part of the invention. On information and belief, MediaTek knows and has known the same to be especially made or especially adapted for use in an

infringement of the '167 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

173. On information and belief, in violation of 35 U.S.C. § 271(f)(1), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 1 of the '167 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

174. On information and belief, in violation of 35 U.S.C. § 271(f)(2), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 1 of the '167 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

175. MediaTek is not licensed or otherwise authorized to practice the claims of the '167 Patent.

176. Thus, by its acts, MediaTek has injured Daedalus and is liable to Daedalus for directly and/or indirectly infringing one or more claims of the '167 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 1.

177. On information and belief, MediaTek has known about the '167 Patent at least since August 23, 2022.⁵¹ At a minimum, MediaTek has knowledge of the '167 Patent at least as of the filing of this Complaint. Accordingly, MediaTek's infringement of the '167 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

178. As a result of MediaTek's infringement of the '167 Patent, Daedalus has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for MediaTek's infringement, but in no event less than a reasonable royalty with interest and costs.

179. On information and belief, MediaTek will continue to infringe the '167 Patent unless enjoined by this Court. MediaTek's infringement of Daedalus' rights under the '167 Patent will continue to damage Daedalus, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

SIXTH COUNT

(Infringement of U.S Patent No. 9,887,838)

180. Daedalus incorporates by reference the allegations set forth in Paragraphs 1-179 of the Complaint as though fully set forth herein.

181. The claims of the '838 Patent are valid and enforceable.

182. On information and belief, in violation of 35 U.S.C. § 271(a), MediaTek has directly infringed and continues to directly infringe one or more claims of the '838 Patent, including at least Claim 9 of the '838 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale,

⁵¹ *Daedalus Prime LLC v. Mazda Motor Corp., et al.*, No. 22-cv-01108 (D. Del. Aug. 23, 2022).

and/or importing into the United States products that embody one or more of the inventions claimed in the '838 Patent, including but not limited to its electronic devices containing SoCs or microprocessors based on or derived from ARMv8.2 architecture, as well as subsequent revisions to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9200 SoCs, and all reasonably similar products (the "'838 Patent Accused Products").

183. Each of the '838 Patent Accused Products comprises a system-on-a-chip. For example, the Dimensity 9200 is a system-on-a-chip (SoC) based on or derived from the ARM Cortex-X3 architecture and the ARM Cortex-A715 architecture.

184. Each of the '838 Patent Accused Products comprises a system-on-a-chip with a security engine that is separate from a processor core of the system-on-a-chip and has a secure memory accessible only by the security engine. Further, the secure memory includes a security key that was encoded in the secure memory during a manufacturing process of the system-on-a-chip.

185. Specifically, the '838 Patent Accused Products comprise a security engine that is separate from a processor core. Specifically, SoCs that include MediaTek's CryptoCore cryptographic module include a sub-chip module which implements the CryptoCore cryptographic module hardware that is separate from a processor core of the system-on-a-chip:⁵²

⁵² MediaTek, Inc., *MediaTek CryptoCore HW v1.0, FW v1.0 FIPS 140-2 Non-Proprietary Security Policy Version 2.2* (June 2018), <https://csrc.nist.gov/CSRC/media/projects/cryptographic-module-validation-program/documents/security-policies/140sp3264.pdf>.

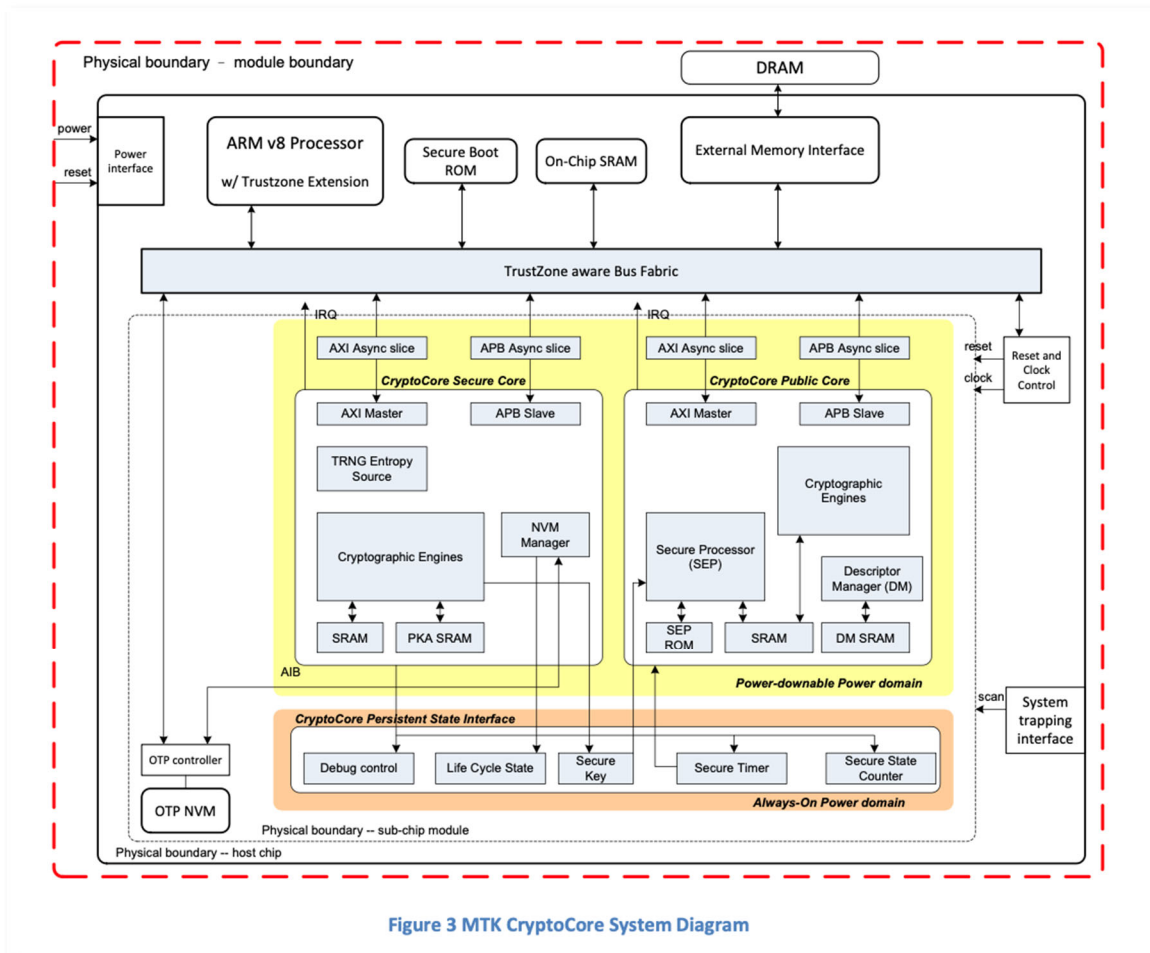


Figure 3 MTK CryptoCore System Diagram

186. Further, the '838 Patent Accused Products comprise secure memory accessible only by the security engine. Specifically, SoCs that include MediaTek's CryptoCore cryptographic module include secure memory⁵³:

⁵³ MediaTek, Inc., *supra* note 44.

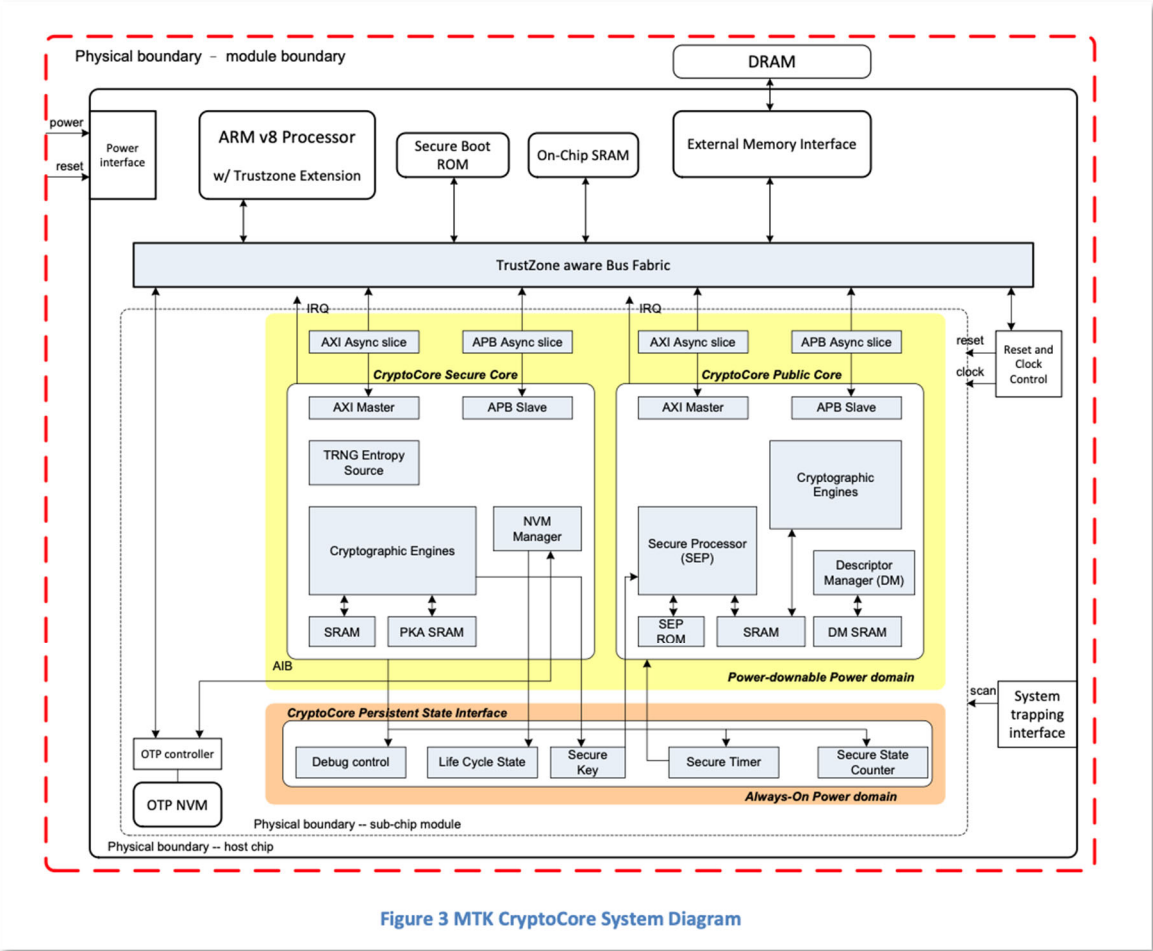


Figure 3 MTK CryptoCore System Diagram

2.1.1 Module Embodiment

MTK CryptoCore is integrated in MediaTek Helio X30 mobile SoC, where the host processor runs two separate operational environments: a Secure OS, also called a Trusted Execution Environment (TEE), and a Public OS as a Rich Execution Environment (REE). The hardware isolation technology enforces data and control separation between the two environments containing the hardware and the firmware components. The firmware in the TEE is generally compact and strictly controlled. The firmware in the REE consists of a fixed kernel and an application layer, and the latter usually modifiable by the user. Therefore hardware and firmware belonging to the TEE are collectively called “Secure World” while those belonging to the REE are called “Public World”. The REE OS for this certification is Linux kernel and Android operating system. The module’s hardware and firmware reflects the system architecture, with the components largely divided into two “cores” — the Secure Core and Public Core. The cores communicate via a Persistent State Interface registers to synchronize their state and pass parameters. The module’s high level diagram, divided between the TEE and REE “worlds”, is shown in Figure 1

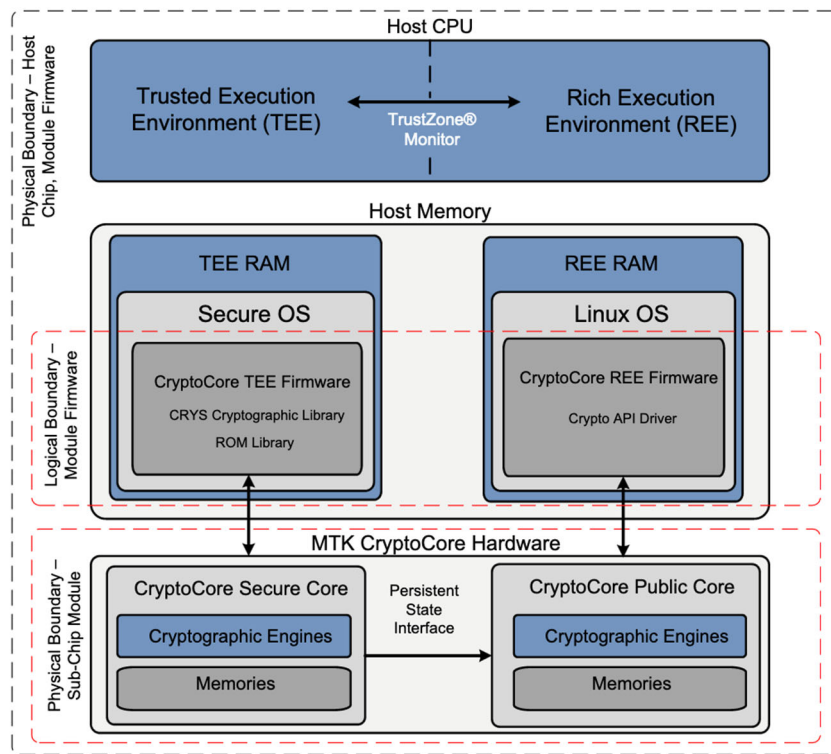


Figure 1 MTK CryptoCore High Level Diagram

187. Further, the '838 Patent Accused Products comprise a security key that was encoded in the secure memory during a manufacturing process of the system-on-a-chip. Specifically, SoCs that include MediaTek’s CryptoCore cryptographic module include secure

memory in the security engine that includes One Time Programmable Non-Volatile Memory (OTP NVM). Secure keys are encoded into the OTP during manufacturing⁵⁴:

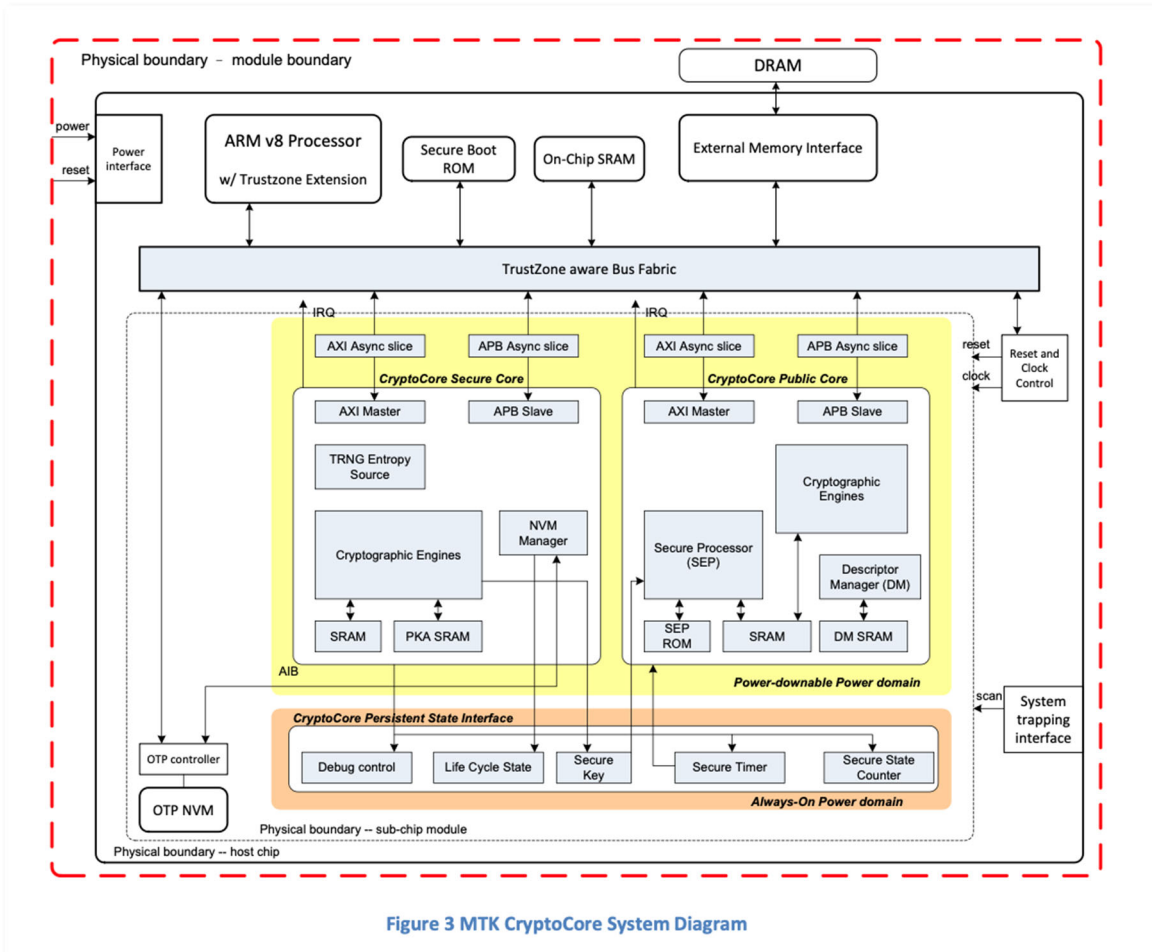


Figure 3 MTK CryptoCore System Diagram

Name	Type	Length (bit)	Purpose	Storage	Entry and Derivation	Zerization
Device Root Key	AES	256	Derivation of device specific keys	OTP and HW register Plaintext	Provisioned into OTP during manufacturing	OTP: command to set key to all 1's Register: via power-on reset
Provisioning Root Master Key	AES	128	Derivation of the Provisioning Master Key	OTP Plaintext	Provisioned into OTP during manufacturing	command to set key to all 1's

⁵⁴ *Id.*

188. Each of the '838 Patent Accused Products comprises a security engine that is capable of generating a random nonce for initiating a request for a secure communication session with a remote server over a network using the nonce. Specifically, to establish a secure session with a remote device, the '838 Accused Products use one or more built-in random number generators that is used to generate one or more cryptographic keys for that session⁵⁵.

