

METHOD AND SYSTEM FOR MULTIPLE 3-D GRAPHIC PIPELINE OVER A PC BUS

Field of the Invention

The present invention relates to a method and system for 3-D (three-dimensional) multiple graphic processing. More particularly, the invention relates to a method and system for improving the parallelization of image processing by Graphic Processing Units (GPUs), based on unified framework of three parallelization methods, which are time division, image division and object division methods.

Definitions, Acronyms and Abbreviations

Throughout this specification, the following definitions are employed:

GPU: GPU (Graphic Processing Unit) like the CPU (Central Processing Unit), a GPU is a single-chip processor which is used primarily for computing 3-D functions. This includes tasks such as, lighting effects, object transformations, and 3-D motion. These are mathematically-intensive tasks, which otherwise, would put quite a strain on the CPU, but since the specialized GPU can handle most of the 3-D calculations, it helps the computer to perform these tasks more efficiently, and, of course, faster.

Polygon: Polygons in 3-D graphics are two-dimensional primitives, allowing generating and displaying of 3-D complex graphical objects. The polygons are the faces of the object and are composed from N vertices. Actually, a polygon is a closed plane figure, bounded by three or more line segments.

Frame Buffer: a Frame Buffer (FB) is a buffer that stores the contents of an image, pixel by pixel. Generally, the portion of memory is reserved for holding the complete bit-mapped image that is sent to the monitor, for display.

Typically the frame buffer is stored in the memory chips on the video adapter. In some instances, however, the video chipset is integrated into the motherboard design, and the frame buffer is stored in the general main memory.

Object compositing unit: performs re-composition of multiple three-dimensional rasters into final image. The merged data is resolved for the closest pixel to the viewer in 3-D space, based on the depth value of pixels. The new method, based on autonomous associative decision, allows the use of multiple GPUs for any frame complexity.

Display list: a Display List is a description of the 3-D scene through a list of graphic primitives, such as polygons and attributes. The display list provides intermediate image storage for quick image retrieval.

Vertex array: a Vertex Array is an array of vertices describing the 3-D scene. A Vertex Array provides intermediate image storage for quick image retrieval.

Alpha blending: Alpha blending controls the way in which the graphic information is displayed, such as levels of transparency, or opacity.

Background of the Invention

The three dimensional graphic pipeline architecture breaks-down into segmented stages of CPU, Bus, GPU vertex processing and GPU fragment (pixel) processing. A given pipeline is only as strong as the weakest link of one of the above stages, thus the main bottleneck determines the overall throughput. Enhancing performance is all that required for reducing or eliminating bottlenecks. The major bottleneck strongly depends on the application. Extreme cases are CAD-like (Computer Aided Design) applications, characterized by an abundance of polygons (vertices), vs. video-

game applications having a small polygon count but intensive fragment activity (e.g., texturing). The first class suffers from vertex processing bottlenecks, while the second class suffers from fragment bottlenecks. Both are frequently jammed over the PC bus. Many applications have mixed characteristics, where bottlenecks may randomly alternate between extremes, on a single frame basis.

The only way to improve the performance of the GPU is by means of parallelizing multiple GPUs according to one of the bottleneck solving methods. There are two predominant methods for rendering graphic data with multiple GPUs. These methods include time division (time domain composition), in which each GPU renders the next successive frame, and image division (screen space composition), in which each GPU renders a subset of the pixels of each frame. The third one, much less popular, is the object division (polygon decomposition) method.

In the time division method each GPU renders the next successive frame. It has the disadvantage of having each GPU render an entire frame. Thus, the speed at which each frame is rendered is limited to the rendering rate of a single GPU. While multiple GPUs enable a higher frame rate, a delay can be imparted in the response time (latency) of the system to a user's input. This occurs because, while at any given time, only one GPU is engaged in displaying a rendered frame, each of the GPUs is in the process of rendering one of a series of frames in a sequence. To maintain the high frame rate, the system delays the user's input until the specific GPU, which first received the signal cycles through the sequence, is again engaged in displaying its rendered frame. In practical applications, this condition serves to limit the number of GPUs that are used in a system. With large data sets, there is another bottleneck, due to the fact that each GPU must be able to access all the data. This requires

either maintaining multiple copy operations of large data sets or possible conflicts in accessing the single copy operation.

Image division method splits the screen between N GPUs, such that each one displays $1/N$ of the image. The entire polygon set is transferred to each GPU for processing, however, the pixel processing is significantly reduced to the window size. Image division has no latency issues, but it has a similar bottleneck with large data sets, since each GPU must examine the entire database to determine which graphic elements fall within the portion of the screen allocated to said GPU. Image division method suits applications with intensive pixel processing.

