

PCT

- (51) **International Patent Classification:**
G09G 5/00 (2006.01)

(21) **International Application Number:**
PCT/US2009/003316

(22) **International Filing Date:**
1 June 2009 (01.06.2009)

(25) **Filing Language:** English

(26) **Publication Language:** English

(30) **Priority Data:**
61/057,483 30 May 2008 (30.05.2008) US

(71) **Applicant** (*for all designated States except US*): **ADVANCED MICRO DEVICES, INC.** [US/US]; One AMD Place, Sunnyvale, CA 94088 (US).

(72) **Inventors; and**

(75) **Inventors/Applicants** (*for US only*): **MANTOR, Michael, J.** [US/US]; 1620 Pinar Drive, Orlando, FL 32825 (US). **BRADY, Jeffrey, T.** [US/US]; 14151 Tanja King Boulevard, Orlando, FL 32828 (US). **FOWLER, Mark, C.** [US/US]; 19 Nicholas Road, Hopkinton, MA 01748 (US). **ZINI, Marcos, P.** [US/US]; 2836 Aloma Lake Run, Oviedo, FL 32765 (US).

(74) **Agents:** **WOOD, Theodore, A.** @ et al.; Sterne, Kessler, Goldstein & Fox P.L.L.C., 1100 New York Avenue, N.W., Washington, DC 20005-3934 (US).

(81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report (Art. 21(3))*

Published:

— with international search report (Art. 21(3))

(54) Title: SCALABLE AND UNIFIED COMPUTE SYSTEM




FIG. 1

(S7) Abstract: A Scalable and Unified Compute System performs scalable, repairable general purpose and graphics shading operations, memory load/store operations and texture filtering. A Scalable and Unified Compute Unit Module comprises a shader pipe array, a texture mapping unit, and a level one texture cache system. The Scalable and Unified Compute Unit Module accepts ALU instructions, input/output instructions, and texture or memory requests for a specified set of pixels, vertices, primitives, surfaces, or general compute work items from a shader program and performs associated operations to compute the programmed output data. The texture mapping unit accepts source data addresses and instruction constants in order to fetch, format, and perform instructed filtering interpolations to generate formatted results based on the specific corresponding data stored in a level one texture cache system. The texture mapping unit consists of an address generating system, a pre-formatter module, interpolator module, accumulator module and a format module. A method for a Scalable and Unified Compute System is also presented.

SCALABLE AND UNIFIED COMPUTE SYSTEM

BACKGROUND

Field of the Invention

[0001] The present invention generally relates to computing operations performed by computing systems, and more particularly to graphics processing tasks performed by computing systems.

Related Art

[0002] A graphics processing unit (GPU) is a complex integrated circuit that is specially configured to carry out graphics processing tasks. A GPU can, for example, execute graphics processing tasks required by an end-user application, such as a video game application. For such applications, for example, there are several layers of software between the end-user application and the GPU.

[0003] The end-user application communicates with an application programming interface (API). An API allows the end-user application to output graphics data and commands in a standardized format, rather than in a format that is dependent on the GPU. Several types of APIs are commercially available, including DirectX® developed by Microsoft Corp. and OpenGL® developed by Silicon Graphics, Inc. The API communicates with a driver. The driver translates standard code received from the API into a native format of instructions understood by the GPU. The driver is typically written by the manufacturer of the GPU. The GPU then executes instructions from the driver.

[0004] A GPU produces, by carrying out a process known as "rendering" creates individual pixels that together form an image based on a higher level description of image components. A GPU typically carries out continuous rendering using pipelines to process pixel, texture, and geometric data. These pipelines are often referred to as a collection of fixed function special purpose pipelines such as rasterizers, setup engines, color blenders, hierarchical depth, texture mapping and programmable stages that can be accomplished in shader pipes or shader pipelines, "shader" being a term in computer

graphics referring to a set of software instructions used by a graphic resource primarily to perform rendering effects. In addition, GPU's can also employ multiple programmable pipelines in a parallel processing design to obtain higher throughput. A multiple of shader pipelines can also be referred to as a shader pipe array.

[0005] As a shader pipe array performs its ongoing rendering process, manufacturing defects in a pipeline can be become known and and subsequent failures can occur within a pipeline. A small defect or failure in a system without any logic repair is fatal and can render the device defective.

