



White Paper


Subject: 3D Graphical Processing
Author: PowerVR (PowerVR)
Copyright: © PowerVR 2000
File Name: 3Dgraphical processing White Paper.doc
Issue Number: 1.0.5
Issue Date: 14 Nov 00

The background to 3D Graphics

3D graphics refer to the systems used to create and manipulate a modelled "world" or scene on a computer and then to display the world on the 2D computer screen in a realistic way. The world is typically constructed from objects made up from meshes of adjoining triangles or polygons each defined by its vertices. Each vertex has a number of properties including position in 3D space (x, y, z) and colour. Each polygon additionally has some global properties such as texture.

To allow users to interact with the 3D world, either to view it from a new position or to change objects within it, the entire world must be processed to produce a new image to be displayed on a screen "viewport". This processing consists of three main steps: transform and lighting, hidden surface removal and texturing and shading.

3D GRAPHICAL PROCESSING



The Basics of 3D Graphical Processing

Transformation alters the world as objects move within it or as a user's point of view changes. Each change can affect the position of any or all of the vertices within the world. Lighting then performs the calculations necessary to simulate the effect of different lights on objects in the scene, affecting the colour of each vertex as a result.

Finally texturing and shading determines the colour of each pixel in the scene by taking into account both the colour of the polygons and their texture. Textures are stored in memory and the relevant pixel from the texture map (called texels) are retrieved and used to texture each pixel before it is written into the display memory. Depending on the texturing technique, each output pixel requires a different number of texels.

The simplest technique, called point sample texturing, uses a single texel from the texture map, more advanced features interpolate between many texels to more accurately estimate the texture for the output pixel.

Conventional 3D Systems

The conventional algorithmic approach to 3D processing was invented in the 1960s and implemented in high performance workstations for still frame rendering in CAD, architectural engineering and film special effects. This approach was also used in real time flight simulators for the defence industry and later in commercial pilot training.


In such a system each polygon is processed through the 3D pipeline in turn, in the order that they are sent to the hardware. This is referred to as "immediate mode" 3D processing and means that the hardware cannot know whether any part of the scene is complete until the last polygon for that scene has been processed.


Conventional systems use a z-buffer technique, where for each screen pixel a depth value is stored. After texturing and shading, each pixel of each polygon is tested against the stored depth value to determine whether it is closer to the viewer than any pixel drawn in the scene so far; if so the pixel is drawn (possibly blending with the previous pixel) and the depth value updated. During this process many pixels can be textured, shaded and drawn only to be overdrawn later in the scene. The z-buffer memory is the same resolution as the display resolution and is typically 16 or 24bits deep depending on the accuracy required.

PowerVR 3D

The PowerVR approach to 3D graphics starts from the premise that taking a different algorithmic approach to 3D processing can eliminate redundant processing and memory bottlenecks by doing only what is absolutely needed thereby minimising costly accesses to off chip memory.

PowerVR is a "display list renderer", that is, groups of polygons are batched together (into a display list) before being processed by the 3D rendering hardware. This is fundamentally different to the approach used by conventional systems, since it allows a scene to be partitioned into small "tiles" or "regions" each of which is rendered independently, leading to fourkey benefits.





Firstly, because the region is only a small subset of the whole scene, PowerVR can implement key operations on-chip without frequent access to external memory. Hidden surface removal (z-buffering) is done with an on-chip z-buffer and display pixel processing and blending also uses an on-chip "frame buffer tile" as a local storage area. This means that the majority of external memory accesses normal to a conventional 3D system are eliminated. All the on-chip processing is performed at high depth and pixel accuracy at full clock rate without waiting for z-buffer or frame buffer memory accesses that slow down conventional 3D systems.

Secondly, "deferred texturing"; in PowerVR systems the hidden surface removal is completed in the first phase of the pipeline before texturing and shading. This has the benefit that only the visible pixels to be finally drawn in the display memory are textured and shaded eliminating both the redundant work performed and most importantly, the redundant texture fetches from memory required by conventional 3D systems. Thirdly accuracy and image quality; since z-buffering and pixel blending are done entirely on-chip they can be performed at higher precision with no performance degradation. In PowerVR all pixel blend operations are performed with true colour precision, irrespective of the number of translucent layers or the bit-depth of the frame buffer (Internal True Colour™) resulting in high image quality without performance loss.


Finally scalability; because PowerVR systems split the screen into tiles which are independent of each other it is easy to add more processing elements (either within a chip or using multiple chips) in order to increase performance.

Why Tile Based Rendering Is The Future Of Graphics Processing

1. Efficient use of memory bandwidth
2. Memory density is rapidly increasing while bus widths are reaching their limits. In the past graphics performance increases have been supported by the use of wider memory buses. A quick look at the history of 3D graphics controllers shows that the industry has migrated from 32 bit single data rate interfaces, to 64 bit, then 128 bit and now 128 bit double data rate interfaces to support the bandwidth demand. This trend cannot continue for two reasons. Firstly the increasing density of individual memory devices means that frame buffer size would become unsustainable for wide bus implementations and the cost of the graphics subsystem will be unacceptable. Secondly graphics devices will be limited in the number of pins they can use for the bus interface.
3. The continuing drive to lower system costs coupled with increasing memory densities is increasingly encouraging shared or unified memory systems where the graphics device is required to share memory with the processor, making memory bandwidth an even rarer commodity.
4. The trend towards moving transformation and lighting (T&L) into the graphics processor adds to the problem. A T&L engine integrated in the graphics controller also competes for a share of the bandwidth available to the graphics processor.
5. The demand for increasing resolution in modern systems, including HDTV, will further add to bandwidth requirements.