Object division method is based on distribution of data subsets between multiple GPUs. The data subsets are rendered in the GPU pipeline, and converted to Frame Buffer (FB) of fragments (sub-image pixels). The multiple FB's sub-images have to be merged (composited) to generate the final image to be displayed. Object division delivers parallel rendering on the level of a single frame of very complex data consisting of large amount of polygons. The input data is decomposed in the polygon level and re-composed in the pixel level. A proprietary driver intelligently distributes data streams, which are generated by the application, between all GPUs. The rasters, generated by the GPUs, are composited into final raster, and moved to the display. The object division method well suits applications that need to render a vast amount of geometrical data. Typically, these are CAD, Digital Content Creation, and comparable visual simulation applications, considered as "viewers," meaning that the data has been pre-designed such that their three-dimensional positions in space are not under the interactive control of the user. However, the user does have interactive control over the viewer's position, the direction of view, and the scale of the graphic data. The user also may have control over the selection of a subset of the data and the method by which it is rendered.

This includes manipulating the effects of image lighting, coloration, transparency and other visual characteristics of the underlying data.

In above applications, the data tends to be very complex, as it usually consists of massive amount of geometrical entities at the display list or vertex array. Therefore, the construction time of a single frame tends to be very long (e.g., typically 0.5 sec for 20 million polygons), which in turn slows down the overall system performance.

Therefore, there is a need to provide a system which can guarantee the best system performance, being exposed to high traffic over the PC (Personal Computer) Bus.

It is an object of the present invention to provide an amplified strength of the GPU by means of parallelizing multiple GPUs.

It is another object of the present invention to provide a system, wherein the construction time of a single frame does not slow down the overall system response.

It is still another object of the present invention to provide a system and method, wherein the graphic pipeline bottlenecks of vertex processing and fragment processing are transparently and intelligently resolved.

It is still a further object of the present invention to provide a system and method that has high scalability and unlimited scene complexity.

It is still a further object of the present invention to provide a process overcoming difficulties that are imposed by the data decomposition, which is partition of data and graphic commands between GPUs.

It is still a further object of the present invention to provide a method and system for an intelligent decomposition of data and graphic commands, preserving the basic features of graphic libraries as state machines and complying with graphic standards.

Other objects and advantages of the invention will become apparent as the description proceeds.

Summary of the Invention

The present invention is directed to a system for improving the parallelization of image processing, using one or more parallelization modes, wherein the image that is displayed on at least one computer screen by one or more Graphic Processing Units, which comprises:

one or more software applications, for issuing graphic commands;

one or more graphic libraries, for storing data used to implement the graphic commands;

one or more Software Hub Drivers, for controlling a Hardware Hub, for interacting with the operation system of the computer and the graphic libraries, for performing real-time analysis of a data stream, from which frames of the image are generated, for determining the parallelization mode of each GPU, and for forwarding the data stream or a portion thereof to each GPU;

one or more GPU Drivers, for allowing the GPUs to interact with the graphic libraries; and

at least one I/O module for interconnecting between the Software module and the Hardware Hub,

wherein, the Hardware Hub distributes between the GPUs, for each frame, graphic commands and the data stream or a portion thereof, according to their relative complexity within the image, and defines the complexity The Software

Hub Driver also composites a graphics output for display, using the outputs obtained from at least one GPU, while alternating, whenever required, the parallelization mode for the each frame.

Parallelization is based on an object division mode or on an image division mode or on a time division mode or on any combination thereof. The hardware hub comprises a compositing unit for composing a complete frame from processed portions of the data stream. The hardware hub comprises a hub router for routing polygonal data, for routing graphic command stream, for routing pixel data and for routing the results of composition, while operating in the object division mode or in the image division mode or in the time division mode or in any combination thereof. The hardware hub comprises a control unit for receiving commands from the Software Hub Driver within the I/O module. The hardware hub comprises a memory unit for storing intermediate processing results of one or more GPUs and data required for composition and transferring the processed data for display.

Preferably, the Software Hub Driver is capable of performing the following operations:

- interception of the graphic commands from the standard graphic library by means of the OS interface and utilities;

- forwarding and creating graphic commands to the GPU Driver by means of the OS interface and utilities;

- controlling the Hardware Hub, registry and installation operations by means of the OS interface and utilities;

- maintaining the consistency of graphic machine states across the GPUs, based on the input graphic commands stream, while using state monitoring;

- estimating the type of graphic load and overload in the executed application graphic context, while using application and graphic resources analysis;

load estimation of the GPUs load balance based on graphic commands stream and time measurements, while using application and graphic resources analysis;

adjusting the load distribution between GPUs according to feedback received from each GPU regarding the load balance, while using application and graphic resources analysis;

performing manipulation in graphic functions according to the current parallelization mode; and

controlling the distributed graphic functions, while modifying the graphic commands and the data stream according to the current parallelization mode.