CAVP Cert	Algorithm	Standard	Mode / Method	Key Lengths, Curves or Moduli	Use	Public Core	Secure Core
1661	DRBG (Deterministic)	[SP800-90A]	CTR-DRBG	256	Random Number Generation		✓
	Random Bit Generator)						
Vendor-affirmed	CKG (Crypto Key Generation)	[SP800-133]		128	Symmetric Key Generation		✓
	Random Bit Generator)						
Vendor-affirmed	CKG (Crypto Key Generation)	[SP800-133]		128	Symmetric Key Generation		✓

189. For example, the security engine included in the '838 Accused Products is designed to provide key management and high-throughput cryptographic operations for secure communications including, for example, secure playback of DRM-protected media, IPsec VPNs, TLS/SSL link protection and more:⁵⁶

2.1 Module Overview

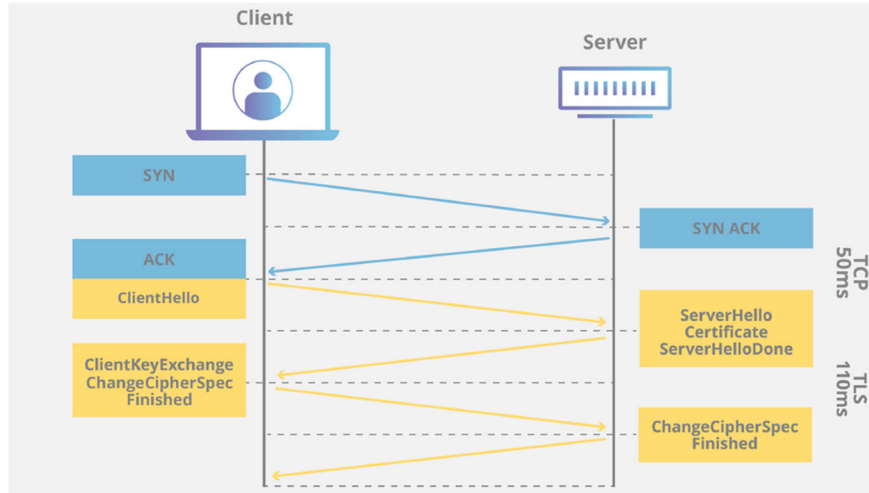
The MTK CryptoCore cryptographic module (hereafter referred to as “the module”) is designed to provide foundational security services for the platform, including secure boot, secure life cycle state, platform identity and key management. It offers high-throughput cryptography operations, suitable for a diverse set of use cases, such as secure playback of DRM-protected media content, IPsec VPNs, TLS/SSL link protection, drive encryption and more.

⁵⁵ *Id.*

⁵⁶ *Id.*

190. Further, the security engine included in the '838 Accused Products is capable of generating a random nonce to initiate a request for a secure communication session with a remote server over a network using the nonce. For example, when initiating a secure session with Transport Layer Security (“TLS”), MediaTek’s CryptoCore cryptographic module generates a random nonce (called a “client random” below) to initiate a request for a secure communication session with a remote server over a network using the nonce during the TLS handshake⁵⁷:

What is a TLS handshake?



TLS is an encryption and authentication protocol designed to secure Internet communications. A TLS handshake is the process that kicks off a communication session that uses TLS. During a TLS handshake, the two communicating sides exchange messages to acknowledge each other, verify each other, establish the cryptographic algorithms they will use, and agree on session keys. TLS handshakes are a foundational part of [how HTTPS works](#).

⁵⁷ *What Happens in a TLS Handshake?* | *SSL Handshake*, <https://www.cloudflare.com/learning/ssl/what-happens-in-a-tls-handshake/>.

When does a TLS handshake occur?

A TLS handshake takes place whenever a user navigates to a website over HTTPS and the browser first begins to query the website's [origin server](#). A TLS handshake also happens whenever any other communications use HTTPS, including [API calls](#) and [DNS over HTTPS](#) queries.

TLS handshakes occur after a [TCP](#) connection has been opened via a TCP handshake.

What are the steps of a TLS handshake?

TLS handshakes are a series of datagrams, or messages, exchanged by a client and a server. A TLS handshake involves multiple steps, as the client and server exchange the information necessary for completing the handshake and making further conversation possible.

The exact steps within a TLS handshake will vary depending upon the kind of key exchange algorithm used and the cipher suites supported by both sides. The RSA key exchange algorithm, while now considered not secure, was used in versions of TLS before 1.3. It goes roughly as follows:

1. **The 'client hello' message:** The client initiates the handshake by sending a "hello" message to the server. The message will include which TLS version the client supports, the cipher suites supported, and a string of random bytes known as the "client random."

What is different about a handshake in TLS 1.3?

TLS 1.3 does not support RSA, nor other cipher suites and parameters that are vulnerable to attack. It also shortens the TLS handshake, making a TLS 1.3 handshake both faster and more secure.

The basic steps of a TLS 1.3 handshake are:

- **Client hello:** The client sends a client hello message with the protocol version, the client random, and a list of cipher suites. Because support for insecure cipher suites has been removed from TLS 1.3, the number of possible cipher suites is vastly reduced. The client hello also includes the parameters that will be used for calculating the premaster secret. Essentially, the client is assuming that it knows the server's preferred key exchange method (which, due to the simplified list of cipher suites, it probably does). This cuts down the overall length of the handshake — one of the important differences between TLS 1.3 handshakes and TLS 1.0, 1.1, and 1.2 handshakes.

191. Each of the '838 Patent Accused Products comprises a security engine that is capable of performing a cryptographic key exchange with the remote server.

192. For example, the security engine included in the '838 Accused Products is designed to provide key management and high-throughput cryptographic operations for secure communications including, for example, secure playback of DRM-protected media, IPSec VPNs, TLS/SSL link protection and more⁵⁸:

2.1 Module Overview

The MTK CryptoCore cryptographic module (hereafter referred to as “the module”) is designed to provide foundational security services for the platform, including secure boot, secure life cycle state, platform identity and key management. It offers high-throughput cryptography operations, suitable for a diverse set of use cases, such as secure playback of DRM-protected media content, IPSec VPNs, TLS/SSL link protection, drive encryption and more.

193. Further, the security engine included in the '838 Accused Products is capable of performing a cryptographic key exchange with the remote server, for example, when performing a TLS handshake⁵⁹:

⁵⁸ MediaTek, Inc., *supra* note 44.

⁵⁹ *What Happens in a TLS Handshake?*, *supra* note 49; E. Rescorla, *The Transport Layer Security (TLS) Protocol Version 1.3*, RFC 8446, (August 2018), at p. 10, <https://datatracker.ietf.org/doc/html/rfc8446>.

What are the steps of a TLS handshake?

TLS handshakes are a series of datagrams, or messages, exchanged by a client and a server. A TLS handshake involves multiple steps, as the client and server exchange the information necessary for completing the handshake and making further conversation possible.

The exact steps within a TLS handshake will vary depending upon the kind of key exchange algorithm used and the cipher suites supported by both sides. The RSA key exchange algorithm, while now considered not secure, was used in versions of TLS before 1.3. It goes roughly as follows:

1. **The 'client hello' message:** The client initiates the handshake by sending a "hello" message to the server. The message will include which TLS version the client supports, the cipher suites supported, and a string of random bytes known as the "client random."

What is different about a handshake in TLS 1.3?

TLS 1.3 does not support RSA, nor other cipher suites and parameters that are vulnerable to attack. It also shortens the TLS handshake, making a TLS 1.3 handshake both faster and more secure.

The basic steps of a TLS 1.3 handshake are:

- **Client hello:** The client sends a client hello message with the protocol version, the client random, and a list of cipher suites. Because support for insecure cipher suites has been removed from TLS 1.3, the number of possible cipher suites is vastly reduced. The client hello also includes the parameters that will be used for calculating the premaster secret. Essentially, the client is assuming that it knows the server's preferred key exchange method (which, due to the simplified list of cipher suites, it probably does). This cuts down the overall length of the handshake — one of the important differences between TLS 1.3 handshakes and TLS 1.0, 1.1, and 1.2 handshakes.

TLS supports three basic key exchange modes:

- (EC)DHE (Diffie-Hellman over either finite fields or elliptic curves)
- PSK-only
- PSK with (EC)DHE

194. Further, the security engine included in the '838 Accused Products is capable of providing key establishment services based on (EC)DHE⁶⁰:

7.4 Key Establishment

Key establishment and key derivation services offered by the module are detailed in Section 4.2. The module offers key establishment services based on the approved Elliptic Curve Diffie-Hellman and Finite Field Diffie-Hellman [SP800-56A], and the non-approved Elliptic Curve Integrated Encryption Scheme [IEEE1363]. The key establishment services are only allowed when using keys allowed by [SP800-131A].

195. Further, in the TLS 1.3 key exchange, the client sends a ClientHello message which contains the random nonce and a list of symmetric cipher/HKDF hash pairs⁶¹:

In the Key Exchange phase, the client sends the ClientHello ([Section 4.1.2](#)) message, which contains a random nonce (ClientHello.random); its offered protocol versions; a list of symmetric cipher/HKDF hash pairs; either a set of Diffie-Hellman key shares (in the "key_share" ([Section 4.2.8](#)) extension), a set of pre-shared key labels (in the "pre_shared_key" ([Section 4.2.11](#)) extension), or both; and potentially additional extensions. Additional fields and/or messages may also be present for middlebox compatibility.

The server processes the ClientHello and determines the appropriate cryptographic parameters for the connection. It then responds with its own ServerHello ([Section 4.1.3](#)), which indicates the negotiated connection parameters. The combination of the ClientHello and the ServerHello determines the shared keys. If (EC)DHE key establishment is in use, then the ServerHello contains a "key_share" extension with the server's ephemeral Diffie-Hellman share; the server's share MUST be in the same group as one of the client's shares. If PSK key establishment is in use, then the ServerHello contains a "pre_shared_key" extension indicating which of the client's offered PSKs was selected. Note that implementations can use (EC)DHE and PSK together, in which case both extensions will be supplied.

196. Each of the '838 Patent Accused Products comprises a security engine that is capable of generating a symmetric session key, based on the cryptographic key exchange, to encrypt messages sent to the remote server and decrypt messages received from the remote server during the secure communication session.

⁶⁰ MediaTek, Inc., *supra* note 44.

⁶¹ Rescorla, *supra* note 51.

197. For example, the security engine included in the '838 Accused Products is designed to provide key management and high-throughput cryptographic operations for secure communications including, for example, secure playback of DRM-protected media, IPSec VPNs, TLS/SSL link protection and more⁶²:

2.1 Module Overview

The MTK CryptoCore cryptographic module (hereafter referred to as “the module”) is designed to provide foundational security services for the platform, including secure boot, secure life cycle state, platform identity and key management. It offers high-throughput cryptography operations, suitable for a diverse set of use cases, such as secure playback of DRM-protected media content, IPSec VPNs, TLS/SSL link protection, drive encryption and more.

198. Further, the security engine included in the '838 Accused Products is capable of generating a symmetric session key, based on the cryptographic key exchange, to encrypt messages sent to the remote server and decrypt messages received from the remote server during the secure communication session, for example, when performing a TLS handshake⁶³:

What happens during a TLS handshake?

During the course of a TLS handshake, the client and server together will do the following:

- Specify which version of TLS (TLS 1.0, 1.2, 1.3, etc.) they will use
- Decide on which cipher suites (see below) they will use
- Authenticate the identity of the server via the server’s public key and the SSL certificate authority’s digital signature
- Generate session keys in order to use symmetric encryption after the handshake is complete

⁶² MediaTek, Inc., *supra* note 44.

⁶³ *What Happens in a TLS Handshake?*, *supra* note 49.

199. Further, in the TLS 1.3 key exchange, a symmetric key is generated based on the (EC)DHE and symmetric cipher/HKDF hash pair negotiated during the key exchange.⁶⁴

200. Further, TLS 1.3 requires support of at least the symmetric cipher AES 128 GCM and HMAC-based Extract-and-Expand Key Derivation Function (HKDF) based on SHA 256. *Id.* at 102. Both of these (and others supported by TLS 1.3) are supported by MediaTek’s CryptoCore cryptographic module⁶⁵:

Service Name	Purpose	Security Functions	Keys and CSPs	User Role	CO Role	Approved	Public Core	Secure Core
AES Public Core Approved	Encryption Decryption	AES-128, 192, 256, modes: ECB, CBC, CTR, OFB, CMAC, XTS, CTS (CBC-CS1), GCM, GMAC	Input: User keys (I)	✓		✓	✓	

CAVP Cert	Algorithm	Standard	Mode / Method	Key Lengths, Curves or Moduli	Use	Public Core	Secure Core
3190	HMAC	[FIPS198-1,RFC2104]	HMAC SHA-1, SHA-224, SHA-256		Message authentication	✓	
3191	HMAC	[FIPS198-1,RFC2104]	HMAC SHA-1, SHA-224, SHA-256, SHA-384, SHA-512		Message authentication		✓

201. Each of the ’838 Patent Accused Products comprises a security engine that is capable of encrypting the symmetric session key based on the security key. For example, the Secure Key Mechanism in MediaTek’s CryptoCore cryptographic module maintains control over keys such as the claimed symmetric session key by encrypting them using AES-128 CCM⁶⁶:

⁶⁴ Rescorla, *supra* note 51, pp. 26, 93.

⁶⁵ MediaTek, Inc., *supra* note 44, p. 19; *Id.* at pp. 4-5.

⁶⁶ *Id.* at p. 15.

2.3.4.7 Secure Key Mechanism

The Secure Key mechanism enables cryptographic operations with the Public Core’s high throughput while the Secure Core takes full control of the keys. The mechanism uses AES-128 CCM authenticated encryption for key wrapping with the Session Key, and a dedicated hardware channel for data passing between the Secure Core and the Public Core. The Session Key is passed from Secure Core to Public Core. The Secure Core uses the Session Key to encrypt a data structure containing a user key and parameters controlling its usage. On the Public Core side, the SEP subsystem unwraps the data using the Session Key, checks the parameters, and loads the key directly into the Public Core encryption engines. The keys are left inaccessible to the Public Core Firmware. The encrypted Secure Key packages are returned to user, and are passed from Secure Core RAM to Public Core RAM by an out-of-band mechanism, such as a TEE application. From there, Secure Key packages can be passed as the key parameter to Public Core functions supporting Secure Key.

Secure Key packages are rendered inaccessible when Session Key is regenerated, either on power-on or on demand.

202. Further, the internal Session Key used by the Secure Key Mechanism is derived from the Device Root Key (security key) that was encoded in the secure memory during a manufacturing process of the system-on-a-chip⁶⁷:

Service Name	Purpose	Security Functions	Keys and CSPs	User Role	CO Role	Approved	Public Core	Secure Core
NIST SP800-108 Key Derivation Function Approved	Key-Based Key derivation function (KBKDF) [SP800-108]	KDF in Counter Mode. AES-128 or 256 in CMAC mode	Input: User keys (I) Device Root Key (R) Device Root Key (R) OEM Key (I) Output: User Keys (O) Session Key (W)	✓		✓		✓
Session Key Generation	Derive Session Key from Device Root Key and DRBG output	Uses the NIST SP800-108 KBKDF; Uses DRBG	Input: Device Root Key (R) Output: Session Key (W)	✓		✓		✓

Name	Type	Length (bit)	Purpose	Storage	Entry and Derivation	Zeroization
Session Key	AES	128	Encryption of Secure Key packages Encryption of RAM backups	HW register Plaintext	Derived using KBKDF from a 96-bit random value and CSP in OTP	via power-on reset

⁶⁷ *Id.* at pp. 19, 24; *Id.* at pp. 27-28.

203. Each of the '838 Patent Accused Products comprises a security engine that is capable of storing the encrypted session key in the secure memory⁶⁸:

2.3.4.7 Secure Key Mechanism

The Secure Key mechanism enables cryptographic operations with the Public Core's high throughput while the Secure Core takes full control of the keys. The mechanism uses AES-128 CCM authenticated encryption for key wrapping with the Session Key, and a dedicated hardware channel for data passing between the Secure Core and the Public Core. The Session Key is passed from Secure Core to Public Core. The Secure Core uses the Session Key to encrypt a data structure containing a user key and parameters controlling its usage. On the Public Core side, the SEP subsystem unwraps the data using the Session Key, checks the parameters, and loads the key directly into the Public Core encryption engines. The keys are left inaccessible to the Public Core Firmware. The encrypted Secure Key packages are returned to user, and are passed from Secure Core RAM to Public Core RAM by an out-of-band mechanism, such as a TEE application. From there, Secure Key packages can be passed as the key parameter to Public Core functions supporting Secure Key.

Secure Key packages are rendered inaccessible when Session Key is regenerated, either on power-on or on demand.

204. Each of the '838 Patent Accused Products is a system-on-a-chip that is capable of establishing the secure communication session with the remote server over the network using the session key⁶⁹:

2.1 Module Overview

The MTK CryptoCore cryptographic module (hereafter referred to as "the module") is designed to provide foundational security services for the platform, including secure boot, secure life cycle state, platform identity and key management. It offers high-throughput cryptography operations, suitable for a diverse set of use cases, such as secure playback of DRM-protected media content, IPSec VPNs, TLS/SSL link protection, drive encryption and more.

205. Users of the '838 Patent Accused Products directly infringe at least Claim 9 of the '838 Patent when they use the '838 Patent Accused Products in the ordinary, customary, and intended way. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly

⁶⁸ *Id.* at p. 15.

⁶⁹ *Id.* at p. 1.

inducing consumers to use the '838 Patent Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the '838 Patent Accused Products to consumers within the United States and instructing and encouraging such customers to use the '838 Patent Accused Products in the ordinary, customary, and intended way, which MediaTek knew infringes at least Claim 9 of the '838 Patent, or, alternatively, was willfully blind to the infringement.

206. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the '838 Patent Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the '838 Patent Accused Products in the United States, which MediaTek knew infringes at least Claim 9 of the '838 Patent, or, alternatively, was willfully blind to the infringement.

207. On information and belief, in violation of 35 U.S.C. § 271(c), MediaTek's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of at least Claim 9 of the '838 Patent, constituting a material part of the invention. On information and belief, MediaTek knows and has known the same to be especially made or especially adapted for use in an infringement of the '838 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

208. On information and belief, in violation of 35 U.S.C. § 271(f)(1), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the

United States all or a substantial portion of the components of the patented invention of at least Claim 9 of the '838 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

209. On information and belief, in violation of 35 U.S.C. § 271(f)(2), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 9 of the '838 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

210. MediaTek is not licensed or otherwise authorized to practice the claims of the '838 Patent.

211. Thus, by its acts, MediaTek has injured Daedalus and is liable to Daedalus for directly and/or indirectly infringing one or more claims of the '838 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 9.

212. On information and belief, MediaTek has known about the '838 Patent at least since August 23, 2022.⁷⁰ At a minimum, MediaTek has knowledge of the '838 Patent at least as of the filing of this Complaint. Accordingly, MediaTek's infringement of the '838 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case

⁷⁰ *Daedalus Prime LLC v. Mazda Motor Corp., et al.*, No. 22-cv-01108 (D. Del. Aug. 23, 2022).

warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

213. As a result of MediaTek's infringement of the '838 Patent, Daedalus has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for MediaTek's infringement, but in no event less than a reasonable royalty with interest and costs.

214. On information and belief, MediaTek will continue to infringe the '838 Patent unless enjoined by this Court. MediaTek's infringement of Daedalus' rights under the '838 Patent will continue to damage Daedalus, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

SEVENTH COUNT

(Infringement of U.S Patent No. 10,705,960)

215. Daedalus incorporates by reference the allegations set forth in Paragraphs 1-214 of the Complaint as though fully set forth herein.