[0006] In addition, GPUs also support texture mapping. Texture mapping is a process used to determine the texture color for a texture mapped pixel through the use of the colors of nearby pixels of the texture, or texels. The process is also referred to as texture smoothing or texture interpolation. However, high image quality texture mapping requires a high degree of computational complexity.

[0007] Given the ever increasing complexity of new software applications and API shader language advancements, the demands on GPUs to provide high quality rendering, texture mapping and generalize memory access for advanced rendering and generalized compute, computation complexities are further increasing.

[0008] In addition, GPUs equipped with a Unified Shader also simultaneously support many types of shader processing, from pixel, vertex, primitive, surface and generalized compute are raising the demand for higher performance generalized memory access capabilities.

[0009] What are needed, therefore, are systems and/or methods to alleviate the aforementioned deficiencies. Particularly, what is needed is a scalable Unified Shader design approach for a shader pipe with a multipurpose texture mapping system capable of generalized load/store operations with the ability to overcome the effects of a defective internal subcircuit with minimal impact on overall system performance.

SUMMARY OF THE INVENTION

[0010] This section is for the purpose of summarizing some aspects of the present invention and to briefly introduce some preferred embodiments.

Simplifications or omissions may be made to avoid obscuring the purpose of the section. Such simplifications or omissions are not intended to limit the scope of the present invention.

[0011] Consistent with the principles of the present invention as embodied and broadly described herein, the present invention includes method and apparatus related to a row based Scalable and Unified Compute Unit Module. The Scalable and Unified Compute Unit Module includes a shader pipe array and texture mapping unit with a level one cache system to perform texture mapping and general load/store accesses with the ability to process shader pipe data destined to a defective shader pipe. The Scalable and Unified Compute System comprises a sequencer, Scalable and Unified Compute Unit Module with access to a level two texture cache system and thus an external memory system. The Scalable and Unified Compute System is configured to accept an executing shader program instruction including input, output, ALU and texture or general memory load/store requests with address data from the shader pipes and program constants to generate the return texel or memory data based on state data controlling the pipelined address and filtering operations for a specific pixel or thread. The texture filter system is configured based on the shader program instruction and constant to generate a formatted interpolation based on texel data stored in the cache system for the addresses stored in the shader pipeline.

[0012] In an embodiment of the invention, the row based shader pipe Scalable and Unified Compute System further comprises a redundant shader pipe system. The redundant shader pipe system is configured to process shader pipe data destined to a defective shader pipes in the shader pipe array.

[0013] In another embodiment the row based Scalable and Unified Compute System further comprises a level two texture cache system. The level two texture cache system can be read and written to by any level one row based texture cache system.

[0014] In another embodiment the texture filter in the texture mapping unit in the Scalable and Unified Compute Unit Module further comprises a pre-formatter module, an interpolator module, an accumulator module, and a format module. The pre-formatter module is configured to receive texel data and convert it to a normalized fixed point format. The interpolator module is

configured to perform an interpolation on the normalized fixed point texel data from the pre-formatter module and generate re-normalized floating point texel data. The accumulator module is configured to accumulate floating point texel data from the interpolator module to achieve the desired level of bilinear, trilinear, and anisotropic filtering. The format module is configured to convert texel data from the accumulator module into a standard floating point representation.

[0015] Further features and advantages of the invention, as well as the structure and operation of various embodiments of the present invention, are described in detail below with reference to the accompanying drawings. It is noted that the invention is not limited to the specific embodiments described herein. Such embodiments are presented herein for illustrative purposes only. Additional embodiments will be apparent to persons skilled in the relevant art(s) based on the teachings contained herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] The accompanying drawings, which are incorporated in and constitute part of the specification, illustrate embodiments of the invention and, together with the general description given above and the detailed description of the embodiment given below, serve to explain the principles of the present invention. In the drawings:

[0017] FIG. 1 is a system diagram depicting an implementation of a Scalable and Unified Compute System.

[0018] FIG. 2 is a system diagram depicting an implementation of a Scalable and Unified Compute System illustrating the details of the shader pipe array.

[0019] FIG. 3 is a system diagram depicting an implementation of a Scalable and Unified Compute System illustrating the details of the texture mapping unit.