The present invention is directed to a method for improving the parallelization of image processing, using one or more parallelization modes, wherein the image that is displayed on at least one computer screen by one or more Graphic Processing Units. Software applications are provided for issuing graphic command and graphic libraries are provided for storing data used to implement the graphic commands. A Software Hub Drivers is provided for controlling a Hardware Hub, for interacting with the operation system of the computer and the graphic libraries, for performing real-time analysis of a data stream, from which frames of the image are generated, for determining the parallelization mode of each GPU, and for forwarding the data stream or a portion thereof to each GPU. GPU Drivers are provided for allowing the GPUs to interact with the graphic libraries and an I/O module is provided for interconnecting between the Software module and the Hardware Hub. graphic commands and the data stream or a portion thereof are distributed between the GPUs for each frame by the Hardware Hub, according to their relative complexity within the image, wherein the complexity is defined by the Software Hub Driver. The Software Hub Driver also composites a graphics output for display, using the outputs

obtained from at least one GPU, while alternating, whenever required, the parallelization mode for the each frame.

Whenever the parallelization mode is an Object division parallelization mode, the following steps are performed:

for each frame, generating a stream of graphic operations and polygonal data;

marking the polygonal data and graphic commands by means of the Software Hub Driver for distribution between multiple GPUs;

sending the marked data to the Hardware Hub;

distributing the marked data via the Hub Router to the multiple GPUs;

rendering the data by means of GPUs;

retrieving the data from the Frame Buffers and forwarding the retrieved data to the compositing unit via the Hub Router;

compositing the content of the Frame Buffers into a single Frame Buffer; and

forwarding the content of the single Frame Buffer to at least one designated GPU for display.

Whenever the parallelization mode is an Image division parallelization mode, the following steps are performed:

subdividing the screen to portions and assigning different viewports to GPUs by means of the Software Hub Driver;

moving the entire polygonal data and graphic commands to the Hub Router;

transmitting the entire polygonal data and graphic commands to GPUs, wherein each GPU receives the same data;

rendering the data by means of GPUs;

forwarding a portion of the content stored in the Frame Buffers to compositing unit in Hardware Hub for the complete image creation; and

forwarding the image to at least one designated GPU for display.

Whenever the parallelization mode is a Time division parallelization mode, the following steps are performed:

forwarding to each one of the multiple GPUs the entire amount of polygons for rendering;

redirecting the entire polygonal data and graphic commands by means of Software Hub Driver to all GPUs, while alternating between them;

rendering the data by means of GPUs;

transferring rendered data from at least one GPU via the Hub Router; and

redirecting the resulting content of the Frame Buffer via Hub Router to at least one designated GPU for display.

The distribution of polygons between multiple GPUs is performed by distributing blocks of data between multiple GPUs and by testing each graphic operation for blocking mode, in which one or more parallelization modes are carried out, thereafter. The data is redirected in regular non-blocking path to at least one designated GPU, This process is repeated until a blocking operation is detected. Then GPUs are synchronized by performing a flush operation in order to terminate rendering and clean up the internal pipeline in each GPU; performing a composition operation for merging the contents of the Frame Buffers into a single Frame Buffer and bytransmitting the single Frame Buffer back to all GPUs. Then the composited complete frame is terminated at all GPUs, except one or more designated GPUs, whenever a Swap operation is detected and displaying the image by means of the one or more designated GPUs. the same data is processed by all GPUs, as long as the blocking mode is active and the Swap operation is not detected. whenever the blocking mode is inactive, the designated data is further processed by multiple GPUs,.

Brief Description of the Drawings

In the drawings:

- Fig. 1 is a block diagram of a multiple GPU architecture system, according to an embodiment of the present invention;
- Fig. 2 is a block diagram of Hardware Hub components, according to an embodiment of the present invention;
- Fig. 3 is a block diagram of Object division parallelization mode, according to an embodiment of the present invention;
- Fig. 4 is a block diagram of Image division parallelization mode, according to an embodiment of the present invention;
- Fig. 5 is a block diagram of Time division parallelization mode, according to an embodiment of the present invention;
- Fig. 6 is a schematic block diagram of a possible integration of the Software Hub Driver into the operating system environment, according to an embodiment of the present invention;
- Fig. 7 is a functional block diagram presenting the main tasks of the Software Hub Driver, according to an embodiment of the present invention;
- Fig. 8 is a flow chart presenting an process for distribution of the polygons between the multiple GPUs, according to an embodiment of the present invention; and
- Fig. 9 discloses a sample configuration of the system, employing 8 GPUs, according to an embodiment of the present invention.

It will be appreciated that for simplicity and clarity of illustration, elements shown in the figures have not necessarily been drawn to scale. For example, the dimensions of some of the elements may be exaggerated relative to other elements for clarity. Further, where considered appropriate, reference

numerals may be repeated between the figures to indicate corresponding or analogous elements.

Detailed Description of the Preferred Embodiments

The Multiple 3-D Graphic Pipeline

The current invention calls for the introduction of an extended PC graphic architecture including novel operational component, the 3-D pipeline Hub.