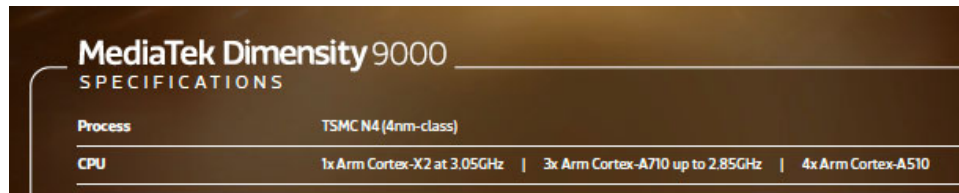
216. The claims of the '960 Patent are valid and enforceable.

217. On information and belief, in violation of 35 U.S.C. § 271(a), MediaTek has directly infringed and continues to directly infringe one or more claims of the '960 Patent, including at least Claim 15 of the '960 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '960 Patent, including but not limited to its electronic devices containing SoCs or microprocessors based on or derived from ARMv8.2 architecture, as well as subsequent revisions to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9000 SoCs, and all reasonably similar products (the "'960 Patent Accused Products").

218. Each of the '960 Patent Accused Products comprises a system. For example, the Dimensity 9000 SoC is a system-on-a-chip.

219. Each of the '960 Patent Accused Products comprises a plurality of cores, the plurality of cores comprising symmetric multi-threaded cores.

220. For example, Dimensity 9000 SoCs include multiple clusters with several clusters having several identical core architectures with identical performance specifications, including 4 x Cortex-A510 cores, and 3 x Cortex-A710 cores. These are symmetric since they contain multiple processors cores having identical specifications and configurations⁷¹:



The infographic displays the specifications for the MediaTek Dimensity 9000. It features a dark background with white text. The title 'MediaTek Dimensity 9000' is prominently displayed at the top, followed by 'SPECIFICATIONS' in a smaller font. Below this, there are two rows of specifications, each with a label on the left and the details on the right. The first row is labeled 'Process' and specifies 'TSMC N4 (4nm-class)'. The second row is labeled 'CPU' and lists three core configurations: '1x Arm Cortex-X2 at 3.05GHz', '3x Arm Cortex-A710 up to 2.85GHz', and '4x Arm Cortex-A510'.

MediaTek Dimensity 9000 SPECIFICATIONS	
Process	TSMC N4 (4nm-class)
CPU	1x Arm Cortex-X2 at 3.05GHz 3x Arm Cortex-A710 up to 2.85GHz 4x Arm Cortex-A510

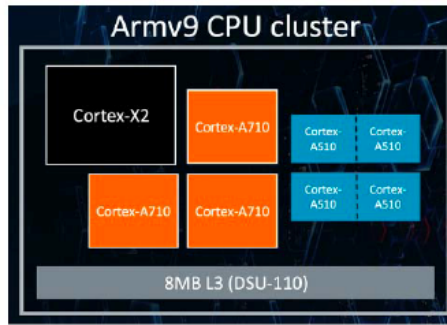
⁷¹ <https://mediatek-marketing.files.svcdn.com/production/documents/Dimensity-9000-Infographic.pdf>; <https://i.mediatek.com/dimensity-9000>; Nayak, et al., *supra* note 2; ARM® Cortex -A Series, Version 4.0 Programmer's Guide (2011-2013), <https://developer.arm.com/documentation/den0013/d/Multi-core-processors/Symmetric-multi-processing>; ARM® Cortex -A Series, Version 1.0 Programmer's Guide for ARMv8-A (2015), <https://developer.arm.com/documentation/den0024/a/Multi-core-processors/Multi-processing-systems/Symmetric-multi-processing>; <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8>, page 18.

arm

WORLD'S 1ST CORTEX-X2 IN A SMARTPHONE CHIP

The Dimensity 9000 uses new Armv9 architecture CPUs and GPU to deliver unparalleled performance. Its octa-core CPU includes an Arm Cortex-X2 that bursts to epic 3GHz, while new LPDDR5X memory makes data immediately available, eliminating the wait to give immediate responsiveness in any app, whatever you're doing.

- Ultra-Core - 1x Arm Cortex-X2 at 3.05GHz
- Super-Cores - 3x Arm Cortex-A710 up to 2.85GHz
- Efficiency Cores - 4x Arm Cortex-A510
- World's first Arm Mali-G710 MC10 graphics processor
- Big caches - 8MB L3 cache + 6MB system cache
- LPDDR5X 7500Mbps support – 20% more power efficient than LPDDR5



Core	Focus	Microarchitectural Highlights
Cortex-X2	Peak Performance	Branch prediction : Decoupled from fetch, improved accuracy and prediction SVE 128b vector length implementation bfloat16, int8 matmul data types support for ML apps Reduced a pipeline stage at dispatch Larger Out of order window size, load-store window/structures. d-TLB Data prefetch enhancements for accuracy and coverage
A710	Sustained Performance	Branch prediction : improved accuracy, capacity doubled for BTB, GHB Increased capacity of L1 instruction TLB Improved efficiency with reduction of mid core width Reduced a pipeline stage at dispatch Reduced DSU access, DRAM refills Data prefetch enhancements for accuracy and coverage
A510	Power Efficiency	Two cores grouped in a complex, multiple complexes per cluster L2 cache, L2 TLB, vector paths data path shared across a complex Fine-grained scheduling between the cores in the complex 3-wide decode and issue, 3 integer ALU pipeline state of art multi-stage branch prediction Advance data prefetchers for accuracy and coverage

Figure 2.5.2: ARMv9 CPU cluster.

L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.

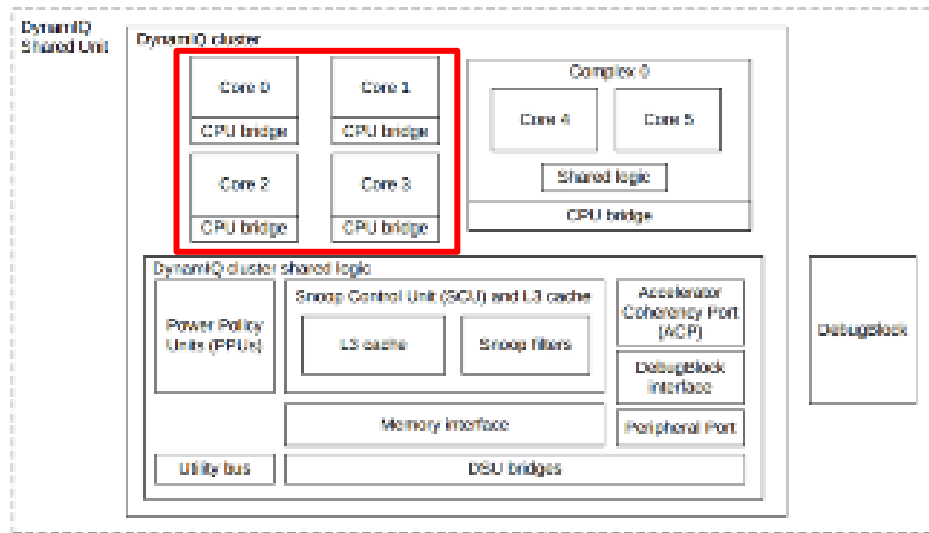
All processor cores in the CPU subsystem incorporate the ARMv9 instruction set with key architectural advances. Memory Tagging Extension (MTE) enables greater security by locking data in the memory using a tag which can only be accessed by the correct key held by the pointer accessing the memory location, as shown in Fig. 2.5.1. Further, a Scalable Vector Extension 2 (SVE2) allows a scalable vector length in multiples of 128b, up to 2048b, enabling increased DSP and ML vector-processing capabilities, as shown in Fig. 2.5.1.

Symmetric multi-processing

Symmetric multi-processing (SMP) is a software architecture that dynamically determines the roles of individual processors. Each core in the cluster has the same view of memory and of shared hardware. Any application, process or task can run on any core and the operating system scheduler can dynamically migrate tasks between cores to achieve optimal system load.

We expect that readers will be familiar with the fundamental operating principles of an OS, but OS terminology will be briefly reviewed here. An application that executes under an operating system is known as a *process*. It performs many operations through calls to the system library that provides certain functions from library code, but also acts a wrapper for system calls to kernel operations. Individual processes have associated resources, including stack, heap and constant data areas, and properties such as scheduling priority settings. The kernel view of a process is called a *task*.

Figure 2-1: DSU-110 DynamIQ™ cluster



221. Further, ARM cores are multi-threaded by design as indicated in the ARM DynamIQ™ specification, and are capable of simultaneously executing two or more processing threads⁷²:

The DSU-110 DynamIQ™ cluster debug system supports both single and multi-threaded cores.

The Arm architecture allows for cores to be single, or multi-threaded. A *Processing Element (PE)* performs a thread of execution. A single-threaded core has one PE and a multi-threaded core has two or more PEs. Because the debugging system allows individual threads to be debugged, the term PE is used throughout this chapter. Where a reference to a core is made, the core can be a single, or multi-threaded core.

⁷² <https://documentation-service.arm.com/static/61ba2d8176bb7f0e683c35cc> Pages 168, 379; ARM® Cortex® - A710 Core, Technical Reference Manual (2019-2021), <https://developer.arm.com/documentation/101800/latest>.

Processing element
 The Arm architecture allows for cores to support multiple PEs.
 The *DynamIQ™ Shared Unit-110 (DSU-110)* supports cores with multiple PEs. Where a reference to a core is made, the core could be a core with only one PE (single-threaded core) or multiple PEs (multi-threaded core).

Bits	Name	Description	Reset
[24]	MT	Multithreaded PE. Performance of PEs at the lowest affinity level is largely independent. If ERRDLVAFF-AHO is not valid, this bit is not valid and reads as UNKNOWN.	0b0

222. Further, other cores in the '960 Patent Accused products contain specific support for the use of multiple threads.

223. Further, the ARMv8 programmer's guide describes features from the architecture that are designed to assist in the execution of multiple threads⁷³:

The additions to the architecture mean that a single physical core can execute code from both the Normal world and the Secure world in a time-sliced fashion, although this depends on the availability of interrupt-generating peripherals that can be configured to be accessible only by the Secure World. For example, a Secure timer interrupt could be used to guarantee some execution time for the Secure world, in a manner resembling preemptive multitasking. Such peripherals may or may not be available, depending on the level of security and use cases that the platform designer intends to support.

hardware. Any application, process, or task can run on any core and the operating system scheduler can dynamically migrate tasks between cores to achieve optimal system load. A multi-threaded application can run on several cores at once. The operating system can hide much of the complexity from applications.

224. Further, MediaTek has announced that its SoCs are optimized and ready to go for the Android operating system software.⁷⁴

225. Further, Android operating system software compatible with the '960 Patent Accused Products contains support for more than one thread of execution within a process, wherein

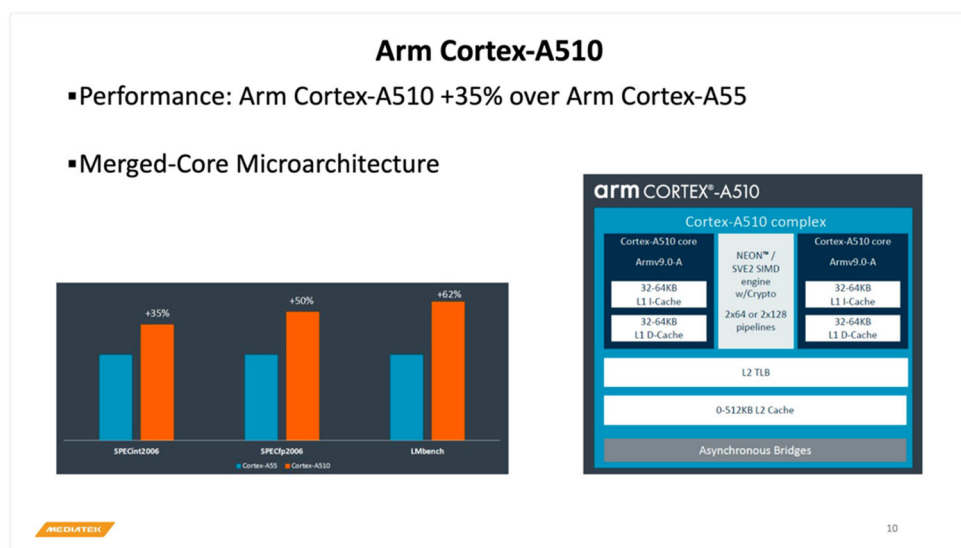
⁷³ ARM®, *supra* note 63.

⁷⁴ MediaTek, *MediaTek SoCs Are Optimized and Ready For Android Oreo (Go Edition)* (December 7, 2017), <https://corp.mediatek.com/news-events/press-releases/mediatek-socs-are-optimized-and-ready-for-android-oreo-go-edition>.

those threads run concurrently, and contains support for the use of multiple threads within programs such as worker threads.⁷⁵

226. Each of the '960 Patent Accused Products comprises a cache subsystem, the cache subsystem comprising a plurality of first level caches and at least one higher level distributed cache comprising a plurality of distributed cache portions that are physically distributed across a die and shared by the plurality of cores, each first level cache integral to one of the plurality of cores and each distributed cache portion accessible to each of the plurality of cores.

227. For example, the MediaTek Dimensity 9000 includes a cache subsystem comprising L1, L2 and L3 caches, including a plurality of first level (L1) caches that are integral to one of the plurality of cores⁷⁶:



⁷⁵ <https://developer.android.com/courses/extras/multithreading>;
<https://developer.android.com/guide/components/processes-and-threads>.

⁷⁶ Rosinger & Pradhan, *supra* note 14, pp. 8, 10; Nayak, et al., *supra* note 2; Arm® Cortex®-A510 Core Technical Reference Manual, p. 33; Arm® Cortex®-A710 Core Technical Reference Manual, p. 32.

CPU Highlights vs. Dimensity 1200

- Arm Cortex-X2 +40% integer performance over Arm Cortex-A78
- Arm Cortex-A510 +35% integer performance over Arm Cortex-A55
- -50% CPU power @ iso-performance

- Geekbenchv5 single-thread 1278 (+36%) , multi-thread 4400 (+33%)

Cortex-X2 1MB L2	Cortex-A7 10 512KB L2	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10
Cortex-A7 10 512KB L2	Cortex-A7 10 512KB L2	Shared 512KB L2\$ Shared VPU		Shared 512KB L2\$ Shared VPU	
DynamIQ Shared Unit equipped 8MB L3\$					



8

L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.

Figure 3-1: Cortex®-A510 core components

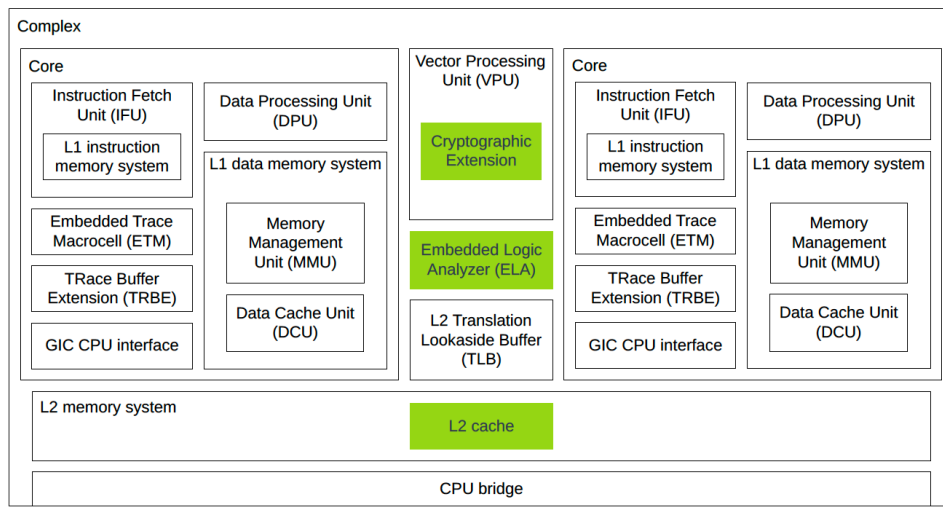
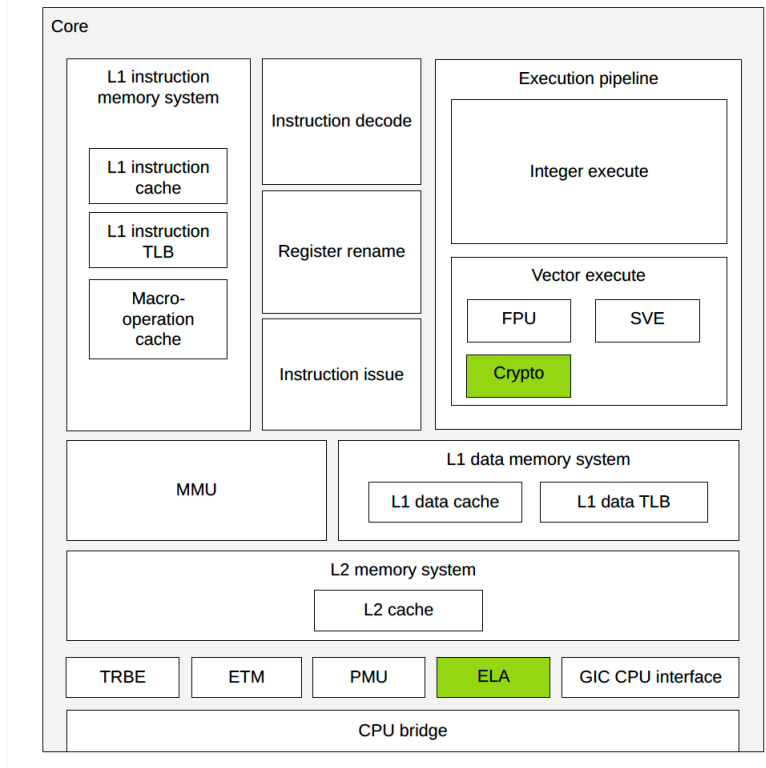


Figure 3-1: Cortex®-A710 core components



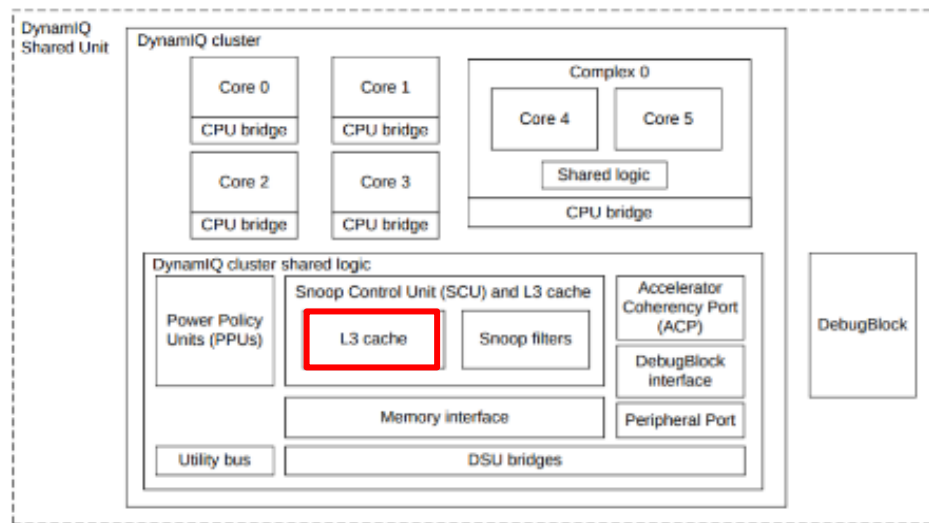
228. Further, the MediaTek Dimensity 9000 also includes a shared L3 cache (e.g., the higher level cache) as part of the DynamIQ Shared Unit - 110. The L3 cache is split into L3 cache slices (or a plurality of distributed cache portions), which are distributed across the SoC die and accessible to each of the plurality of cores⁷⁷:

2 The DynamIQ™ Shared Unit-110

The DynamIQ™ Shared Unit-110 (DSU-110) provides a shared L3 memory system, snoop control and filtering, and other control logic to support a cluster of A-class architecture cores. The cluster is called the DSU-110 DynamIQ™ cluster. Additionally, all the external interfaces to System on Chip (SoC) are provided through the DSU-110.