PowerVR performs hidden surface removal (the depth test) and pixel blending entirely on-chip, reducing memory bandwidth requirements significantly.

Texturing though is the operation which dominates use of memory bandwidth. PowerVR's deferred texturing - only texturing pixels which are used on screen - cuts the bandwidth required for texturing and this benefit is becoming increasingly important as 3D applications become better at representing realistic environments with increased freedom of movement in open space. This trend coupled with the rapid increase in scene complexity and polygon count is leading to higher depth complexity (or overdraw - the average number of times a pixel is drawn in a frame). Several years ago the depth of complexity of most games was 1.2 to 1.5. Today's games have a depth complexity of between 2 and 3. This factor is set to increase to between 4 and 5 as the games become more realistic and graphics hardware performance increases. Hardware T&L also tends to increase depth complexity by allowing the generation of even more polygons and restricting the ability of an application to use scene management to restrict depth complexity. In a conventional system, depth complexity is a measure of how much effort is wasted on texturing and drawing pixels which are subsequently overdrawn. PowerVR's deferred texturing means that unnecessary pixels are never drawn with a resulting huge saving in memory bandwidth requirements.



It may be argued that memory bandwidth issues can be resolved by using on-chip DRAM for frame buffer, z buffer and textures. However the drive for higher resolutions, increased colour and depth accuracy and increased texture space means that on chip DRAM cannot satisfy the demand for memory in the foreseeable future. For example a graphics device which supports resolutions up to 1920x1440 pixels would require (1920x1440x2x4) bytes for a double buffered frame buffer, (1920x1440x4) bytes for the z-buffer, as well as texture memory, i.e. nearly 32MB of memory plus texture memory. A tile based renderer can use on-chip DRAM for texture caching alone.

The graph overleaf shows how combinations of high resolution and other features in a 3D application can quickly drive the memory bandwidth required by conventional 3D to levels unachievable even with 128 bit DDR interfaces whereas PowerVR requires less bandwidth and does not require all the bandwidth available even at 1600x1200 resolutions. The assumptions for the application in the graph are:

- Depth complexity of four
- 8-sample anisotropic texture filtering
- Two textures per pixel
- On-chip caching of texture samples to reduce texture fetches by a factor of 1.5
- Screen refresh rate of 72Hz
- It is clear that PowerVR requires around one third of the bandwidth compared to traditional 3D.


Image Quality

PowerVR performs a number of key operations entirely on chip leading to a number of image quality benefits:

Hidden surface removal: In conventional systems, the number of bits used in the z buffer for the depth testing is often tied to the colour depth being rendered and so when rendering to 16bpp displays, the z buffer accuracy is also only 16bpp leading to visual artifacts due to lack of accuracy. On PowerVR, the z buffer for hidden surface removal is on chip and always operates at full 32bit Z/Stencil accuracy, irrespective of the render mode.

Stencil operations: In a similar manner, most conventional systems only support stencil buffer operations when rendering to 32bpp. Again, on PowerVR the stencil buffer is on chip and stencil operations are supported whatever the render mode.

Blending operations: Because traditional 3D systems are immediate mode and render polygons in the order that they are sent, each polygon has to be written out to memory upon completion. If subsequent polygons are blended on those already rendered e.g. to create an explosion effect, the previous ones must be read back into the chip, blended, and the result written back to memory. Apart from consuming memory bandwidth, this can also lead to serious image quality degradation when rendering in 16bpp modes as each polygon is rendered internally at 32bpp, dithered to 16bpp when written to frame buffer and then read back in for use in further blend operations leading to cumulative image degradation as though a video tape copy has been used to make further copies. PowerVR's Internal True Colour™ performs all blending operations on all the pixels in each tile at the full 32bpp colour resolution before performing a single high quality error diffusion dither to 16bpp on output to the frame buffer if necessary. In fact, PowerVR's quality of rendering at 16bpp is the equal or better of many systems quality at 32bpp.



Scalability

PowerVR's tile base architecture is scalable in a number of ways:

Efficient memory usage: Because of its amazingly efficient usage of memory bandwidth, PowerVR can be used in unified memory architectures where the memory bandwidth is shared with the processor - ideal for low power, low cost implementations.

Tile based scalability: Because PowerVR is tile based, it is simple and efficient to share the task of rendering between many processors. Indeed even in single chip solutions, PowerVR performs parallel depth tests in parallel due to its on chip z buffer.

Silicon scalability: Because PowerVR's z buffer, blending and stencil buffers are on chip, the rate at which these operations are performed are all at silicon rate, without the need to wait for fetches from or writes to memory. This means that PowerVR's performance will scale with advances in silicon technology in line with Moore's law, whereas traditional 3D systems, with a memory bandwidth bottleneck, will advance at the rate of increasing memory bandwidth, historically far slower.

Conclusion

PowerVR's tile based architecture provides far more efficient use of memory bandwidth than traditional 3D systems. This along with its inherent scalability and image quality benefits means that it brings key benefits across the spectrum of 3D applications from low power / low memory bandwidth applications to the highest performance systems available. PowerVR systems provide fully featured and standards compatible solutions with wide application support in PC, console and arcade systems today.