Fig. 1 presents multiple GPU architecture system 100, according to an embodiment of the present invention. The hub 110 is located in a unique position on the PC bus, between the CPU (Central Processing Unit) and a cluster of GPUs 130. The multiple graphic pipeline architecture, as described in Fig. 1, consists of Software Application 121, Graphic Library 122, Software Hub Driver 123, GPU Driver 124, Hardware Hub 110, cluster of GPUs 130, and display(s) 140. Usually, one of the GPUs is designated as a display unit. It should be noted, that it is possible to have more than one display unit, or include a display unit directly inside the Hardware Hub. A display unit can drive multiple screens, as well.

The Hub mechanism consists of a Hardware Hub component 110, located on the PC bus between I/O (Input/Output) chipset 160, being a I/O module, and a cluster of GPUs 130, and a Software module comprising Software Hub Driver 123, executed by the PC.

The Hardware Hub 110 carries on at least the following action:

- distributes decomposed polygon stream and graphic commands between GPUs;
- composites graphics output for display according to different parallel modes;

- 13 -

- makes cache of data; and
- alternates modes of parallelism.

The Software Hub Driver 123, besides controlling the Hardware Hub 110, also carries on at least the following actions:

- interacts with the OS (Operation System) and graphic library, such as OpenGL, DirectX;
- performs real-time analysis of the data stream;
- determines the parallelization mode; and
- decomposes the data and command stream.

One advantage of this method is that the unique location of the Hardware Hub 110 allows it to control the graphic pipeline, while being transparent to the application 121. The application 121, along with Graphic Library 122, such as OpenGL, keeps working as if it was a single GPU.

Another advantage of this method is that the unique location of the Hardware Hub 110 allows it to control a graphic pipeline between the User Interface 150 and Display 140, while being transparent to the GPU. Each GPU of the cluster keeps working as if it is the only graphic processor hooked on the I/O chipset 160.

Still another advantage of this method is that the unique location of the Hardware Hub 110 allows it to control the graphic pipeline for different parallelization modes: image division mode, time division mode or object division mode.

Further advantage of this method is that the unique location of the Hardware Hub 110 allows it to sense in real-time the varying parameters of application's

load, such as polygon count, texture volume, human interaction, and to intelligently decide and carry on the current optimal parallelization method.

It should be noted that according to all embodiments of the present invention, the display(s) 140 may be connected directly or indirectly to the Hardware Hub 110.

Reference is now made to Fig. 2, which discloses the components of the Hardware Hub 110, according to an embodiment of the present invention. Control Unit 210, accepts proprietary commands from the Software Hub Driver over the bus, and accordingly controls the other units. Hub Router 230 routes polygonal data and graphic command stream from left to right, routes pixel data from right to compositing unit, and routes compositing results to the right. Compositing Unit 240 performs various compositing schemes according to parallelization mode. Auxiliary Memory block 220 is used for storing intermediate processing results of one or more GPUs, data required for composition and transferring the processed data for display.

The Hardware Hub 110 utilizes its units according to parallelization modes: an Object division mode, an Image division mode, a Time division mode. These modes are adaptively handled by the Hardware Hub 110 according to application needs.

Fig. 3 discloses the operation of the Object division parallelization mode, according to an embodiment of the present invention. CPU executes the 3-D graphic application 310, which along with standard graphic library generate a stream of graphic operations and polygonal data. They are typically organized in data blocks either as Display List, Vertex Array or free polygons, which are polygons that are neither organized in Display List nor in Vertex Array. The Software Hub Driver at step 320 marks the polygonal data and graphic

commands for distribution between multiple GPUs, in a way that the overall load is balanced. The marked data is forwarded to Hardware Hub. At step 330 it is distributed via the Hub Router to multiple GPUs. After the rendering process in GPUs at step 340 is accomplished, the Frame Buffers are retrieved and forwarded via the Hub Router to the compositing unit at step 350. Here the frame buffers are composited to a single frame buffer, which is forwarded to the designated GPU for display. The single frame buffer is displayed at step 360.

Fig. 4 discloses the operation of the Image division parallelization mode, according to an embodiment the present invention. In this mode the Software Hub Driver assigns different viewports to GPUs in order to subdivide the screen between them. The viewports aspects are set according to load balancing considerations, to keep the overall GPU load evenly balanced. In step 420, the entire polygonal data and graphic commands are moved to the Hub Router at Hardware Hub, and they are transmitted at step 430 to multiple GPUs. All GPUs receive the same data. After rendering at step 440, the partial frame buffers are brought to compositing unit in Hardware Hub for the full image creation at step 450, and then this image is moved out to designated GPU for display. The full image is displayed at step 460.