⁷⁷ Arm DynamIQ Shared Unit-110 Technical Reference Manual at 18, 101; Arm® Cortex®-A710 Core Technical Reference Manual, p. 32.

Figure 2-1: DSU-110 DynamiQ™ cluster




7.6 Cache slices and portions

The L3 cache of the *DynamiQ™ Shared Unit-110* (DSU-110) can be divided into up to eight identical slices, each containing between 256KB and 2MB of the cache. A cache slice consists of the data, tag, victim, and snoop filter RAMs and associated logic. A portion is a further subdivision of RAM in a cache slice.

For each cache slice, both the data RAM and tag RAM is subdivided into two portions.

Figure 3-1: Cortex®-A710 core components

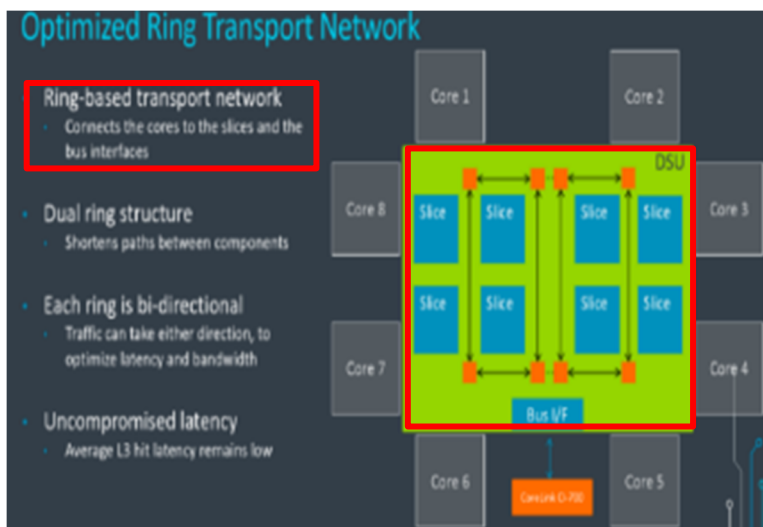


229. Further, in the MediaTek Dimensity 9000, the ARMv9 L3 cache comprises slices that are distributed across the SoC or processor die. The DSU-110 includes an interconnect bus, such as a ring-based transport network, which enables the processor cores to access the L3 cache slices⁷⁸:

Cache Slices

To be able to accommodate the new cores and future cores, Arm went with a large amount of cache. But at the same time, they are looking to considerably increase the bandwidth which means a lot more simultaneous cache accesses taking place in parallel. Their solution was to split it up into slices. Each slice makes up a portion of the L3 cache, includes part of the snoop filter along with the associated control logic. The actual number of cache slices is configurable on the DSU-110 – with up to eight slices supported (as the number of cores). Depending on the bandwidth requirements and the market, Arm partners can choose to configure it to best suit their needs. During cache access from a core, the target cache slice is chosen based on the hash of the address. It's, therefore,

⁷⁸ David Schor, *Arm Launches The DSU-110 For New Armv9 CPU Clusters*, WikiChip Fuse (May, 25, 2021), <https://fuse.wikichip.org/news/5270/arm-launches-the-dsu-110-for-new-armv9-cpu-clusters/>; ARM DynamIQ Shared Unit-110 Technical Reference Manual at 27, 101, 102.



Internal Interconnects

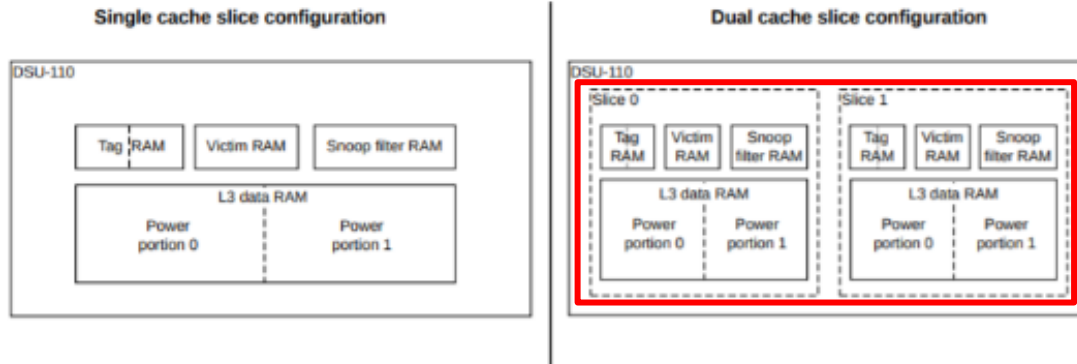
With the larger cache support and a high number of transactions, the interconnect implementation in the DSU was designed to allow data to flow quickly within the cluster. On the previous DSU, Arm used a hybrid crossbar implementation. Addressing this, the new DSU-110 uses a ring-based transport network to connect the cores to the slices and to the bus interface. But interestingly, Arm discovered that a single traditional ring was insufficient to meet the performance targets they wanted. For that

7.6 Cache slices and portions

The L3 cache of the *DynamiQ™ Shared Unit-110* (DSU-110) can be divided into up to eight identical slices, each containing between 256KB and 2MB of the cache. A cache slice consists of the data, tag, victim, and snoop filter RAMs and associated logic. A portion is a further subdivision of RAM in a cache slice.

For each cache slice, both the data RAM and tag RAM is subdivided into two portions.

The following figure shows the differences between a single and a dual cache slice configuration.

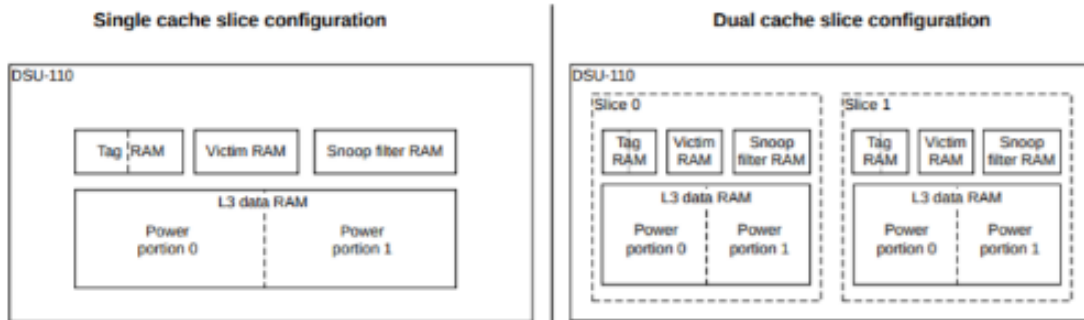
Figure 7-2: Comparison between a single and dual L3 cache slice configuration

Splitting the L3 cache into slices provides the following advantages:

- Improving the physical floorplan when implementing the macrocell, by ensuring that the RAMs are located close to the logic that is controlling them.
- Increasing the bandwidth because the slices can be accessed in parallel.

When a core type can be defined as part of a complex, then all instances of that core type (in the cluster) are implemented as complexes. This is either as part of a single-core complex or dual-core complex. Having all instances of a core type formed into complexes within the cluster, ensures consistent clock and power management control.

Within a dual-core complex, logic such as a Vector Processing Unit, L2 Translation Lookaside Buffer (TLB), and L3 cache logic is shared between the cores and is collectively known as Shared logic. In a single-core complex, the same logic resides outside the core but is collectively known as Dedicated logic.

Figure 7-2: Comparison between a single and dual L3 cache slice configuration

Splitting the L3 cache into slices provides the following advantages:

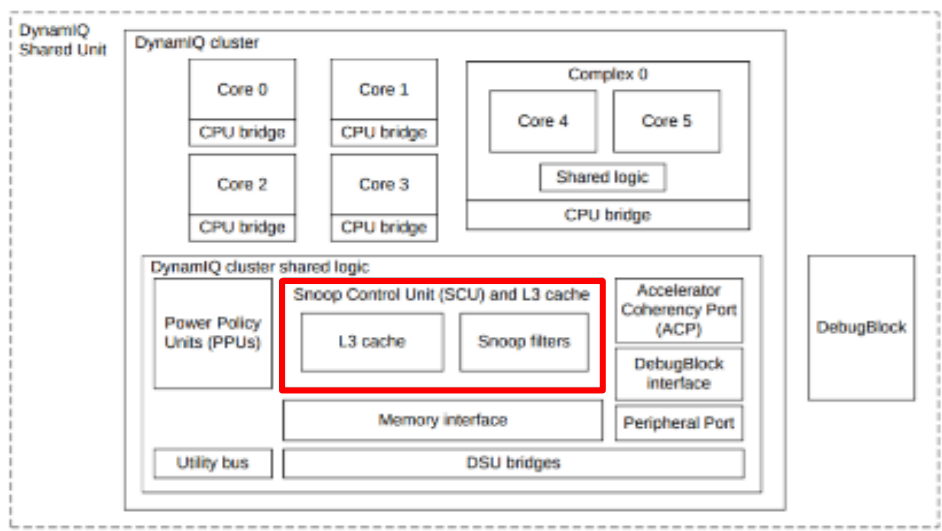
- Improving the physical floorplan when implementing the macrocell, by ensuring that the RAMs are located close to the logic that is controlling them.
- Increasing the bandwidth because the slices can be accessed in parallel.

230. Each of the '960 Patent Accused Products comprises cache management circuitry operative to provide coherent, non-uniform access to the plurality of distributed cache portions by the plurality of cores.

231. For example, the Dimensity 9000 SoCs include DynamIQ cluster shared logic that includes a Snoop Control Unit (SCU), which is the primary cache management module (e.g., the cache management circuitry) for the system. The SCU interfaces with at least the L3 cache interconnect network to manage the L3 cache and provide coherent access across all DynamIQ clusters⁷⁹:

⁷⁹ Arm DynamIQ Shared Unit-110 Technical Reference Manual at 27, 102; ARM® DynamIQ™ Shared Unit Technical Reference Manual.

Figure 2-1: DSU-110 DynamIQ™ cluster




All cores in the DSU-110 DynamIQ™ cluster, including those in complexes, are coherently connected to an L3 memory system that includes an L3 cache and a *Snoop Control Unit* (SCU). The SCU maintains coherency between caches in the cores and the L3 cache, and includes a snoop filter to optimize coherency maintenance operations. The shared L3 cache simplifies process migration between the cores.

All the external interfaces including those to the cores are provided through the DSU-110 to the *System on Chip* (SoC). Main system transactions are supported through the memory interface which can be implemented as a coherent or non-coherent interface. A Peripheral port is provided to support low latency access to external system components but also can be used as a non-coherent master interface. The *Accelerator Coherency Port* (ACP) provides coherent access for non-cached masters that need I/O coherency with the cluster. The Utility bus is a memory-mapped port that provides a programming interface to the PPU's and some of the other system components.

232. Further, '960 Accused Products comprise L3 cache slices that are physically distributed across the SoC die and are connected via an interconnect bus. Each cache slice is assigned to a group of cores. When a core faces a cache-miss in the associated cache slice portion, it fetches the cache line from a different core slice, if available, which increases the latency. Since the latency associated with access to cache lines in the L3 cache depends upon which cache slice

stores the cache line, the cache management circuitry provides non-uniform access of the cache slices to the cores⁸⁰:



The shared L3 cache of the DSU-110 provides the following functionality:

- A dynamically optimized cache allocation policy, which is typically exclusive. This cache allocation policy means that in normal use, a line is either in the cache of one or more cores (or complexes) or in the L3 cache, but not in both caches. Only Cacheable, shareable memory locations are allocated in the L3 cache. Non-shareable memory locations are not allocated in the L3 cache.
- Groups of cache ways can be partitioned and assigned to processes² by the *Memory System Resource Partitioning and Monitoring (MPAM)* architecture extension. Cache partitioning ensures that each process does not dominate the use of the cache to disadvantage other processes.

7.3 L3 cache partitioning

The L3 cache supports a partitioning scheme that alters the victim selection policy to prevent processes from using the entire L3 cache to the disadvantage of other processes.

Cache partitioning is intended for specialized software where there are distinct classes of processes running with different cache accessing patterns. For example, two processes A and B run on separate cores in the same cluster and therefore share the L3 cache. If process A is more data-intensive than process B, then process A can cause all the cache lines that process B allocates to be evicted. Evicting these allocated cache lines can reduce the performance of process B.

⁸⁰ Schor, *supra* note 70; Arm DynamIQ Shared Unit-110 Technical Reference Manual at 19, 35, 96, 97.

Snoop Control Unit (SCU)

The SCU maintains coherency between all the data caches in the cluster.

The SCU contains buffers that can handle direct cache-to-cache transfers between cores without having to read or write data to the L3 cache. Cache line migration enables dirty lines to be moved between cores. Also, there is no requirement to write back transferred cache line data to the L3 cache.

All cores in the DSU-110 DynamIQ™ cluster, including those in complexes, are coherently connected to an L3 memory system that includes an L3 cache and a Snoop Control Unit (SCU). The SCU maintains coherency between caches in the cores and the L3 cache, and includes a snoop filter to optimize coherency maintenance operations. The shared L3 cache simplifies process migration between the cores.

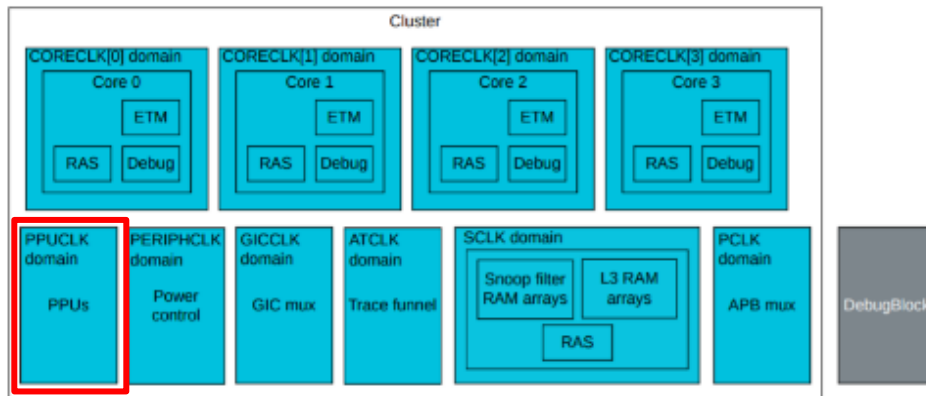
233. Each of the '960 Patent Accused Products comprises power management circuitry operative to enable a first frequency of operation for a first cluster of the plurality of cores.

234. For example, the Dimensity 9000's DSU-110 includes Power Policy Units (PPU, "power management circuitry") and a Power Control Module that provides DVFS control on per-core and per-cluster level. The Power Control Module also manages power consumption of the L3 distributed cache through power-gating of the cores as well as the L3 cache.⁸¹

different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.

⁸¹ Nayak, et al., *supra* note 2; Arm DynamIQ Shared Unit-110 Technical Reference Manual at 49, 51.

Figure 4-2: DSU-110 pin-controlled reset domains



The Power Policy Units (PPUs) for the cluster and each of the cores are used to control the power management features of the cluster and cores using a software interface. This includes managing various power states and transitions between these states. Certain power mode changes, for example powering up the cluster from a powered down state, include implicit resets to internal logic.

5.1 Power management in the DSU-110

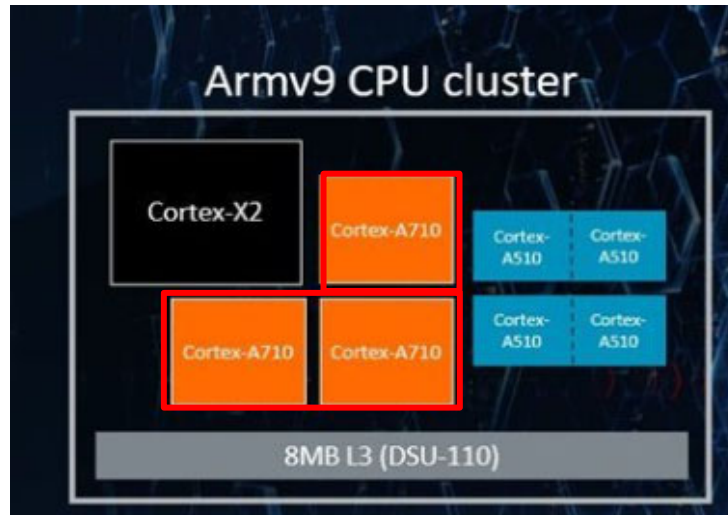
The DynamIQ™ Shared Unit-110 (DSU-110) provides various mechanisms to control both dynamic and static power dissipation. These mechanisms are associated with a set of power domains, power modes, and operational modes. Some of these mechanisms are brought under software control using Power Policy Units (PPUs).

The power management techniques employed by the DSU-110 and cores in the cluster include:

- Internal core clock gating where different internal parts of the core are clock idle
- Per-core Dynamic Voltage and Frequency Scaling (DVFS)
- Powerdown
- Retention, a low-power mode that retains the register and RAM state

235. Further, the '960 Patent Accused Products include processor clusters such as clusters of ARM Cortex-A510 cores (“first cluster”) and clusters of ARM Cortex-A710 cores (“second cluster”). Cores within each ARM cluster are located proximate to each another, as compared to distance from cores from a different cluster. Cores within a given cluster are typically


assigned to the same independent frequency (e.g., first frequency and second frequency) and voltage domains⁸²:



Cluster features


The DSU-110 has the following cluster features:

- Support for Arm[®]v9.0-A architecture cores
- Support for up to three types of core, and a maximum of eight cores in the cluster
- Power Policy Units (PPUs) providing autonomous power management of the L3 cache and the cores
- Support for cores running independently at different frequencies and voltages known as Dynamic Voltage Frequency Scaling (DVFS). For cores in a complex, DVFS is only possible for the whole complex not for individual cores.



⁸² Bedi, *supra* note 20; Arm DynamIQ Shared Unit-110 Technical Reference Manual at 18, 20.

Figure 2-1: DSU-110 DynamIQ™ cluster



A DSU-110 DynamIQ™ cluster consists of between one and eight cores, with up to three different types of cores in the same cluster. Cores can be configured for various performance points during macrocell implementation and run at different frequencies and voltages.