[0020] FIG. 4 is a flowchart depicting an implementation of a method for a Scalable and Unified Compute System.

[0021] The features and advantages of the present invention will become more apparent from the detailed description set forth below when taken in

conjunction with the drawings, in which like reference characters identify corresponding elements throughout. In the drawings, like reference numbers generally indicate identical, functionally similar, and/or structurally similar elements. The drawing in which an element first appears is indicated by the leftmost digit(s) in the corresponding reference number.

DETAILED DESCRIPTION

[0022] The invention will be better understood from the following descriptions of various “embodiments” of the invention. Thus, specific “embodiments” are views of the invention, but each is not the whole invention. In one respect, the invention relates to a Scalable and Unified Compute System whereby a shader pipe array processes shader program instructions on input pixel, vertex, and primitive, surface or compute work items to create output data for each item using generated texel data or memory load/store operations. In embodiments of this invention, bilinear texture mapping, trilinear texture mapping, and anisotropic texture mapping are applied to the texel data that is stored in a multi-level cache system. In another embodiment, a redundant shader system can be added and configured to process shader pipe data directed to defective shader pipes within the shader pipe array to recover devices with a defective sub-circuit in one or more shader pipes.

[0023] Embodiments of this invention that have configurations containing two or more Scalable and Unified Compute Systems, a subset of the Unified Compute Unit System itself can be configured to be a repairable unit. In such an embodiment workloads destined to a defective Unified Compute Unit System will instead be sent to a redundant Unified Compute Unit System to process all ALU, texture, and memory operations. This increases the portion of the device that is covered by repair significantly due to the inclusion of texture mapping unit and L1 cache system and thus significantly improves on the yield of such a device.

[0024] While specific configurations, arrangements, and steps are discussed, it should be understood that this is done for illustrative purposes only. A person skilled in the pertinent art(s) will recognize that other configurations, arrangements, and steps can be used without departing from the spirit and

scope of the present invention. It will be apparent to a person skilled in the pertinent art(s) that this invention can also be employed in a variety of other applications.

[0025] It is noted that references in the specification to “one embodiment,” “an embodiment,” “an example embodiment,” etc., indicate that the embodiment described can include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it would be within the knowledge of one skilled in the art to incorporate such a feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

[0026] While the present invention is described herein with reference to illustrative embodiments for particular applications, it should be understood that the invention is not limited thereto. Those skilled in the art with access to the teachings provided herein will recognize additional modifications, applications, and embodiments within the scope thereof and additional fields in which the invention would be of significant utility.

[0027] FIG. 1 is an illustration of a Scalable and Unified Compute System 100 according to an embodiment of the present invention. System 100 comprises a sequencer 110, a Scalable and Unified Compute Unit Module 120, and a level two cache system 130. Scalable and Unified Compute Unit Module 120 comprises a shader pipe array 122, optional redundant shader pipe array 124, texture mapping unit 126, and level one texture cache system 128.

[0028] Shader pipe array 122 performs ALU operations on input data. Sequencer 110 controls the shader program instruction issue for contained workloads and the flow of data through shader pipe array 122. In addition, in an embodiment where the redundant shader pipe array 124 is present, sequencer 110 reacts to defective shader pipes that occur within shader pipe array 122 by scheduling instructions to the appropriate redundant units.

[0029] Sequencer 110 can issue a texture fetch or load/store operation that will initiate shader pipe array 122 to send addresses with the instruction issued to texture mapping unit 126. In this instance texture mapping unit 126

generates appropriate addresses to the level one texture cache system 128 that contains texel data or memory data associated with the addresses. Level one cache system 128, after receiving the addresses, will return the associated texel or memory data to texture mapping unit 126. When the requested texel or memory data is not present in the level one cache system 128, then the request is forwarded to a level two cache system 130 obtain and return the requested texel data.

[0030] FIG. 2 is an illustration of a Scalable and Unified Compute Unit Module 120 including a more detailed view of the shader pipe array 122 according to an embodiment of the present invention. In this embodiment, shader pipe array 122 comprises one or more shader pipe blocks, here represented as SP_0 through SP_M, where "M" represents a positive integer greater than one.