Fig. 5 discloses the operation of the Time division parallelization mode, according to an embodiment the present invention. In time division the processing of each frame takes N frame time units, while N GPUs (or N clusters of GPUs) are participating. The entire amount of polygons is forwarded to each GPU, for rendering. At each frame time unit, the Software Hub Driver redirects the polygonal data and graphic commands at step 530 to a cluster of GPUs at a time, while alternating between them. The data is transferred to the above cluster of GPUs via the Hub Router, rendered in the GPUs at step 540, and then the resulting frame buffer at step 550 is redirected

via Hub Router to the designated GPU for display. All GPUs are coordinated by Software Hub Driver to create a continuous sequence of frames. The resulting frame buffer is displayed at step 560.

The Hardware Hub competence is its scaling technology: Architecture and cluster of proprietary processes devoted to scale existing GPUs performance in PC based systems, by enabling the use of multiple GPUs in parallel on the level of chip, card or chip IP (Intellectual Property) core, and handling multiple bus paths between the CPU and GPU. The technology achieves linear increase in performance. It is invariant to a graphics vendor and also it is transparent to an application. In the present invention, the graphic pipeline bottlenecks of vertex processing, fragment processing and bus transfer are completely and intelligently resolved. As bottlenecks may shift between frames, the Hardware Hub is designed with a smart real-time feedback system between the Control Unit 210, disclosed in Fig. 2, and Software Hub Driver 123, disclosed in Fig. 1, by means of the bus, utilizing the different parallelization modes to overcome different bottlenecks and maintain maximum performance at the frame level.

The Software Hub Driver

The Software Hub Driver is a software package residing in the PC and coexisting with computer's operating system, standard graphic library, application and Vendor's GPU Driver. Fig. 6 is a schematic block diagram of a possible integration of the Software Hub Driver 630 into the operating system environment according to an embodiment of the present invention. Next to graphic application block 610 there is standard graphic library block 620. The Software Hub Driver 630 is located beneath the standard graphic library 620, intercepting the graphic command and data stream on the way to the Vendor's GPU Driver 640. The Software Hub Driver 630 also controls the Hardware Hub 660.

Fig.7 is a functional block diagram presenting the main tasks of the Software Hub Driver, according to an embodiment of the present invention. OS interface and Utilities block 710 is responsible for interception of the graphic commands from the standard graphic library, forwarding and creating graphic commands to Vendor's GPU Driver, controlling the Hardware Hub, registry and installation, OS services and utilities. State Monitoring block 720 is responsible for maintaining consistency of graphic machine states across the GPUs, based on the input graphic commands stream. Application and graphic resources analysis block 730 is responsible for the application observation - estimating the type of graphic load and bottleneck in the current application graphic context, graphic resources (GPUs) load estimation for load balance based on graphic commands stream and time measurements, handling the feedback from GPUs in regard to load balancing. Parallelism policy management block 740 is based on load analysis. All parallelization modes, which are the Object division mode, Image division mode and Time division mode, are combined together in order to achieve best performance and optimal load handling. Parallelization policy is based on the analysis of the load, and it must preserve the state of the graphic system at all relevant GPUs across the electronic circuit or chip. For example, changing of a state by adding a new light source in the scene at some time point, must affect all subsequent polygons at different GPUs. Parallelism policy management block 740 is responsible for the interpretation of the policy for specific manipulation in graphic functions. Distributed graphic functions control block 750 is responsible for modification of graphic command and data stream based on the parallelization policy.

Object Division Decomposition Process

Object division is a well known concept, but data decomposition (partition of data and graphic commands between GPUs), while being also a known concept, has not been applied yet effectively, as it imposes various great

difficulties. These difficulties are handled successfully by a proposed process and its implementation according to the present invention.

The decomposition, and more importantly, the composition, must be accurate and efficient. Certain operations must be performed in the order they are submitted by the application. For example, in case of semi-transparency, the commands and polygon stream must keep a certain order for creating a correct graphic result.

Intelligent decomposition of data and graphic commands is needed, preserving the basic features of graphic libraries as state machines, and complying with the graphic standards. The proposed decomposition process, according to the present invention, is performed by the Software Hub Driver. CPU runs the 3-D graphic application, generating flow of graphic commands and data. They are typically organized in blocks, such as Display Lists or Vertex Arrays, stored in the system memory.

According to the present invention, the Software Hub Driver, running in the CPU, decomposes the set of scene polygons (or vertices). Their physical distribution is performed by the Hardware Hub.

The polygons are rendered in the GPU, while maintaining the resulting Frame Buffer in local memory. All FBs are transferred, via Hub Router, to compositing unit in Hardware Hub, to be merged into single FB. Finally, the composited FB is forwarded for display.