Cluster features

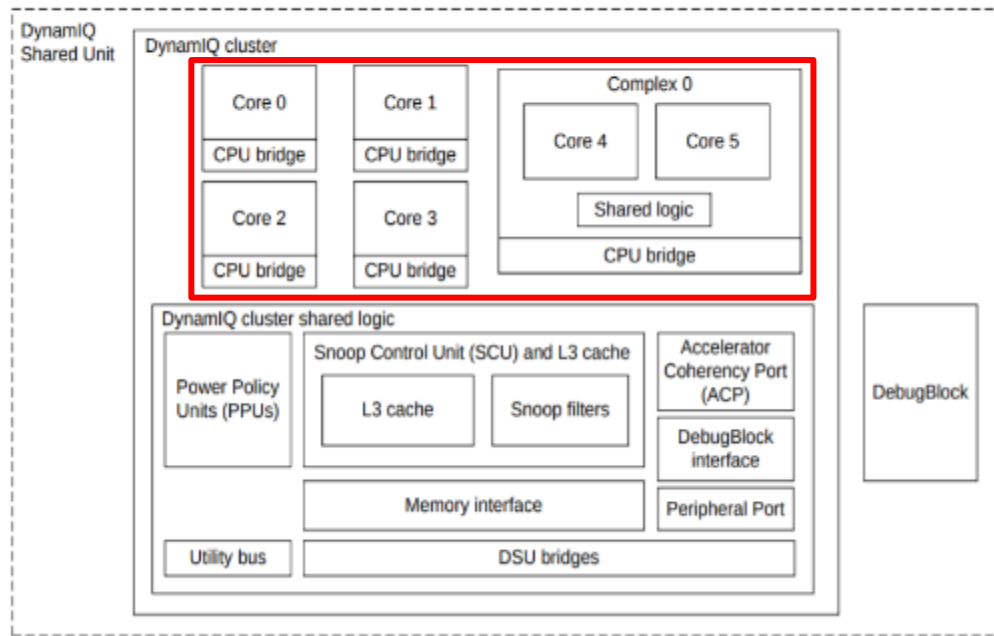
The DSU-110 has the following cluster features:

- Support for Arm®v9.0-A architecture cores
- Support for up to three types of core, and a maximum of eight cores in the cluster
- Power Policy Units (PPUs) providing autonomous power management of the L3 cache and the cores
- Support for cores running independently at different frequencies and voltages known as Dynamic Voltage Frequency Scaling (DVFS). For cores in a complex, DVFS is only possible for the whole complex not for individual cores.
- The DSU-110 has an internal transport mechanism that is responsible for all communication between components in the design. The topology of the transport is defined by the number of cores and number of L3 cache slices.

However, as with all of Arm's "LITTLE" CPUs, efficiency is still king. Not only does Cortex-A510 boost power efficiency by up to 20 percent (ISO process)⁴ through the 3-wide in-order design, but it also provides industry-leading area efficiency. An innovation that makes this possible is merged core microarchitecture. This allows two Cortex-A510 CPUs to be grouped into a complex, with multiple complexes per CPU cluster.

236. Each of the '960 Patent Accused Products comprises cores that, within any given cluster, are closer to each other than they are to cores from other clusters. Accordingly, the average distance between cores in at least one of the clusters will be less than the average distance between the plurality of cores⁸³:

Figure 2-1: DSU-110 DynamIQ™ cluster



⁸³ Arm DynamIQ Shared Unit-110 Technical Reference Manual at 18, 20; Bedi, *supra* note 20.

However, as with all of Arm's "LITTLE" CPUs, efficiency is still king. Not only does Cortex-A510 boost power efficiency by up to 20 percent (ISO process)⁸⁴ through the 3-wide in-order design, but it also provides industry-leading area efficiency. An innovation that makes this possible is merged core microarchitecture. This allows two Cortex-A510 CPUs to be grouped into a complex, with multiple complexes per CPU cluster.

237. Each of the '960 Patent Accused Products employs power management circuitry that is operative to selectively gate power to the first cluster of the plurality of cores and distributed cache portions of the at least one higher level distributed cache that correspond to the first cluster and/or the second cluster of the plurality of cores and distributed cache portions of the at least one higher level distributed cache that correspond to the second cluster.

238. For example, the Dimensity 9000's PPU ("power management circuitry") provides advanced power management features including selectively reducing power to individual CPU cluster cores as well as the L2 cache through DVFS. Additionally, individual L3 cache slices can also be partially powered down by the PPU⁸⁴:

different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.

⁸⁴ Rosinger & Pradhan, *supra* note 14, pp. 8, 10; Nayak, et al., *supra* note 2; Arm DynamIQ Shared Unit-110 Technical Reference Manual at 49, 51, 58.






Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

Arm DynamIQ Shared Unit-110


- DSU-110 for Cache Coherency & Shared L3\$
- Optimized Ring Transport Network
 - Bi-directional, Dual-ring
- 2X Bandwidth and 25% Lower Leakage
 - Support Partial SRAM & Logic Shutdown
- Cache Partition for QoS



MEDIATEK

11

Figure 4-2: DSU-110 pin-controlled reset domains



The *Power Policy Units (PPUs)* for the cluster and each of the cores are used to control the power management features of the cluster and cores using a software interface. This includes managing various power states and transitions between these states. Certain power mode changes, for example powering up the cluster from a powered down state, include implicit resets to internal logic.

5.1 Power management in the DSU-110

The *DynamiQ™ Shared Unit-110 (DSU-110)* provides various mechanisms to control both dynamic and static power dissipation. These mechanisms are associated with a set of power domains, power modes, and operational modes. Some of these mechanisms are brought under software control using *Power Policy Units (PPUs)*.

The power management techniques employed by the DSU-110 and cores in the cluster includes:

- Internal core clock gating where different internal parts of the core are clock idle
- Per-core *Dynamic Voltage and Frequency Scaling (DVFS)*
- Powerdown
- Retention, a low-power mode that retains the register and RAM state

5.4 L3 RAM power control

In addition to retention features, the *DynamiQ™ Shared Unit-110 (DSU-110)* can further reduce static leakage power using two powerdown features. Firstly, optionally power down of all but one of the L3 cache slices. Secondly, within each L3 cache slice powerdown a portion of the L3 cache RAM that the cache slice contains.

5.4.1 L3 cache RAM powerdown

The L3 cache RAMs typically contribute to a large proportion of the total leakage power, particularly for large cache sizes. Therefore, it is beneficial to be able to power down the RAMs when only some of the L3 cache is required, but it also results in reducing cache capacity. Parts of the L3 cache RAM can be independently powered down to reduce RAM leakage power when not in use. L3 cache powerdown is controlled by the cluster *Power Policy Unit (PPU)*.

The L3 cache RAM powerdown feature allows the RAMs to be powered down in groups of ways, giving options of 100%, 50%, or 0% of the L3 cache capacity. When a workload is making light use of the L3 cache, then this can be detected and the L3 cache capacity reduced without significant impact on the performance. For example, this can occur when the L3 cache has a relatively small memory footprint that mostly fits within the L2 cache.

239. Further, the Dimensity 9000's PPU controls chip-level power consumption by reducing/increasing/gating power delivered to clusters, L1 & L2 caches and L3 cache slices based on system requirements and operating conditions (“selectively gating power”)⁸⁵:

Cluster features

The DSU-110 has the following cluster features:

- Support for Arm®v9.0-A architecture cores
- Support for up to three types of core, and a maximum of eight cores in the cluster
- **Power Policy Units (PPUs) providing autonomous power management of the L3 cache and the cores**

The DSU-110 DynamIQ™ cluster can be implemented with various power domains to target power performance levels. These power domains are managed through the *Power Policy Units* (PPUs). The DSU-110 DynamIQ™ cluster supports many mechanisms to reduce static and dynamic power dissipation. For example, placing the cores and L3 cache into retention and powering down parts of the L3 cache.

6.6 Programming sequences for the cluster and the core

Example *Power Policy Unit* (PPU) programming sequences are provided for both the cluster and the cores. One of these sequences uses the static mode policy to demonstrate programming using this policy. However, because static power management can require considerable activity from the System Control Processor, Arm strongly recommends that you use dynamic power management for normal operation of the cluster.

6.6.1 Programming sequence to bring the cluster and cores from Off to On mode

Use the following steps, to program the *Power Policy Unit* (PPU) for the DSU-110 DynamIQ™ cluster and each of the cores to request a change of PPU mode from Off mode to On mode.

6.6.2 Programming sequence to bring the cluster and cores from On to Off mode

Use the following steps, to program the *Power Policy Unit* (PPU) for the DSU-110 DynamIQ™ cluster and each of the cores to request a change of PPU mode from On to Off.

⁸⁵ Arm DynamIQ Shared Unit-110 Technical Reference Manual at 20, 89.

240. Further, the L3 cache slices can be partially powered down based on the system workload⁸⁶:

Cache features

The DSU-110 has the following cache features:

- Optional unified 16-way set-associative L3 cache, configurable from 256KB to 16MB
- 64-byte cache lines
- L3 cache slice support, for improved bandwidth and cache RAM layout, up to eight slices supported
- L3 cache powerdown based either on cache slices or cache ways

5.4 L3 RAM power control

In addition to retention features, the *DynamIQ™ Shared Unit-110* (DSU-110) can further reduce static leakage power using two powerdown features. Firstly, optionally power down of all but one of the L3 cache slices. Secondly, within each L3 cache slice powerdown a portion of the L3 cache RAM that the cache slice contains.

5.4.1 L3 cache RAM powerdown

The L3 cache RAMs typically contribute to a large proportion of the total leakage power, particularly for large cache sizes. Therefore, it is beneficial to be able to power down the RAMs when only some of the L3 cache is required, but it also results in reducing cache capacity. Parts of the L3 cache RAM can be independently powered down to reduce RAM leakage power when not in use. L3 cache powerdown is controlled by the cluster *Power Policy Unit* (PPU).

5.4.2 L3 cache slice powerdown

In addition to powering down the L3 cache RAMs, you can gain further leakage savings by powering down some of the L3 cache slice control logic as well. Control of powering up or powering down L3 cache slices is performed by the cluster *Power Policy Unit* (PPU).

⁸⁶ Arm DynamIQ Shared Unit-110 Technical Reference Manual at 19, 20, 54, 58, 62.

5.3 Cluster power modes

The DSU-110 DynamiQ™ cluster and each of the cores and complexes in the cluster have a defined set of power modes and corresponding legal transactions between these modes.

The following table shows the supported power modes for the DSU-110 DynamiQ™ cluster.

Power mode	Short name	Description
On mode	ON	On mode is the normal mode of operation where all cluster functionality is available.
Off mode	OFF	In Off mode, power is removed from the cluster logic and all the RAMs.
Functional retention mode	FUNC_RET	In Functional retention mode, the L3 cache RAMs and snoop filter RAMs are placed in a retention state. Data is retained in these RAMs. The rest of the DynamiQ™ cluster shared logic remains powered up.

5.4.1 L3 cache RAM powerdown

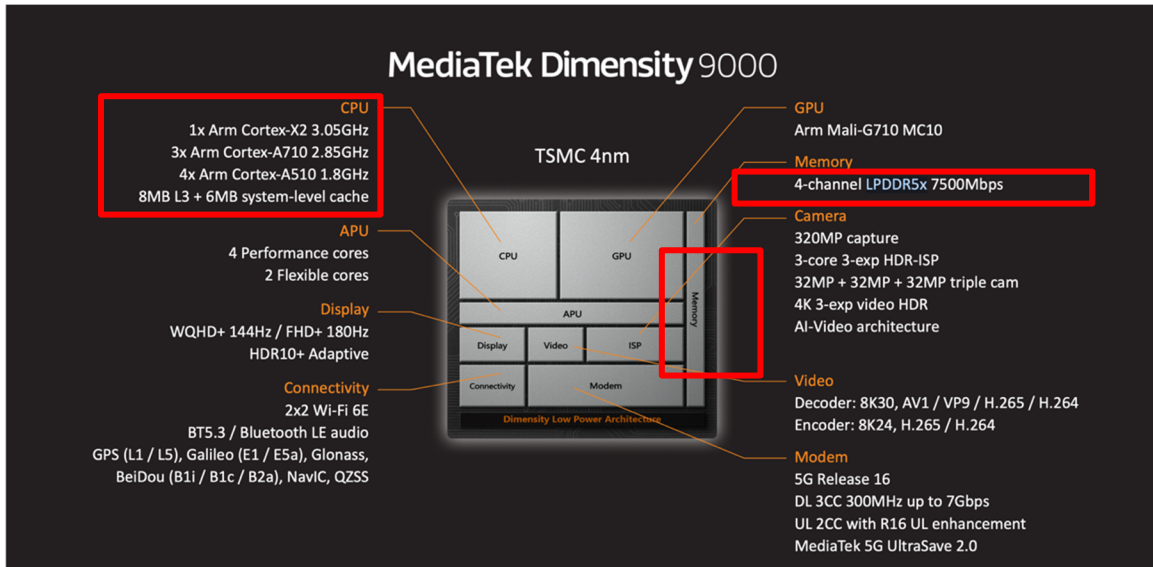
The L3 cache RAMs typically contribute to a large proportion of the total leakage power, particularly for large cache sizes. Therefore, it is beneficial to be able to power down the RAMs when only some of the L3 cache is required, but it also results in reducing cache capacity. Parts of the L3 cache RAM can be independently powered down to reduce RAM leakage power when not in use. L3 cache powerdown is controlled by the cluster *Power Policy Unit* (PPU).

The L3 cache RAM powerdown feature allows the RAMs to be powered down in groups of ways, giving options of 100%, 50%, or 0% of the L3 cache capacity. When a workload is making light use of the L3 cache, then this can be detected and the L3 cache capacity reduced without significant impact on the performance. For example, this can occur when the L3 cache has a relatively small memory footprint that mostly fits within the L2 cache.

241. Each of the '960 Patent Accused Products comprises a first integrated memory controller coupled with the symmetric multi-threaded cores described above.

242. For example, the Dimensity 9000 SoCs include memory controllers for four channels of LPDDR5x memory⁸⁷:

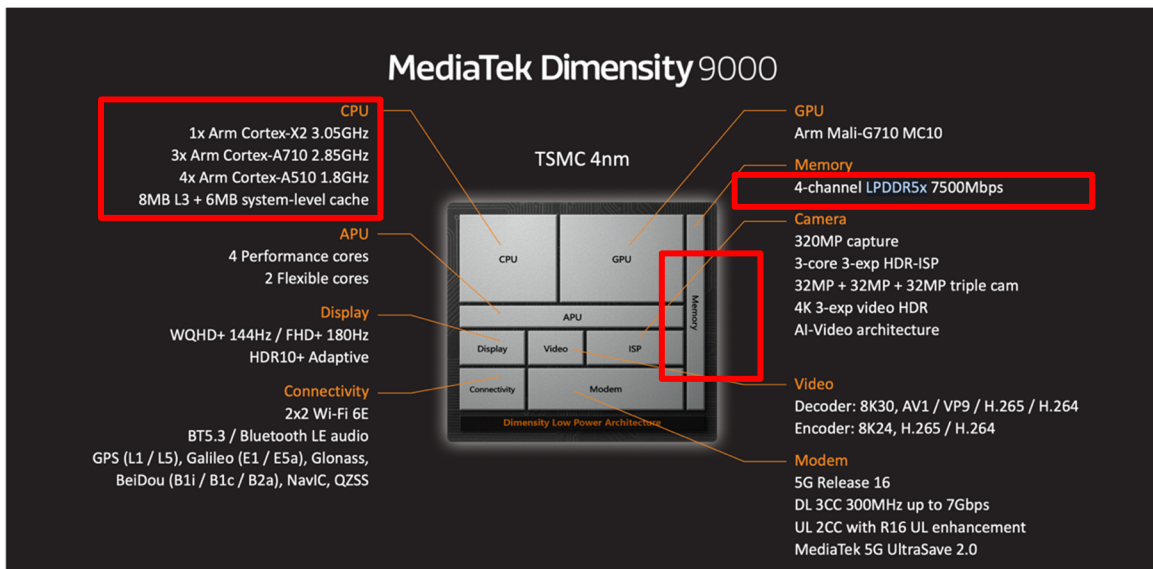
⁸⁷ Rosinger & Pradhan, *supra* note 14, p. 3; Nayak, et al., *supra* note 2.



4K video at 60fps; camera resolutions up to 320MPixels are supported. LPDDR5-6400/LPDDR5X-7500 memory interfaces facilitate up to 24GB of external SDRAM over four 16b channels for a peak transfer rate of 0.46Tb/s.

243. Each of the '960 Patent Accused Products comprises a second integrated memory controller coupled with the symmetric multi-threaded cores.

244. For example, the Dimensity 9000 SoCs include memory controllers for four channels of LPDDR5x memory⁸⁸:



⁸⁸ Rosinger & Pradhan, *supra* note 14, p. 3; Nayak, et al., *supra* note 2.

4K video at 60fps; camera resolutions up to 320MPixels are supported. LPDDR5-6400/LPDDR5X-7500 memory interfaces facilitate up to 24GB of external SDRAM over four 16b channels for a peak transfer rate of 0.46Tb/s.

245. Further, on information and belief, MediaTek has actively induced and/or contributed to infringement of at least Claim 15 of the '960 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

246. Users of the '960 Patent Accused Products directly infringe at least Claim 15 of the '960 Patent when they use the '960 Patent Accused Products in the ordinary, customary, and intended way. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the '960 Patent Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the '960 Patent Accused Products to consumers within the United States and instructing and encouraging such customers to use the '960 Patent Accused Products in the ordinary, customary, and intended way, which MediaTek knew infringes at least Claim 15 of the '960 Patent, or, alternatively, was willfully blind to the infringement.

247. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the '960 Patent Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the '960 Patent Accused Products in the United States, which MediaTek knew infringes at least Claim 15 of the '960 Patent, or, alternatively, was willfully blind to the infringement.

248. On information and belief, in violation of 35 U.S.C. § 271(c), MediaTek's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of at least Claim 15 of the '960 Patent, constituting a material part of the invention. On information and belief, MediaTek knows and has known the same to be especially made or especially adapted for use in an infringement of the '960 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

249. On information and belief, in violation of 35 U.S.C. § 271(f)(1), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 15 of the '960 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

250. On information and belief, in violation of 35 U.S.C. § 271(f)(2), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 15 of the '960 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

251. MediaTek is not licensed or otherwise authorized to practice the claims of the '960 Patent.

252. Thus, by its acts, MediaTek has injured Daedalus and is liable to Daedalus for directly and/or indirectly infringing one or more claims of the '960 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 15.

253. On information and belief, MediaTek has known about the '960 Patent at least since August 23, 2022.⁸⁹ At a minimum, MediaTek has knowledge of the '960 Patent at least as of the filing of this Complaint. Accordingly, MediaTek's infringement of the '960 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

254. As a result of MediaTek's infringement of the '960 Patent, Daedalus has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for MediaTek's infringement, but in no event less than a reasonable royalty with interest and costs.

255. On information and belief, MediaTek will continue to infringe the '960 Patent unless enjoined by this Court. MediaTek's infringement of Daedalus' rights under the '960 Patent will continue to damage Daedalus, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

EIGHTH COUNT

(Infringement of U.S Patent No. 10,725,919)

256. Daedalus incorporates by reference the allegations set forth in Paragraphs 1-255 of the Complaint as though fully set forth herein.