[0031] In an embodiment where redundant shader pipe array 124 is present, if sequencer 110 identifies, as an example, that the shader pipe located in shader pipe block SP_1 is defective, then the shader pipe data destined to the defective pipe would be sent to redundant shader pipe array 124 via the input stream by the input module and processed by redundant shader pipe array 124. All texture mapping requests would be intercepted by redundant shader pipe array 124 when instructed via horizontal control path 211 from sequencer 110. Once redundant shader pipe array 124 processes the shader pipe data initially destined to the defective shader pipe, the processed redundant shader pipe array 124 data would be transferred from redundant shader pipe array 124 back to the output stream of shader pipe 122 and realigned in an output unit (not shown).

[0032] In an embodiment, redundant shader pipe array 124 consists of a single block, and therefore can only process shader pipe data destined to a single defective shader pipe at a time. In another embodiment wherein redundant shader pipe array 124 comprises multiple redundant shader blocks, then redundant shader pipe array 124 can process shader pipe data destined to more than one defective shader pipe simultaneously.

[0033] FIG. 3 illustrates a more detailed view of texture mapping unit 126 according to an embodiment of the present invention. In this embodiment, shader pipe array 122 generates a texture or memory load/store request to

texture mapping unit 126 that comprises an address generator system 318, a pre-formatter module 310, an interpolator module 312, an accumulator module 314, and a format module 316. The texture mapping unit 126 receives a request from shader arrays 122 and 124 and sequencer 110 respectively, and processes the instruction in address generator system 318 to determine the actual addresses to service. Resulting texel data received from the level one texture cache system 128 and, after the data is processed in pre-formatter module 310, interpolator module 312, accumulator module 314, and format module 316. The resultant texel data is sent back to the requesting resource in shader pipe array 122 and/or redundant shader pipe array 124. Pre-formatter module 310 is configured to receive texel data and perform a block normalization thereby generating normalized fixed point texel data. Interpolator module 312 receives the normalized fixed point texel data from pre-formatter module 310 and performs one or more interpolations, each of that are accumulated in accumulator module 314, to achieve the desired level of bilinear, trilinear, and anisotropic texture mapping. Format module 316 converts the accumulated texel data in accumulator module 314 to a standard floating point representation for the requesting resource, shader pipe array 122. For general load/store data pre-formatter module 310, interpolator module 312, accumulator module 314, and format module 316 pass the requested return data unmodified.

[0034] An embodiment represented by FIG. 3 also illustrates the use of a level two cache system 130. The level two cache system is additional memory that is available to Scalable and Unified Computer Unit Module 120 when it is necessary or desirable to read and/or write data from and to the level one cache system 128.

[0035] FIG. 4 is a flowchart depicting a method 400 for texture mapping using a Scalable and Unified Compute System. Method 400 begins at step 402. In step 404, a shader pipe receives set texture requests from a sequencer for a specific set of pixels, vertices, primitives, surfaces, or computer work items. In step 406 the shader pipe generates data set addresses based on shader program instructions for the specified set of pixels, vertices, primitives, surfaces, or compute work items. In step 408, a texture mapping unit retrieved stored texel data from a level one and/or level two texture cache system.

Moreover, in step 410 the texture mapping unit calculates a formatted accumulated interpolation based on the retrieved texel data and the originating shader instruction. Method 400 ends at step 412.

[0036] The functions, processes, systems, and methods outlined in FIGs. 1, 2, 3, and 4 can be implemented in software, firmware, or hardware, or using any combination thereof. If programmable logic is used, such logic can execute on a commercially available processing platform or a special purpose device.

[0037] As would be apparent to one skilled in the relevant art, based on the description herein, embodiments of the present invention can be designed in software using a hardware description language (HDL) such as, for example, Verilog or VHDL. The HDL-design can model the behavior of an electronic system, where the design can be synthesized and ultimately fabricated into a hardware device. In addition, the HDL-design can be stored in a computer product and loaded into a computer system prior to hardware manufacture.

[0038] It is to be appreciated that the Detailed Description section, and not the Summary and Abstract sections, is intended to be used to interpret the claims. The Summary and Abstract sections may set forth one or more but not all exemplary embodiments of the present invention as contemplated by the inventor(s), and thus, are not intended to limit the present invention and the appended claims in any way.