The Software Hub Driver carries out the following process of distribution of the polygons between the multiple GPUs. It is assumed, that the regular way the graphic application works, remains unchanged. Per frame, a typical application generates a stream of graphic calls that includes blocks of graphic

data; each block consists of a list of geometric operations, such as single vertex operations or buffer based operations (vertex array). Typically, the decomposition process splits the data between GPUs preserving the blocks as basic data units. Geometric operations are attached to the block(s) of data, instructing the way the data is handled. A block is directed to designated GPUs. However, there are operations belonging to the group of Blocking Operations, such as Flush, Swap, Alpha blending, which affect the entire graphic system, setting the system to blocking mode. Blocking operations are exceptional in that they require a composed valid FB data, thus in the parallel setting of the present invention, they have an effect on all GPUs. Therefore, whenever one of the Blocking operations is issued, all the GPUs must be synchronized. Each frame has at least 2 blocking operations: Flush and Swap, which terminate the frame.

Fig. 8 is a flow chart presenting the process for distribution of the polygons between the multiple GPUs, according to an embodiment the present invention. The frame activity starts with distributing blocks of data between GPUs. Each graphic operation is tested for blocking mode at step 820. In a regular (non-blocking) path, data is redirected to the designated GPU at the step 830. This loop is repeated until a blocking operation is detected.

When the blocking operation is detected, all GPUs must be synchronized at step 840 by at least the following sequence:

- performing a flush operation in order to terminate rendering and clean up the internal pipeline (flushing) in GPU;
- performing a composition in order to merge the contents of FBs into a single FB; and
- transmitting the contents of said single FB back to all GPUs, in order to create a common ground for continuation.

The Swap operation activates the double buffering mechanism, swapping the back and front color buffers. If Swap is detected at step 850, it means that a composited complete frame must be terminated at all GPU, except GPU0. All GPUs have the final composited contents of a FB designated to store said contents, but only the one connected to the screen (GPU0) displays the image at step 860.

Another case is operations that are applied globally to the scene and need to be broadcasted to all the GPUs. If one of the other blocking operations is identified, such as Alpha blending for transparency, then all GPUs are flushed as before at step 840, and merged into a common FB. This time the Swap operation is not detected (step 850), and therefore all GPUs have the same data, and as long as the blocking mode is on (step 870), all of them keep processing the same data (step 880). If the end of the block mode is detected at step 870, GPUs return working on designated data (step 830).

Adaptive handling of graphic load by combining three division methods

In addition, the present invention introduces a dynamic load-balancing technique that combines the object division method with the image division and time division methods in image and time domains, based on the load exhibits by previous processing stages. Combining all the three parallel methods into a unified framework dramatically increases the effectiveness of our invention.

Parallel processing is implemented by a pipeline, such as any common GPU allows the data to be processed in parallel, in time, image and object domains. The processing performed on the graphical processing system, either in parallel on multi-GPU or sequential, results in a sequence of complete raster images stored in a frame buffer, and sent to the display unit. These images are

referred as frames in short. A frame consists of fragments. A fragment is an extended pixel stored in memory, which consists of attributes such as color, alpha, depth, stencil, etc. When processing is performed in parallel in the time domain, typically each GPU is responsible for the production of a complete frame. In the other two domains, which are the image and object domains, all GPU operate in parallel to produce a single frame. Screen-space parallel-processing implies that each GPU renders a subset of the fragments of each frame, and object parallel-processing implies that the input data for each frame, in particular the geometric data (e.g., the polygon set representing the scene) is distributed between the multi GPUs.

Each one of three domains (time, image and object domains) has advantages and disadvantages. The effectiveness of each discipline is a dynamic function based on input data. Moreover, in many cases no single discipline is superior. In these cases a combination of two or even all the three disciplines may yield the most optimum results.

The present invention describes the specification for a parallel-processing system of three-dimensional data. It provides a novel process for object parallel-processing that consists of efficient decomposition of the data between the different GPU, and then the composition of the frames produced on the various GPUs into a final frame ready to be rendered.

The present invention provides a method to integrate all the three parallel modes dynamically into a unified framework to achieve maximum load balancing. At each frame, the set of available GPUs can be reconfigured based on the time it took to render the previous frames, and the bottlenecks exhibited during the processing of these frames.

Fig. 9 discloses a sample configuration of the system, employing 8 GPUs, according to an embodiment of the present invention. According to the above sample configuration, a balanced graphic application is assumed. The GPUs are divided into two groups for time division parallelism. GPUs indexed with 1, 2, 3, and 4 are configured to process even frames and GPUs indexed with 5, 6, 7, and 8 are configured to process odd frames. Within each group, two GPU subgroups are set for image division: the GPUs with the lower indexes (1,2 and 5,6 respectively) are configured to process half of the screen, and the high-indexed GPU (3,4 and 7,8 respectively) are configured to process the other half. Finally, for the object division, GPUs indexed with 1, 3, 5 and 7 are fed with half of the objects, and GPUs indexed with 2, 4, 6 and 8 are fed with the other half of the objects.