257. The claims of the '919 Patent are valid and enforceable.

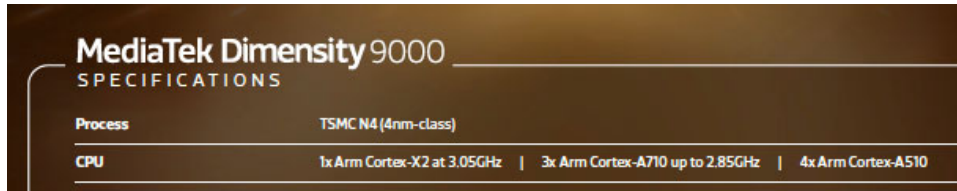
⁸⁹ *Daedalus Prime LLC v. Mazda Motor Corp., et al.*, No. 22-cv-01108 (D. Del. Aug. 23, 2022).

258. On information and belief, in violation of 35 U.S.C. § 271(a), MediaTek has directly infringed and continues to directly infringe one or more claims of the '919 Patent, including at least Claim 16 of the '919 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '919 Patent, including but not limited to its electronic devices containing SoCs or microprocessors based on or derived from ARMv8.2 architecture, as well as subsequent revisions to the ARM architecture such as the ARMv9 architecture, such as the Dimensity 9000 SoCs, and all reasonably similar products (the "'919 Patent Accused Products").

259. Each of the '919 Patent Accused Products comprises a processor. For example, the Dimensity 9000 SoC contains one or more microprocessors based on or derived from the ARM Cortex-X2 architecture, the ARM Cortex-A710 architecture, and the ARM Cortex-A510 architecture.

260. Each of the '919 Patent Accused Products comprises a plurality of cores, the plurality of cores comprising symmetric multi-threaded cores.

261. For example, Dimensity 9000 SoCs include multiple clusters with several clusters having several identical core architectures with identical performance specifications, including 4 x Cortex-A510 cores, and 3 x Cortex-A710 cores. These are symmetric since they contain multiple processors cores having identical specifications and configurations⁹⁰:



MediaTek Dimensity 9000	
SPECIFICATIONS	
Process	TSMC N4 (4nm-class)
CPU	1x Arm Cortex-X2 at 3.05GHz 3x Arm Cortex-A710 up to 2.85GHz 4x Arm Cortex-A510


⁹⁰ <https://mediatek-marketing.files.svcdn.com/production/documents/Dimensity-9000-Infographic.pdf>; <https://i.mediatek.com/dimensity-9000>; Nayak, et al., *supra* note 2; ARM®, *supra* note 63; ARM®, *supra* note 63; <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8>, page 18.

arm

WORLD'S 1ST CORTEX-X2 IN A SMARTPHONE CHIP

The Dimensity 9000 uses new Armv9 architecture CPUs and GPU to deliver unparalleled performance. Its octa-core CPU includes an Arm Cortex-X2 that bursts to epic 3GHz, while new LPDDR5X memory makes data immediately available, eliminating the wait to give immediate responsiveness in any app, whatever you're doing.

- Ultra-Core - 1x Arm Cortex-X2 at 3.05GHz
- Super-Cores - 3x Arm Cortex-A710 up to 2.85GHz
- Efficiency Cores - 4x Arm Cortex-A510
- World's first Arm Mali-G710 MC10 graphics processor
- Big caches - 8MB L3 cache + 6MB system cache
- LPDDR5X 7500Mbps support – 20% more power efficient than LPDDR5



Core	Focus	Microarchitectural Highlights
Cortex-X2	Peak Performance	Branch prediction : Decoupled from fetch, improved accuracy and prediction SVE 128b vector length implementation bfloat16, int8 matmul data types support for ML apps Reduced a pipeline stage at dispatch Larger Out of order window size, load-store window/structures. d-TLB Data prefetch enhancements for accuracy and coverage
A710	Sustained Performance	Branch prediction : improved accuracy, capacity doubled for BTB, GHB Increased capacity of L1 instruction TLB Improved efficiency with reduction of mid core width Reduced a pipeline stage at dispatch Reduced DSU access, DRAM refills Data prefetch enhancements for accuracy and coverage
A510	Power Efficiency	Two cores grouped in a complex, multiple complexes per cluster L2 cache, L2 TLB, vector paths data path shared across a complex Fine-grained scheduling between the cores in the complex 3-wide decode and issue, 3 integer ALU pipeline state of art multi-stage branch prediction Advance data prefetchers for accuracy and coverage

Figure 2.5.2: ARMv9 CPU cluster.

L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.


All processor cores in the CPU subsystem incorporate the ARMv9 instruction set with key architectural advances. Memory Tagging Extension (MTE) enables greater security by locking data in the memory using a tag which can only be accessed by the correct key held by the pointer accessing the memory location, as shown in Fig. 2.5.1. Further, a Scalable Vector Extension 2 (SVE2) allows a scalable vector length in multiples of 128b, up to 2048b, enabling increased DSP and ML vector-processing capabilities, as shown in Fig. 2.5.1.

Symmetric multi-processing

Symmetric multi-processing (SMP) is a software architecture that dynamically determines the roles of individual processors. Each core in the cluster has the same view of memory and of shared hardware. Any application, process or task can run on any core and the operating system scheduler can dynamically migrate tasks between cores to achieve optimal system load.

We expect that readers will be familiar with the fundamental operating principles of an OS, but OS terminology will be briefly reviewed here. An application that executes under an operating system is known as a *process*. It performs many operations through calls to the system library that provides certain functions from library code, but also acts a wrapper for system calls to kernel operations. Individual processes have associated resources, including stack, heap and constant data areas, and properties such as scheduling priority settings. The kernel view of a process is called a *task*.

Figure 2-1: DSU-110 DynamIQ™ cluster



262. Further, ARM cores are multi-threaded by design as indicated in the ARM DynamIQ™ specification, and are capable of simultaneously executing two or more processing threads⁹¹:

The DSU-110 DynamIQ™ cluster debug system supports both single and multi-threaded cores.

The Arm architecture allows for cores to be single, or multi-threaded. A *Processing Element (PE)* performs a thread of execution. A single-threaded core has one PE and a multi-threaded core has two or more PEs. Because the debugging system allows individual threads to be debugged, the term PE is used throughout this chapter. Where a reference to a core is made, the core can be a single, or multi-threaded core.

⁹¹ <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8?token=> Pages 168, 379; <https://developer.arm.com/documentation/101800/latest>.

Processing element
 The Arm architecture allows for cores to support multiple PEs.
 The *DynamiQ™ Shared Unit-110 (DSU-110)* supports cores with multiple PEs. Where a reference to a core is made, the core could be a core with only one PE (single-threaded core) or multiple PEs (multi-threaded core).

Bits	Name	Description	Reset
[24]	MT	Multithreaded PE. Performance of PEs at the lowest affinity level is largely independent. If ERRDLVAFF.AHO is not valid, this bit is not valid and reads as UNKNOWN.	0b0

263. Further, other cores in the '919 Patent Accused products contain specific support for the use of multiple threads.

264. Further, the ARMv8 programmer's guide describes features from the architecture that are designed to assist in the execution of multiple threads⁹²:

The additions to the architecture mean that a single physical core can execute code from both the Normal world and the Secure world in a time-sliced fashion, although this depends on the availability of interrupt-generating peripherals that can be configured to be accessible only by the Secure World. For example, a Secure timer interrupt could be used to guarantee some execution time for the Secure world, in a manner resembling preemptive multitasking. Such peripherals may or may not be available, depending on the level of security and use cases that the platform designer intends to support.

hardware. Any application, process, or task can run on any core and the operating system scheduler can dynamically migrate tasks between cores to achieve optimal system load. A multi-threaded application can run on several cores at once. The operating system can hide much of the complexity from applications.

265. Further, MediaTek has announced that its SoCs are optimized and ready to go for the Android operating system software.⁹³


266. Further, Android operating system software compatible with the '919 Patent Accused Products contains support for more than one thread of execution within a process, wherein

⁹² ARM®, *supra* note 63.
⁹³ MediaTek, *supra* note 66.

those threads run concurrently, and contains support for the use of multiple threads within programs such as worker threads.⁹⁴

267. Each of the '919 Patent Accused Products comprises a cache subsystem, the cache subsystem comprising a plurality of first-level caches and at least one higher-level distributed cache comprising a plurality of distributed cache portions that are physically distributed across a die and shared by the plurality of cores, each first-level cache integral to one of the plurality of cores and each distributed cache portion accessible to each of the plurality of cores.

268. For example, the MediaTek Dimensity 9000 includes a cache subsystem comprising L1, L2 and L3 caches, including a plurality of first level (L1) caches that are integral to one of the plurality of cores⁹⁵:



⁹⁴ <https://developer.android.com/courses/extras/multithreading>;
<https://developer.android.com/guide/components/processes-and-threads>.

⁹⁵ Rosinger & Pradhan, *supra* note 14, pp. 8, 10; Nayak, et al., *supra* note 2; Arm® Cortex®-A510 Core Technical Reference Manual, p. 33; Arm® Cortex®-A710 Core Technical Reference Manual, p. 32.

CPU Highlights vs. Dimensity 1200

- Arm Cortex-X2 +40% integer performance over Arm Cortex-A78
- Arm Cortex-A510 +35% integer performance over Arm Cortex-A55
- -50% CPU power @ iso-performance
- Geekbenchv5 single-thread 1278 (+36%) , multi-thread 4400 (+33%)

Cortex-X2 1MB L2	Cortex-A7 10 512KB L2	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10	Cortex-A5 10
Cortex-A7 10 512KB L2	Cortex-A7 10 512KB L2	Shared 512KB L2\$ Shared VPU		Shared 512KB L2\$ Shared VPU	
DynamIQ Shared Unit equipped 8MB L3\$					



8

L1 instruction cache, 64KB L1 data cache, and a 1MB private L2 cache. The 2nd gear consists of three Balanced Performance (BP) cores utilizing the ARMv9 Cortex-A710 architecture, each with a 64KB L1 instruction cache, a 64KB L1 data cache, and a 512KB private L2 cache. The 3rd gear features four High Efficiency (HE) ARMv9 Cortex-A510 cores [1], with each core using a 64KB L1 instruction cache, 64KB L1 data cache. Further, the HE CPU cores are implemented in pairs to facilitate the sharing of a 512KB L2 cache, floating-point and vector hardware between two CPUs cores, improving area and power efficiency, maintaining full v9 compatibility, without sacrificing performance of key workloads. Finally, an 8MB L3 cache is shared across all the cores of the CPU complex.

Figure 3-1: Cortex®-A510 core components





Figure 3-1: Cortex®-A710 core components



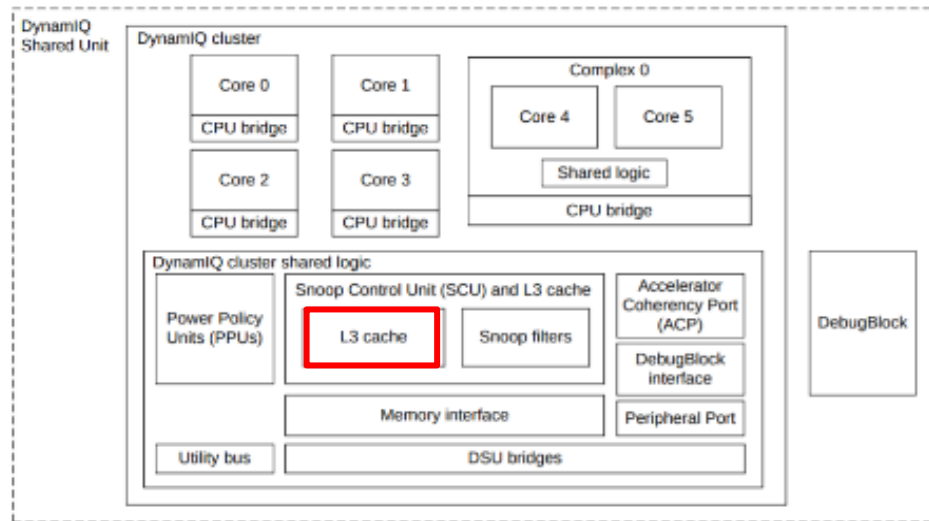
269. Further, the MediaTek Dimensity 9000 also includes a shared L3 cache (e.g., the higher level cache) as part of the DynamIQ Shared Unit - 110. The L3 cache is split into L3 cache slices (or a plurality of distributed cache portions), which are distributed across the SoC die and accessible to each of the plurality of cores⁹⁶:

2 The DynamIQ™ Shared Unit-110

The DynamIQ™ Shared Unit-110 (DSU-110) provides a shared L3 memory system, snoop control and filtering, and other control logic to support a cluster of A-class architecture cores. The cluster is called the DSU-110 DynamIQ™ cluster. Additionally, all the external interfaces to System on Chip (SoC) are provided through the DSU-110.

⁹⁶ Arm DynamIQ Shared Unit-110 Technical Reference Manual at 18, 101; Arm® Cortex®-A710 Core Technical Reference Manual, p. 32.

Figure 2-1: DSU-110 DynamiQ™ cluster

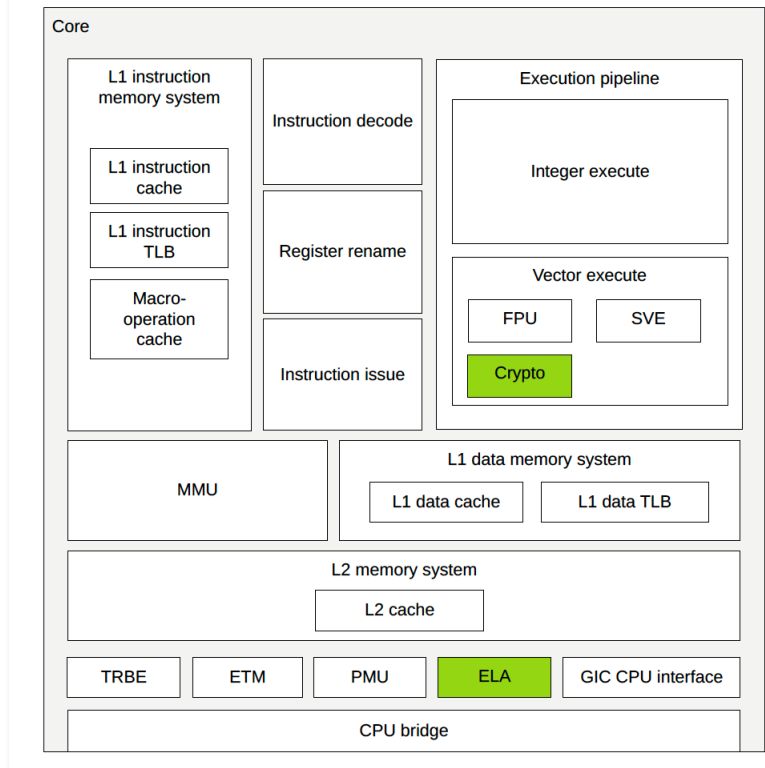


7.6 Cache slices and portions

The L3 cache of the *DynamiQ™ Shared Unit-110* (DSU-110) can be divided into up to eight identical slices, each containing between 256KB and 2MB of the cache. A cache slice consists of the data, tag, victim, and snoop filter RAMs and associated logic. A portion is a further subdivision of RAM in a cache slice.

For each cache slice, both the data RAM and tag RAM is subdivided into two portions.

Figure 3-1: Cortex®-A710 core components

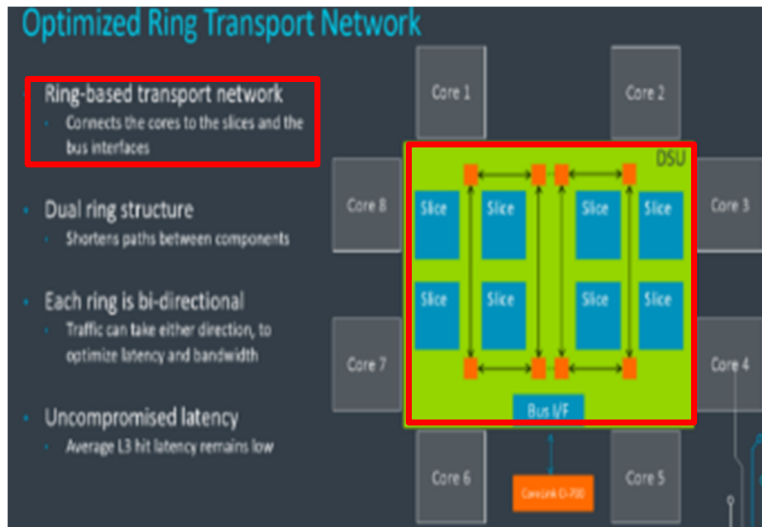


270. Further, in the MediaTek Dimensity 9000, the ARMv9 L3 cache comprises slices that are distributed across the SoC or processor die. The DSU-110 includes an interconnect bus, such as a ring-based transport network, which enables the processor cores to access the L3 cache slices⁹⁷:

Cache Slices

To be able to accommodate the new cores and future cores, Arm went with a large amount of cache. But at the same time, they are looking to considerably increase the bandwidth which means a lot more simultaneous cache accesses taking place in parallel. Their solution was to split it up into slices. Each slice makes up a portion of the L3 cache, includes part of the snoop filter along with the associated control logic. The actual number of cache slices is configurable on the DSU-110 – with up to eight slices supported (as the number of cores). Depending on the bandwidth requirements and the market, Arm partners can choose to configure it to best suit their needs. During cache access from a core, the target cache slice is chosen based on the hash of the address. It's, therefore,

⁹⁷ Schor, *supra* note 70; Arm DynamIQ Shared Unit-110 Technical Reference Manual at 27, 101, 102.



Internal Interconnects

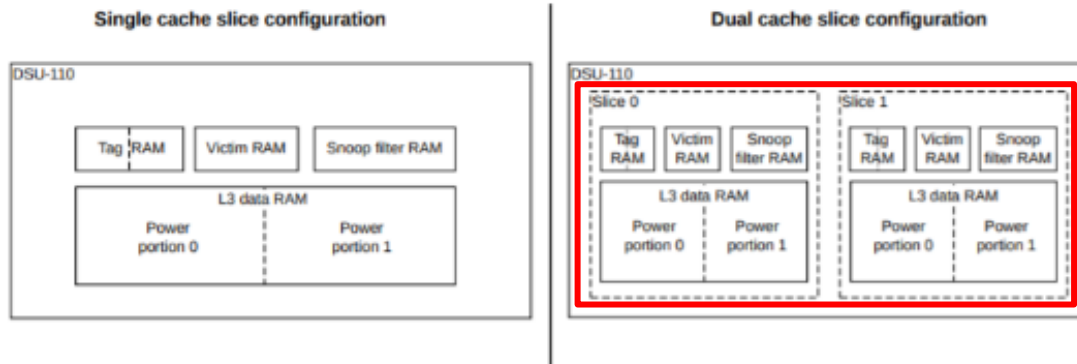
With the larger cache support and a high number of transactions, the interconnect implementation in the DSU was designed to allow data to flow quickly within the cluster. On the previous DSU, Arm used a hybrid crossbar implementation. Addressing this, the new DSU-110 uses a ring-based transport network to connect the cores to the slices and to the bus interface. But interestingly, Arm discovered that a single traditional ring was insufficient to meet the performance targets they wanted. For that

7.6 Cache slices and portions

The L3 cache of the *DynamiQ™ Shared Unit-110* (DSU-110) can be divided into up to eight identical slices, each containing between 256KB and 2MB of the cache. A cache slice consists of the data, tag, victim, and snoop filter RAMs and associated logic. A portion is a further subdivision of RAM in a cache slice.