[0039] The present invention has been described above with the aid of functional building blocks illustrating the implementation of specified functions and relationships thereof. The boundaries of these functional building blocks have been arbitrarily defined herein for the convenience of the description. Alternate boundaries can be defined so long as the specified functions and relationships thereof are appropriately performed.

[0040] The foregoing description of the specific embodiments will so fully reveal the general nature of the invention that others can, by applying knowledge within the skill of the art, readily modify and/or adapt for various applications such specific embodiments, without undue experimentation, without departing from the general concept of the present invention. Therefore, such adaptations and modifications are intended to be within the meaning and range of equivalents of the disclosed embodiments, based on the teaching and guidance presented herein. It is to be understood that the

phraseology or terminology herein is for the purpose of description and not of limitation, such that the terminology or phraseology of the present specification is to be interpreted by the skilled artisan in light of the teachings and guidance.

[0041] While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example only, and not limitation. It will be apparent to persons skilled in the relevant art that various changes in form and detail can be made therein without departing from the spirit and scope of the invention. Thus, the breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

WHAT IS CLAIMED IS:

1. A scalable and unified compute module, comprising:
 - a shader pipe array configured to accept texture instructions and generate output data;
 - a level one texture cache system configured to store the texture output data; and
 - a texture mapping unit configured to receive texture output data from the level one texture cache system and generate formatted interpolation data.
2. The scalable and unified compute module of claim 1, wherein the shader pipe array is configured to accept ALU, load/store, and output instructions
3. The scalable and unified compute module of claim 1, further comprising a redundant shader pipe array configured to process shader pipe data destined to a defective shader pipe of said shader pipe array.
4. The scalable and unified compute module of claim 1, wherein the shader pipe array comprises one or more shader pipe blocks.
5. The scalable and unified compute module of claim 4, wherein a shader pipe block is configured to contain one or more shader pipes.
6. The scalable and unified compute module of claim 1, wherein the level one texture cache system is configured to read and write to a level two cache system.

7. The scalable and unified compute module of claim 1, wherein the texture mapping unit comprises:
 - a pre-formatter module configured to accept texel data and generate normalized fixed point texel data;
 - a interpolator module configured to perform an interpolation on the normalized fixed point texel data from the pre-formatter module and generate re-normalized floating point texel data;
 - an accumulator module configured to accumulate floating point texel data from the interpolator module; and
 - a format module configured to convert texel data from the accumulator module into a standard floating point representation.
8. The scalable and unified compute module of claim 7, wherein the interpolator module is configured to perform one or more interpolations in order to achieve at least one of:
 - a bilinear texture filtering;
 - a trilinear texture filtering; and
 - an anisotropic texture filtering.
9. A method for scalable and unified computing, comprising:
 - receiving a set of texture requests for a specified set of pixels, vertices, primitives, surfaces, or compute work items;
 - generating a data set of addresses based on a shader program instruction for the specified set of pixels, vertices, primitives, surfaces, or compute work items;
 - retrieving stored texel data from a level one cache system; and
 - calculating a formatted accumulated interpolation based on the retrieved texel data and an originating shader instruction.
10. The scalable and unified computing method of claim 9, further comprising:
 - processing shader pipe data destined to a defective shader pipe.

11. The scalable and unified computing method of claim 9, further comprising:

reading from and writing to a level two cache system from the level one texture cache system.

12. The scalable and unified computing method of claim 9, further comprising:

receiving floating point texel data;
generating normalized fixed point texel data from the floating point texel data;
performing an interpolation on the normalized fixed point texel data;
generating re-normalized floating point texel data;
accumulating re-normalized texel data; and
formatting the accumulated re-normalized texel data into a standard floating point representation.

13. The scalable and unified computing method of claim 12, wherein interpolation further comprises:

bilinear texture filtering;
trilinear texture filtering; and
anisotropic texture filtering.

14. The scalable and unified computing method of claim 12, wherein the method is performed by synthesizing hardware description language instructions.

15. A system for scalable and unified computing, comprising:

a processor; and

a memory, in communication with said processor, configured to store a plurality of processing instructions for directing said computing system to:

receive a set of texture requests for a specified set of pixels, vertices, primitives, surfaces, or compute work items;

generate a data set of addresses based on a shader program instruction for the specified set of pixels, vertices, primitives, surfaces, or compute work items;

retrieve stored texel data from a level one cache system; and

calculate a formatted accumulated interpolation based on the retrieved texel data and an originating shader instruction.