If at some point the system detects that the bottlenecks exhibited in previous frames occur at the raster stage of the pipeline, it means that fragment processing dominates the time it takes to render the frames and that the configuration is imbalanced. At that point the GPUs are reconfigured, so that each GPU will render a quarter of the screen within the respective frame. The original partition for time division, between GPUs 1,2,3,4 and between 5,6,7,8 still holds, but GPU 2 and GPU 5 are configured to render the first quarter of screen in even and odd frames respectively. GPUs 1 and GPU 6 - the second quarter, GPU 4 and GPU 7 - the third quarter, and GPU 3 and GPU 8 - the forth quarter. No object division is implied.

In addition, if at some point the system detects that the bottleneck exhibited in previous frames occurs at the geometry stage of the pipe, the GPUs are reconfigured, so that each GPU will process a quarter of the geometrical data within the respective frame. That is, GPU 3 and GPU 5 are configured to process the first quarter of the polygons in even and odd frames respectively.

GPU 1 and GPU 7 - the second quarter, GPU 4 and GPU 6 - the third quarter and GPU 2 and GPU 8 - the forth quarter. No image division is implied.

It should be noted, that taking 8 GPUs is sufficient in order to combine all three parallel modes, which are time, image and object division modes, per frame. Taking the number of GPUs larger than 8, also enables combining all 3 modes, but in a non-symmetric fashion. The flexibility also exists in frame count in a time division cycle. In the above example, the cluster of 8 GPUs was broken down into the two groups, each group handling a frame. However, it is possible to extend the number of frames in a time division mode to a sequence, which is longer than 2 frames, for example 3 or 4 frames.

Taking a smaller number of GPUs still allows the combination of the parallel modes, however the combination of two modes only. For example, taking only 4 GPUs enables to combine image and object division modes, without time division mode. It is clearly understood from Fig. 9, while taking the group of GPU1, GPU2, GPU3 and GPU4, which is the left cluster. Similarly, the group of GPU1, GPU2, GPU5 and GPU6, which is the upper cluster, employs both object and time division modes. Finally, the configuration of the group of GPU2, GPU4, GPU5 and GPU6, which is the middle cluster, employs image and time division modes.

It should be noted, that similarly to the above embodiments, any combination between the parallel modes can be scheduled to evenly balance the graphic load.

It also should be noted, that according to the present invention, the parallelization process between all GPUs may be based on an object division mode or image division mode or time division mode or any combination thereof in order to optimize the processing performance of each frame.

While some embodiments of the invention have been described by way of illustration, it will be apparent that the invention can be put into practice with many modifications, variations and adaptations, and with the use of numerous equivalents or alternative solutions that are within the scope of persons skilled in the art, without departing from the spirit of the invention or exceeding the scope of the claims.

Claims

1. A system for improving the parallelization of image processing, using one or more parallelization modes, wherein said image that is displayed on at least one computer screen by one or more Graphic Processing Units, comprising at least:
 - a. one or more software applications, for issuing graphic commands;
 - b. one or more graphic libraries, for storing data used to implement said graphic commands;
 - c. one or more Software Hub Drivers, for controlling a Hardware Hub, for interacting with the operation system of said computer and said graphic libraries, for performing real-time analysis of a data stream, from which frames of said image are generated, for determining the parallelization mode of each GPU, and for forwarding said data stream or a portion thereof to each GPU;
 - d. one or more GPU Drivers, for allowing said GPUs to interact with said graphic libraries; and
 - e. at least one I/O module for interconnecting between said Software module and said Hardware Hub,wherein, said Hardware Hub distributes, for each frame, between said GPUs, graphic commands and said data stream or a portion thereof, according to their relative complexity within said image, said complexity is defined by said Software Hub Driver; and composites a graphics output for display, using the outputs obtained from at least one GPU, while alternating, whenever required, said parallelization mode for said each frame.
2. System according to claim 1, wherein the parallelization is based on an object division mode or on an image division mode or on a time division mode or on any combination thereof.

3. System according to claim 1, wherein the hardware hub comprises at least one compositing unit at least for composing a complete frame from processed portions of the data stream.
4. System according to claim 2, wherein the hardware hub comprises at least one hub router for routing polygonal data, for routing graphic command stream, for routing pixel data and for routing the results of composition, while operating in the object division mode or in the image division mode or in the time division mode or in any combination thereof.
5. System according to claim 1, wherein the hardware hub comprises at least one control unit for receiving commands from the Software Hub Driver within the I/O module.
6. System according to claim 1, wherein the hardware hub comprises a memory unit for storing intermediate processing results of one or more GPUs and data required for composition and transferring the processed data for display.
7. System according to claim 1, wherein the Software Hub Driver is capable of performing the following operations:
 - a. interception of the graphic commands from the standard graphic library by means of the OS interface and utilities;
 - b. forwarding and creating graphic commands to the GPU Driver by means of the OS interface and utilities;
 - c. controlling the Hardware Hub, registry and installation operations by means of the OS interface and utilities;
 - d. maintaining the consistency of graphic machine states across the GPUs, based on the input graphic commands stream, while using state monitoring;