For each cache slice, both the data RAM and tag RAM is subdivided into two portions.

The following figure shows the differences between a single and a dual cache slice configuration.

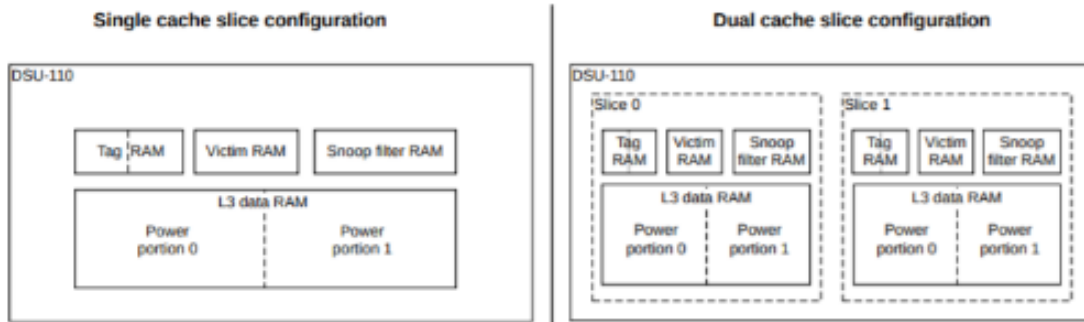
Figure 7-2: Comparison between a single and dual L3 cache slice configuration

Splitting the L3 cache into slices provides the following advantages:

- Improving the physical floorplan when implementing the macrocell, by ensuring that the RAMs are located close to the logic that is controlling them.
- Increasing the bandwidth because the slices can be accessed in parallel.

When a core type can be defined as part of a complex, then all instances of that core type (in the cluster) are implemented as complexes. This is either as part of a single-core complex or dual-core complex. Having all instances of a core type formed into complexes within the cluster, ensures consistent clock and power management control.

Within a dual-core complex, logic such as a Vector Processing Unit, L2 Translation Lookaside Buffer (TLB), and L3 cache logic is shared between the cores and is collectively known as Shared logic. In a single-core complex, the same logic resides outside the core but is collectively known as Dedicated logic.

Figure 7-2: Comparison between a single and dual L3 cache slice configuration

Splitting the L3 cache into slices provides the following advantages:

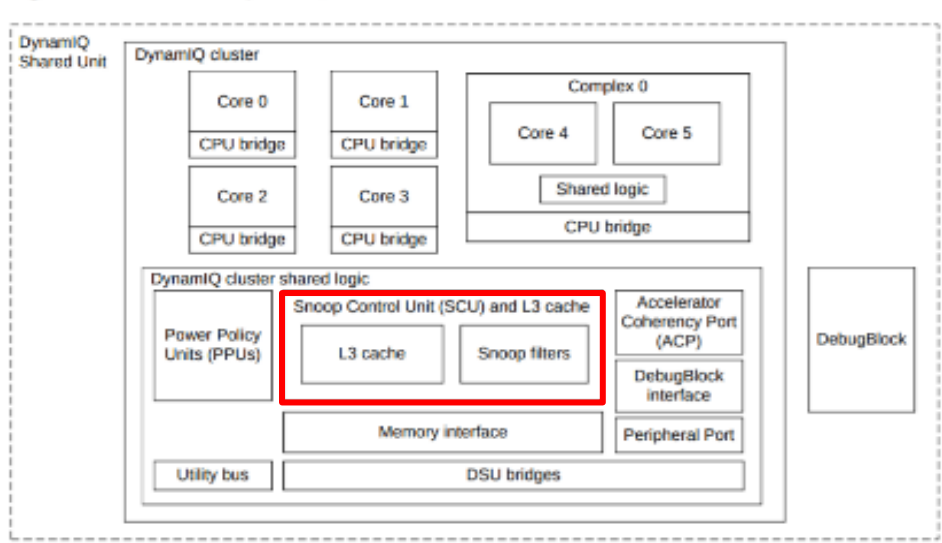
- Improving the physical floorplan when implementing the macrocell, by ensuring that the RAMs are located close to the logic that is controlling them.
- Increasing the bandwidth because the slices can be accessed in parallel.

271. Each of the '919 Patent Accused Products comprises cache management circuitry operative to provide coherent, non-uniform access to the plurality of distributed cache portions by the plurality of cores.

272. For example, the Dimensity 9000 SoCs include DynamIQ cluster shared logic that includes a Snoop Control Unit (SCU), which is the primary cache management module (e.g., the cache management circuitry) for the system. The SCU interfaces with at least the L3 cache interconnect network to manage the L3 cache and provide coherent access across all DynamIQ clusters⁹⁸:

⁹⁸ Arm DynamIQ Shared Unit-110 Technical Reference Manual at 27, 102; ARM® DynamIQ™ Shared Unit Technical Reference Manual.

Figure 2-1: DSU-110 DynamIQ™ cluster

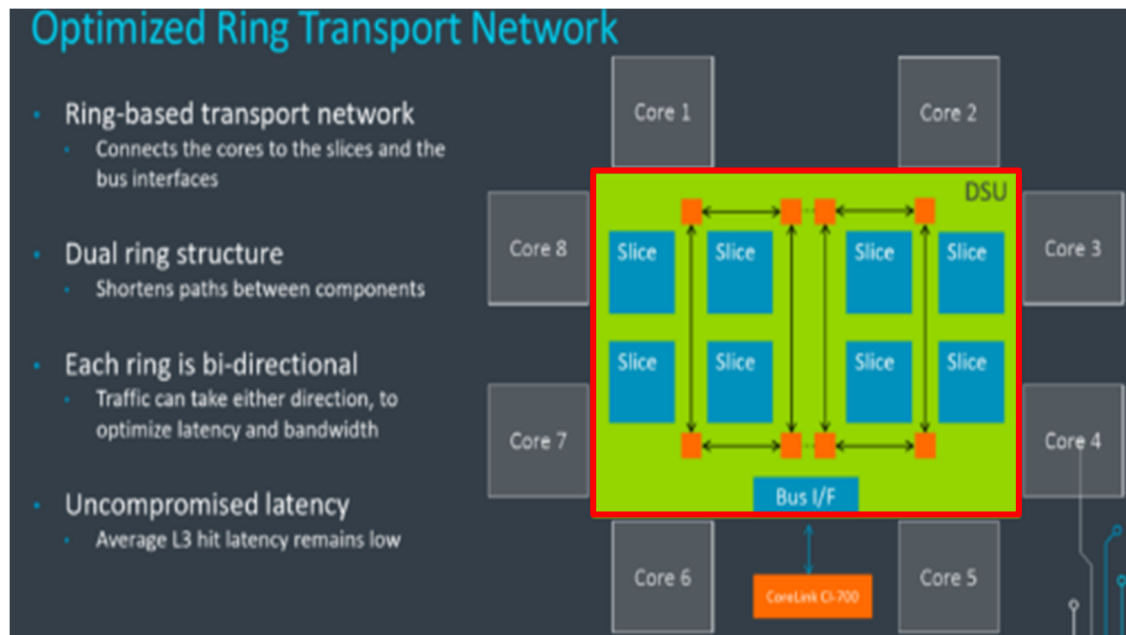


All cores in the DSU-110 DynamIQ™ cluster, including those in complexes, are coherently connected to an L3 memory system that includes an L3 cache and a *Snoop Control Unit* (SCU). The SCU maintains coherency between caches in the cores and the L3 cache, and includes a snoop filter to optimize coherency maintenance operations. The shared L3 cache simplifies process migration between the cores.

All the external interfaces including those to the cores are provided through the DSU-110 to the *System on Chip* (SoC). Main system transactions are supported through the memory interface which can be implemented as a coherent or non-coherent interface. A Peripheral port is provided to support low latency access to external system components but also can be used as a non-coherent master interface. The *Accelerator Coherency Port* (ACP) provides coherent access for non-cached masters that need I/O coherency with the cluster. The Utility bus is a memory-mapped port that provides a programming interface to the PPU's and some of the other system components.

273. Further, the '919 Accused Products comprise L3 cache slices that are physically distributed across the SoC die and are connected via an interconnect bus. Each cache slice is assigned to a group of cores. When a core faces a cache-miss in the associated cache slice portion, it fetches the cache line from a different core slice, if available, which increases the latency. Since the latency associated with access to cache lines in the L3 cache depends upon which cache slice

stores the cache line, the cache management circuitry provides non-uniform access of the cache slices to the cores⁹⁹:



The shared L3 cache of the DSU-110 provides the following functionality:

- A dynamically optimized cache allocation policy, which is typically exclusive. This cache allocation policy means that in normal use, a line is either in the cache of one or more cores (or complexes) or in the L3 cache, but not in both caches. Only Cacheable, shareable memory locations are allocated in the L3 cache. Non-shareable memory locations are not allocated in the L3 cache.
- Groups of cache ways can be partitioned and assigned to processes² by the *Memory System Resource Partitioning and Monitoring (MPAM)* architecture extension. Cache partitioning ensures that each process does not dominate the use of the cache to disadvantage other processes.

7.3 L3 cache partitioning

The L3 cache supports a partitioning scheme that alters the victim selection policy to prevent processes from using the entire L3 cache to the disadvantage of other processes.

Cache partitioning is intended for specialized software where there are distinct classes of processes running with different cache accessing patterns. For example, two processes A and B run on separate cores in the same cluster and therefore share the L3 cache. If process A is more data-intensive than process B, then process A can cause all the cache lines that process B allocates to be evicted. Evicting these allocated cache lines can reduce the performance of process B.

⁹⁹ Schor, *supra* note 70; Arm DynamIQ Shared Unit-110 Technical Reference Manual at 19, 35, 96, 97.

Snoop Control Unit (SCU)

The SCU maintains coherency between all the data caches in the cluster.

The SCU contains buffers that can handle direct cache-to-cache transfers between cores without having to read or write data to the L3 cache. Cache line migration enables dirty lines to be moved between cores. Also, there is no requirement to write back transferred cache line data to the L3 cache.

All cores in the DSU-110 DynamIQ™ cluster, including those in complexes, are coherently connected to an L3 memory system that includes an L3 cache and a Snoop Control Unit (SCU). The SCU maintains coherency between caches in the cores and the L3 cache, and includes a snoop filter to optimize coherency maintenance operations. The shared L3 cache simplifies process migration between the cores.

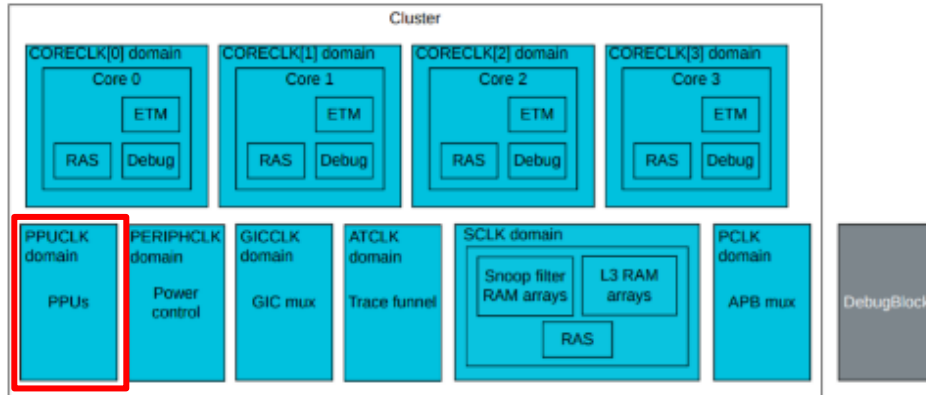
274. Each of the '919 Patent Accused Products comprises power management circuitry operative to enable a first frequency of operation for a first cluster of the plurality of cores.

275. For example, the Dimensity 9000's DSU-110 includes Power Policy Units (PPU, "power management circuitry") and a Power Control Module that provides DVFS control on per-core and per-cluster level. The Power Control Module also manages power consumption of the L3 distributed cache through power-gating of the cores as well as the L3 cache.¹⁰⁰

different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.

¹⁰⁰ Nayak, et al., *supra* note 2; Arm DynamIQ Shared Unit-110 Technical Reference Manual at 49, 51.

Figure 4-2: DSU-110 pin-controlled reset domains



The Power Policy Units (PPUs) for the cluster and each of the cores are used to control the power management features of the cluster and cores using a software interface. This includes managing various power states and transitions between these states. Certain power mode changes, for example powering up the cluster from a powered down state, include implicit resets to internal logic.

5.1 Power management in the DSU-110

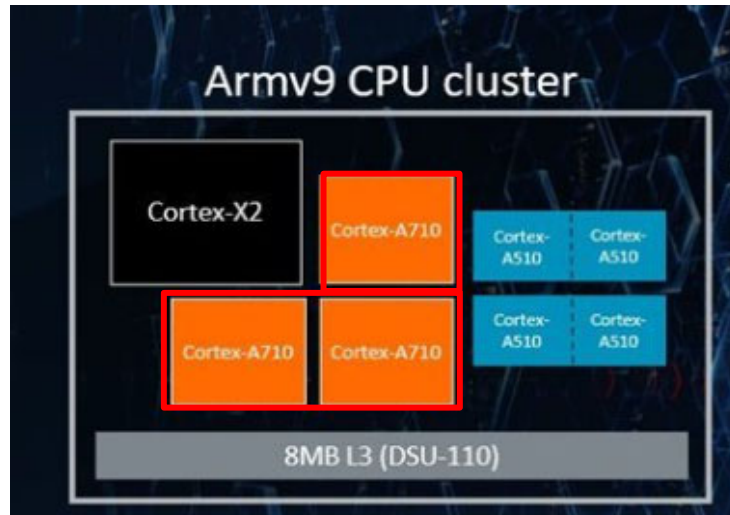
The DynamIQ™ Shared Unit-110 (DSU-110) provides various mechanisms to control both dynamic and static power dissipation. These mechanisms are associated with a set of power domains, power modes, and operational modes. Some of these mechanisms are brought under software control using Power Policy Units (PPUs).

The power management techniques employed by the DSU-110 and cores in the cluster include:

- Internal core clock gating where different internal parts of the core are clock idle
- Per-core Dynamic Voltage and Frequency Scaling (DVFS)
- Powerdown
- Retention, a low-power mode that retains the register and RAM state

276. Further, the '919 Patent Accused Products include processor clusters such as clusters of ARM Cortex-A510 cores (“first cluster”) and clusters of ARM Cortex-A710 cores (“second cluster”). Cores within each ARM cluster are located proximate to each other, as compared to distance from cores from a different cluster. Cores within a given cluster are typically

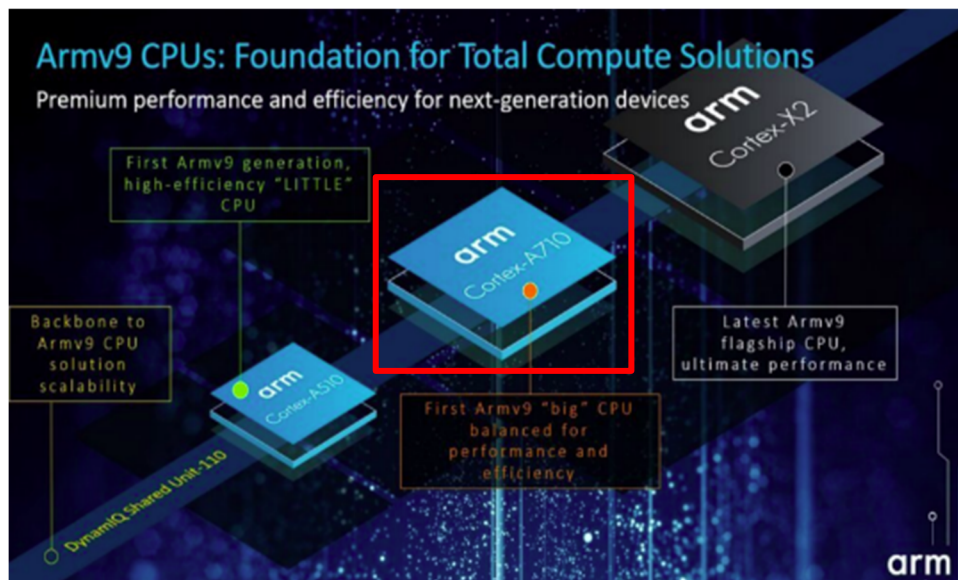
assigned to the same independent frequency (e.g., first frequency and second frequency) and voltage domains¹⁰¹:



Cluster features

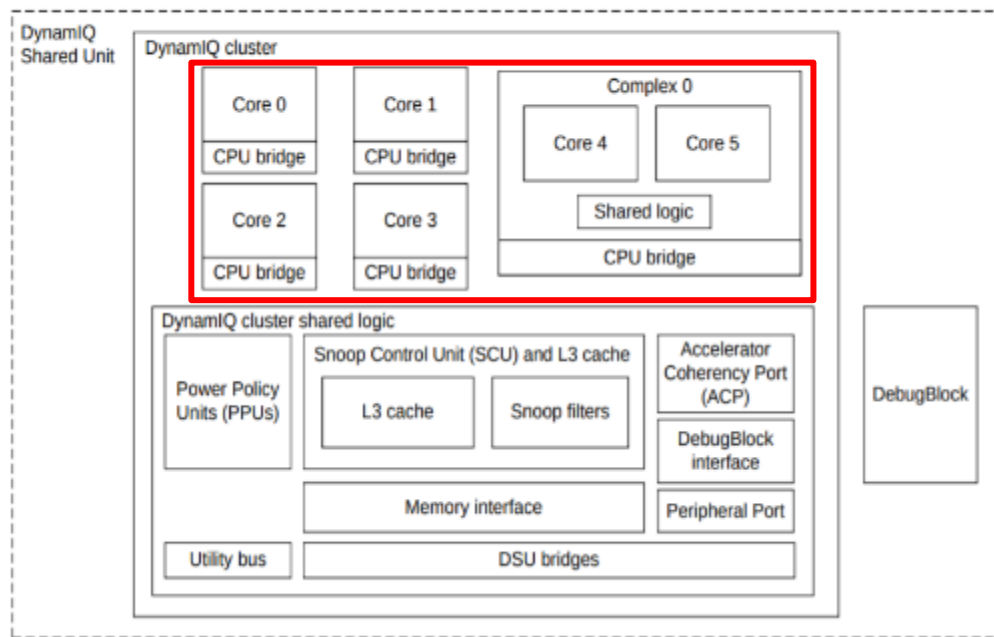
The DSU-110 has the following cluster features:

- Support for Arm[®]v9.0-A architecture cores
- Support for up to three types of core, and a maximum of eight cores in the cluster
- Power Policy Units (PPUs) providing autonomous power management of the L3 cache and the cores
- Support for cores running independently at different frequencies and voltages known as Dynamic Voltage Frequency Scaling (DVFS). For cores in a complex, DVFS is only possible for the whole complex not for individual cores.