16. The scalable and unified computing system of claim 15, further comprising instructions for causing the computing system to:

process shader pipe data destined to a defective shader pipe.

17. The scalable and unified computing system of claim 15, further comprising instructions for causing the computing system to:

read and write to a level two cache system from the level one texture cache system.

18. The scalable and unified computing system of claim 15, further comprising instructions for causing the computing system to:

receive floating point texel data;

generate normalized fixed point texel data from the floating point texel data;

perform an interpolation on the normalized fixed point texel data;

generate re-normalized floating point texel data;

accumulate re-normalized texel data; and

format the accumulated re-normalized texel data into a standard floating point representation.

19. The scalable and unified computing system of claim 15, further comprising instructions for causing said computing system to:

- filter with a bilinear texture filter;
- filter with a trilinear texture filter; and
- filter with an anisotropic texture filter.

20. A system for scalable and unified computing, comprising:

- means for receiving a texture request from a resource for a specified pixel;
- means for generating output texel data based on rendering calculations for the specified pixel;
- means for retrieving texel data from a level one cache system; and
- means for calculating, based on the texel data, a formatted interpolation utilizing a mapping unit.

21. The scalable and unified computing system of claim 20, further comprising:

- means for processing shader pipe data destined to a defective shader pipe.

22. The scalable and unified computing system of claim 20, further comprising:

- means for reading from and writing to a level two cache system from the level one texture cache system.

23. The scalable and unified computing system of claim 20, further comprising:

- means for receiving floating point texel data;
- means for generating normalized fixed point texel data from the floating point texel data;
- means for performing an interpolation on the normalized fixed point texel data;
- means for generating re-normalized floating point texel data;
- means for accumulating re-normalized texel data; and
- means for formatting accumulated re-normalized texel data into a standard floating point representation.

24. The scalable and unified computing system of claim 20, further comprising:

- means for filtering using a bilinear texture filter;
- means for filtering using a trilinear texture filter; and
- means for filtering using an anisotropic texture filter.

25. A computer readable medium carrying one or more sequences of one or more instructions which when executed by one or more processor-based computing devices cause the one or more computing devices to perform a method of scalable and unified computing, comprising:

- receiving a set of texture requests for a specified set of pixels, vertices, primitives, surfaces, or compute work items;
- generating a data set of addresses based on a shader program instruction for the specified set of pixels, vertices, primitives, surfaces, or compute work items;
- retrieving stored texel data from a level one cache system; and
- calculating a formatted accumulated interpolation based on the retrieved texel data and an originating shader instruction.




FIG. 1

200




FIG. 2

300




FIG. 3




FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 09/03316

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G09G 5/00 (2009.01)

USPC - 345/581

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

USPC: 345/581

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

USPC: 345/505; 345/552

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

USPTO WEST (PGPB, USPT, EPAB, JPAB); Google Scholar

Search Terms Used: texel, texture, element, pixel, shader, pipe, array, cache, anisotropic, bilinear, trilinear, filter,

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2008/0094405 A1 (BASTOS et al.) 24 April 2008 (24.04.2008), entire document, especially para [0010], [0024], [0029]-[0033], [0046]-[0048]	1-25
Y	US 6,104,415 A (GOSSETT) 15 August 2000 (15.08.2000), entire document, especially col 2, ln 28-54; col 7, ln 6-67; col 10, ln 33-51; col 11, ln 60 to col 12, ln 56 and Figs 4-5	1-25
A	US 2006/0250409 A1 (BANDO et al.) 09 November 2006 (09.11.2006)	1-25
A	US 7,330,188 B1 (SOLANKI et al.) 12 February 2008 (12.02.2008)	1-25
A	US 7,164,426 B1 (DULUK et al.) 16 January 2007 (16.01.2007)	1-25
A	US 7,136,068 B1 (PRIEM et al.) 14 November 2006 (14.11.2006)	1-25

☐

Further documents are listed in the continuation of Box C.

☐

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

07 July 2009 (07.07.2009)

Date of mailing of the international search report

17 JUL 2009

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774