- e. estimating the type of graphic load and overload in the executed application graphic context, while using application and graphic resources analysis;
 - f. load estimation of the GPUs load balance based on graphic commands stream and time measurements, while using application and graphic resources analysis;
 - g. adjusting the load distribution between GPUs according to feedback received from each GPU regarding said load balance, while using application and graphic resources analysis;
 - h. performing manipulation in graphic functions according to the current parallelization mode; and
 - i. controlling the distributed graphic functions, while modifying said graphic commands and said data stream according to said current parallelization mode.
8. A method for improving the parallelization of image processing, using one or more parallelization modes, wherein said image that is displayed on at least one computer screen by one or more Graphic Processing Units, comprising at least:
- a. providing one or more software applications, for issuing graphic commands;
 - b. providing one or more graphic libraries, for storing data used to implement said graphic commands;
 - c. providing one or more Software Hub Drivers, for controlling a Hardware Hub, for interacting with the operation system of said computer and said graphic libraries, for performing real-time analysis of a data stream, from which frames of said image are generated, for determining the parallelization mode of each GPU, and for forwarding said data stream or a portion thereof to each GPU;

- d. providing one or more GPU Drivers, for allowing said GPUs to interact with said graphic libraries.
 - e. providing at least one I/O module for interconnecting between said Software module and said Hardware Hub; and
 - f. for each frame, distributing between said GPUs and by means of said Hardware Hub, graphic commands and said data stream or a portion thereof, according to their relative complexity within said image, said complexity is defined by said Software Hub Driver; and compositing a graphics output for display, using the outputs obtained from at least one GPU, while alternating, whenever required, said parallelization mode for said each frame.
9. A method according to claim 8, wherein the parallelization mode is an Object division parallelization mode and the following steps are performed:
- a. for each frame, generating a stream of graphic operations and polygonal data;
 - b. marking the polygonal data and graphic commands by means of the Software Hub Driver for distribution between multiple GPUs;
 - c. sending the marked data to the Hardware Hub;
 - d. distributing said marked data via the Hub Router to said multiple GPUs;
 - e. rendering the data by means of GPUs;
 - f. retrieving the data from the Frame Buffers and forwarding the retrieved data to the compositing unit via the Hub Router;
 - g. compositing the content of said Frame Buffers into a single Frame Buffer; and
 - h. forwarding the content of said single Frame Buffer to at least one designated GPU for display.

10. A method according to claim 8, wherein the parallelization mode is an Image division parallelization mode and the following steps are performed:
 - a. subdividing the screen to portions and assigning different viewports to GPUs by means of the Software Hub Driver;
 - b. moving the entire polygonal data and graphic commands to the Hub Router;
 - c. transmitting said entire polygonal data and graphic commands to GPUs, wherein each GPU receives the same data;
 - d. rendering the data by means of GPUs;
 - e. forwarding a portion of the content stored in the Frame Buffers to compositing unit in Hardware Hub for the complete image creation; and
 - f. forwarding said image to at least one designated GPU for display.
11. A method according to claim 8, wherein the parallelization mode is a Time division parallelization mode and the following steps are performed:
 - a. forwarding to each one of the multiple GPUs the entire amount of polygons for rendering;
 - b. redirecting the entire polygonal data and graphic commands by means of Software Hub Driver to all GPUs, while alternating between them;
 - c. rendering the data by means of GPUs;
 - d. transferring rendered data from at least one GPU via the Hub Router; and
 - e. redirecting the resulting content of the Frame Buffer via Hub Router to at least one designated GPU for display.
12. A method according to claim 8, wherein the distribution of polygons between multiple GPUs is performed by:
 - a. distributing blocks of data between multiple GPUs;

- b. testing each graphic operation for blocking mode, in which one or more parallelization modes are carried out;
 - c. redirecting the data in regular non-blocking path to at least one designated GPU;
 - d. repeating step (b) and (c) until a blocking operation is detected;
 - e. synchronizing GPUs by the following sequence:
 - e.1. performing a flush operation in order to terminate rendering and clean up the internal pipeline in each GPU;
 - e.2. performing a composition operation for merging the contents of the Frame Buffers into a single Frame Buffer; and
 - e.3. transmitting said single Frame Buffer back to all GPUs;
 - f. terminating the composited complete frame at all GPUs, except one or more designated GPUs, whenever a Swap operation is detected and displaying the image by means of said one or more designated GPUs;
 - g. processing the same data by means of all GPUs, as long as the blocking mode is active and the Swap operation is not detected; and
 - h. continuing to process the designated data by means of multiple GPUs, whenever the blocking mode is inactive.
13. Method according to claim 8, wherein the parallelization is based on an object division mode or on an image division mode or on a time division mode or on any combination thereof.