¹⁰¹ Bedi, *supra* note 20; Arm DynamIQ Shared Unit-110 Technical Reference Manual at 18, 20.

Figure 2-1: DSU-110 DynamIQ™ cluster



A DSU-110 DynamIQ™ cluster consists of between one and eight cores, with up to three different types of cores in the same cluster. Cores can be configured for various performance points during macrocell implementation and run at different frequencies and voltages.

Cluster features

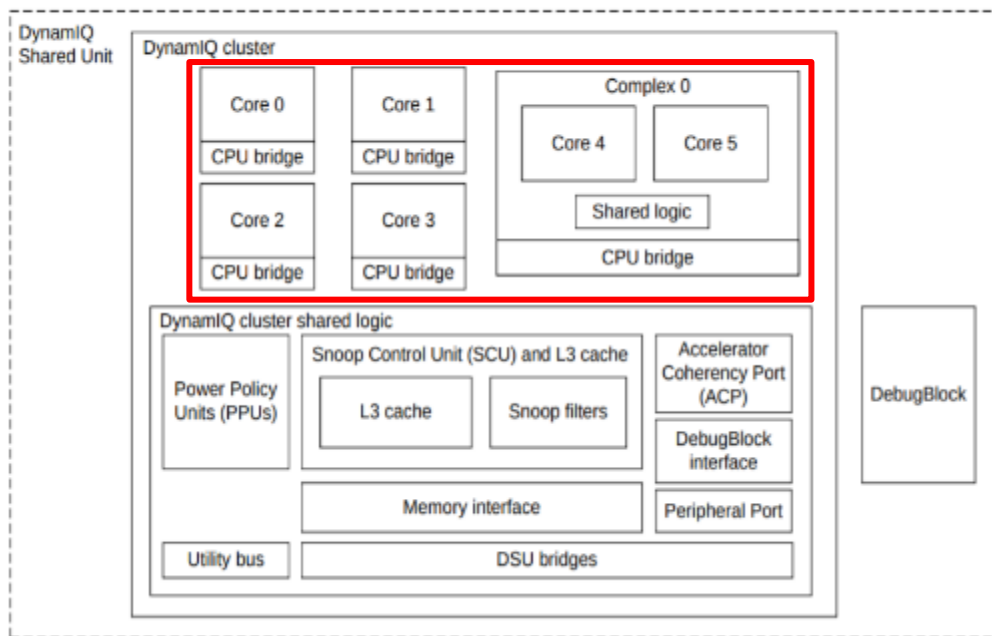
The DSU-110 has the following cluster features:

- Support for Arm®v9.0-A architecture cores
- Support for up to three types of core, and a maximum of eight cores in the cluster
- Power Policy Units (PPUs) providing autonomous power management of the L3 cache and the cores
- Support for cores running independently at different frequencies and voltages known as Dynamic Voltage Frequency Scaling (DVFS). For cores in a complex, DVFS is only possible for the whole complex not for individual cores.
- The DSU-110 has an internal transport mechanism that is responsible for all communication between components in the design. The topology of the transport is defined by the number of cores and number of L3 cache slices.

However, as with all of Arm's "LITTLE" CPUs, efficiency is still king. Not only does Cortex-A510 boost power efficiency by up to 20 percent (ISO process)¹ through the 3-wide in-order design, but it also provides industry-leading area efficiency. An innovation that makes this possible is merged core microarchitecture. This allows two Cortex-A510 CPUs to be grouped into a complex, with multiple complexes per CPU cluster.

277. Each of the '919 Patent Accused Products comprises cores that, within any given cluster, are closer to each other than they are to cores from other clusters. Accordingly, the average distance between cores in at least one of the clusters will be less than the average distance between the plurality of cores¹⁰²:

Figure 2-1: DSU-110 DynamIQ™ cluster



¹⁰² Arm DynamIQ Shared Unit-110 Technical Reference Manual at 18, 20; Bedi, *supra* note 20.

However, as with all of Arm's "LITTLE" CPUs, efficiency is still king. Not only does Cortex-A510 boost power efficiency by up to 20 percent (ISO process)¹ through the 3-wide in-order design, but it also provides industry-leading area efficiency. An innovation that makes this possible is merged core microarchitecture. This allows two Cortex-A510 CPUs to be grouped into a complex, with multiple complexes per CPU cluster.

278. Each of the '919 Patent Accused Products employs power management circuitry that is operative to selectively gate power to the first cluster of the plurality of cores and distributed cache portions of the at least one higher-level distributed cache that correspond to the first cluster and/or the second cluster of the plurality of cores and distributed cache portions of the at least one higher-level distributed cache that correspond to the second cluster.

279. For example, the Dimensity 9000's PPU ("power management circuitry") provides advanced power management features including selectively reducing power to individual CPU cluster cores as well as the L2 cache through DVFS. Additionally, individual L3 cache slices can also be partially powered down by the PPU¹⁰³.

different gears of the CPU subsystem enabling maximum power efficiency. Dynamic voltage and frequency scaling (DVFS) is employed along with adaptive voltage scaling to adjust operating voltage and frequency. Figure 2.5.1 demonstrates the power efficiency of the CPU subsystem achieving 27% improvement in single thread performance of the HP core over the BP core.

¹⁰³ Rosinger & Pradhan, *supra* note 14, pp. 8, 10; Nayak, et al., *supra* note 2; Arm DynamIQ Shared Unit-110 Technical Reference Manual at 49, 51, 58.






Figure 2.5.1: ARMv9 memory tagging, SVE2 performance chart, CPU power efficiency.

Arm DynamIQ Shared Unit-110

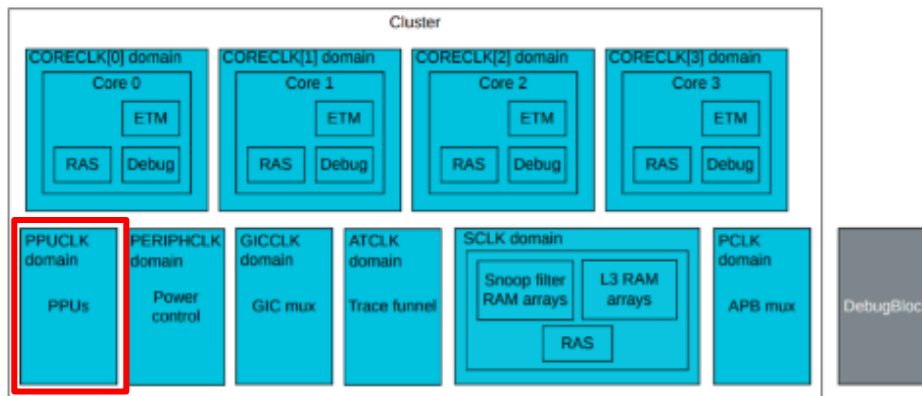
- DSU-110 for Cache Coherency & Shared L3\$
- Optimized Ring Transport Network
 - Bi-directional, Dual-ring
- 2X Bandwidth and 25% Lower Leakage
 - Support Partial SRAM & Logic Shutdown
- Cache Partition for QoS



ARMADA

11

Figure 4-2: DSU-110 pin-controlled reset domains



The *Power Policy Units (PPUs)* for the cluster and each of the cores are used to control the power management features of the cluster and cores using a software interface. This includes managing various power states and transitions between these states. Certain power mode changes, for example powering up the cluster from a powered down state, include implicit resets to internal logic.

5.1 Power management in the DSU-110

The *DynamiQ™ Shared Unit-110 (DSU-110)* provides various mechanisms to control both dynamic and static power dissipation. These mechanisms are associated with a set of power domains, power modes, and operational modes. Some of these mechanisms are brought under software control using *Power Policy Units (PPUs)*.

The power management techniques employed by the DSU-110 and cores in the cluster includes:

- Internal core clock gating where different internal parts of the core are clock idle
- Per-core *Dynamic Voltage and Frequency Scaling (DVFS)*
- Powerdown
- Retention, a low-power mode that retains the register and RAM state

5.4 L3 RAM power control

In addition to retention features, the *DynamiQ™ Shared Unit-110 (DSU-110)* can further reduce static leakage power using two powerdown features. Firstly, optionally power down of all but one of the L3 cache slices. Secondly, within each L3 cache slice powerdown a portion of the L3 cache RAM that the cache slice contains.

5.4.1 L3 cache RAM powerdown

The L3 cache RAMs typically contribute to a large proportion of the total leakage power, particularly for large cache sizes. Therefore, it is beneficial to be able to power down the RAMs when only some of the L3 cache is required, but it also results in reducing cache capacity. Parts of the L3 cache RAM can be independently powered down to reduce RAM leakage power when not in use. L3 cache powerdown is controlled by the cluster *Power Policy Unit (PPU)*.

The L3 cache RAM powerdown feature allows the RAMs to be powered down in groups of ways, giving options of 100%, 50%, or 0% of the L3 cache capacity. When a workload is making light use of the L3 cache, then this can be detected and the L3 cache capacity reduced without significant impact on the performance. For example, this can occur when the L3 cache has a relatively small memory footprint that mostly fits within the L2 cache.

280. Further, the Dimensity 9000's PPU controls chip-level power consumption by reducing/increasing/gating power delivered to clusters, L1 & L2 caches and L3 cache slices based on system requirements and operating conditions ("selectively gating power")¹⁰⁴:

Cluster features

The DSU-110 has the following cluster features:

- Support for Arm®v9.0-A architecture cores
- Support for up to three types of core, and a maximum of eight cores in the cluster
- **Power Policy Units (PPUs) providing autonomous power management of the L3 cache and the cores**

The DSU-110 DynamIQ™ cluster can be implemented with various power domains to target power performance levels. These power domains are managed through the *Power Policy Units (PPUs)*. The DSU-110 DynamIQ™ cluster supports many mechanisms to reduce static and dynamic power dissipation. For example, placing the cores and L3 cache into retention and powering down parts of the L3 cache.

6.6 Programming sequences for the cluster and the core

Example *Power Policy Unit (PPU)* programming sequences are provided for both the cluster and the cores. One of these sequences uses the static mode policy to demonstrate programming using this policy. However, because static power management can require considerable activity from the System Control Processor, Arm strongly recommends that you use dynamic power management for normal operation of the cluster.

6.6.1 Programming sequence to bring the cluster and cores from Off to On mode

Use the following steps, to program the *Power Policy Unit (PPU)* for the DSU-110 DynamIQ™ cluster and each of the cores to request a change of PPU mode from Off mode to On mode.

6.6.2 Programming sequence to bring the cluster and cores from On to Off mode

Use the following steps, to program the *Power Policy Unit (PPU)* for the DSU-110 DynamIQ™ cluster and each of the cores to request a change of PPU mode from On to Off.

¹⁰⁴ Arm DynamIQ Shared Unit-110 Technical Reference Manual at 20, 89.

281. Further, the L3 cache slices can be partially powered down based on the system workload¹⁰⁵:

Cache features

The DSU-110 has the following cache features:

- Optional unified 16-way set-associative L3 cache, configurable from 256KB to 16MB
- 64-byte cache lines
- L3 cache slice support, for improved bandwidth and cache RAM layout, up to eight slices supported
- L3 cache powerdown based either on cache slices or cache ways

5.4 L3 RAM power control

In addition to retention features, the *DynamIQ™ Shared Unit-110* (DSU-110) can further reduce static leakage power using two powerdown features. Firstly, optionally power down of all but one of the L3 cache slices. Secondly, within each L3 cache slice powerdown a portion of the L3 cache RAM that the cache slice contains.

5.4.1 L3 cache RAM powerdown

The L3 cache RAMs typically contribute to a large proportion of the total leakage power, particularly for large cache sizes. Therefore, it is beneficial to be able to power down the RAMs when only some of the L3 cache is required, but it also results in reducing cache capacity. Parts of the L3 cache RAM can be independently powered down to reduce RAM leakage power when not in use. L3 cache powerdown is controlled by the cluster *Power Policy Unit* (PPU).

5.4.2 L3 cache slice powerdown

In addition to powering down the L3 cache RAMs, you can gain further leakage savings by powering down some of the L3 cache slice control logic as well. Control of powering up or powering down L3 cache slices is performed by the cluster *Power Policy Unit* (PPU).

¹⁰⁵ Arm DynamIQ Shared Unit-110 Technical Reference Manual at 19, 20, 54, 58, 62.

5.3 Cluster power modes

The DSU-110 DynamiQ™ cluster and each of the cores and complexes in the cluster have a defined set of power modes and corresponding legal transactions between these modes.

The following table shows the supported power modes for the DSU-110 DynamiQ™ cluster.

Power mode	Short name	Description
On mode	ON	On mode is the normal mode of operation where all cluster functionality is available.
Off mode	OFF	In Off mode, power is removed from the cluster logic and all the RAMs.
Functional retention mode	FUNC_RET	In Functional retention mode, the L3 cache RAMs and snoop filter RAMs are placed in a retention state. Data is retained in these RAMs. The rest of the DynamiQ™ cluster shared logic remains powered up.

5.4.1 L3 cache RAM powerdown

The L3 cache RAMs typically contribute to a large proportion of the total leakage power, particularly for large cache sizes. Therefore, it is beneficial to be able to power down the RAMs when only some of the L3 cache is required, but it also results in reducing cache capacity. Parts of the L3 cache RAM can be independently powered down to reduce RAM leakage power when not in use. L3 cache powerdown is controlled by the cluster *Power Policy Unit* (PPU).

The L3 cache RAM powerdown feature allows the RAMs to be powered down in groups of ways, giving options of 100%, 50%, or 0% of the L3 cache capacity. When a workload is making light use of the L3 cache, then this can be detected and the L3 cache capacity reduced without significant impact on the performance. For example, this can occur when the L3 cache has a relatively small memory footprint that mostly fits within the L2 cache.

282. Further, on information and belief, MediaTek has actively induced and/or contributed to infringement of at least Claim 16 of the '919 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

283. Users of the '919 Patent Accused Products directly infringe at least Claim 16 of the '919 Patent when they use the '919 Patent Accused Products in the ordinary, customary, and intended way. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the '919 Patent Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the '919 Patent Accused Products to consumers within the United States and instructing and encouraging

such customers to use the '919 Patent Accused Products in the ordinary, customary, and intended way, which MediaTek knew infringes at least Claim 16 of the '919 Patent, or, alternatively, was willfully blind to the infringement.

284. On information and belief, MediaTek's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the '919 Patent Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the '919 Patent Accused Products in the United States, which MediaTek knew infringes at least Claim 16 of the '919 Patent, or, alternatively, was willfully blind to the infringement.

285. On information and belief, in violation of 35 U.S.C. § 271(c), MediaTek's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of at least Claim 16 of the '919 Patent, constituting a material part of the invention. On information and belief, MediaTek knows and has known the same to be especially made or especially adapted for use in an infringement of the '919 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

286. On information and belief, in violation of 35 U.S.C. § 271(f)(1), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 16 of the '919 Patent, where such components are uncombined in whole or in part, in such

manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

287. On information and belief, in violation of 35 U.S.C. § 271(f)(2), MediaTek's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 16 of the '919 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

288. MediaTek is not licensed or otherwise authorized to practice the claims of the '919 Patent.

289. Thus, by its acts, MediaTek has injured Daedalus and is liable to Daedalus for directly and/or indirectly infringing one or more claims of the '919 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 16.

290. On information and belief, MediaTek has known about the '919 Patent at least since August 23, 2022.¹⁰⁶ At a minimum, MediaTek has knowledge of the '919 Patent at least as of the filing of this Complaint. Accordingly, MediaTek's infringement of the '919 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

¹⁰⁶ *Daedalus Prime LLC v. Mazda Motor Corp., et al.*, No. 22-cv-01108 (D. Del. Aug. 23, 2022).

291. As a result of MediaTek's infringement of the '960 Patent, Daedalus has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for MediaTek's infringement, but in no event less than a reasonable royalty with interest and costs.

292. On information and belief, MediaTek will continue to infringe the '919 Patent unless enjoined by this Court. MediaTek's infringement of Daedalus' rights under the '919 Patent will continue to damage Daedalus, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

DEMAND FOR JURY TRIAL

Pursuant to Rule 38(b) of the Federal Rules of Civil Procedure, Plaintiff demands a trial by jury in this action for all issues triable by a jury.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff prays for judgment and seeks relief from MediaTek as follows:

293. For judgment that MediaTek has infringed and continues to infringe the claims of the '316, '197, '281, '228, '167, '838, '960 and '919 Patents;

294. For a permanent injunction against MediaTek and its respective officers, directors, agents, servants, affiliates, employees, divisions, branches, subsidiaries, parents, and all other acting in active concert therewith from infringement of the '316, '197, '281, '228, '167, '838, '960 and '919 Patents;

295. For an accounting of all damages sustained by Plaintiff as a result of MediaTek's acts of infringement;

296. For a mandatory future royalty payable on each and every future sale by MediaTek of a product that is found to infringe one or more of the Asserted Patents and on all future products which are not more than colorably different from products found to infringe;

297. For a judgment and order finding that MediaTek's infringement is willful and awarding to Plaintiff enhanced damages pursuant to 35 U.S.C. § 284;

298. For a judgment and order requiring MediaTek to pay Plaintiff's damages, costs, expenses, and pre- and post-judgment interest for its infringement of the '316, '197, '281, '228, '167, '838, '960 and '919 Patents as provided under 35 U.S.C. § 284 and without limitation under 35 U.S.C. § 287;

299. For a judgment and order finding that this is an exceptional case within the meaning of 35 U.S.C. § 285 and awarding to Plaintiff its reasonable attorneys' fees; and

300. For such other and further relief in law and in equity as the Court may deem just and proper.

Dated: April 8, 2024

Respectfully Submitted,

/s/ Garland Stephens, with permission
Charles Everingham IV

Garland Stephens

LEAD ATTORNEY

Texas Bar No. 24053910

garland@bluepeak.law

John Brinkmann

Texas Bar No. 24068091

john@bluepeak.law

Richard Koehl

Texas Bar No. 24115754

richard@bluepeak.law

Robert Magee

California Bar No. 271443 (to be admitted
pro hac vice)

robert@bluepeak.law

Heng Gong

New York Bar No. 4930509 (to be admitted
pro hac vice)

heng@bluepeak.law

BLUE PEAK LAW GROUP LLP
3139 West Holcombe Blvd, PMB 8160
Houston, TX 77025
Telephone: 281-972-3036

Of Counsel:

WARD SMITH & HILL, PLLC
Claire Abernathy Henry
Texas State Bar No. 24053063
claire@wsfirm.com
Charles Everingham IV
Texas Bar No. 00787447
ce@wsfirm.com
Andrea L. Fair
Texas State Bar No. 24078488
E-mail: andrea@wsfirm.com
1507 Bill Owens Pkwy
Longview, Texas 75604
Phone: (903) 757-6400
Fax: (903) 757-2323

Attorneys for Plaintiff Daedalus Prime LLC