



PowerVR Performance Recommendations

The Golden Rules

Public. This publication contains proprietary information which is subject to change without notice and is supplied 'as is' without warranty of any kind. Redistribution of this document is permitted with acknowledgement of the source.

Filename : PowerVR Performance Recommendations.The Golden Rules
Version : PowerVR SDK REL_18.1@5080009a External Issue
Issue Date : 31 May 2018
Author : Imagination Technologies Limited

Contents

The Golden Rules.....	3
1. Do Understand Your Target Device	4
2. Do Profile Your Application	5
3. Do Not Use Alpha Blend Unnecessarily	6
4. Do Perform Clear	7
5. Do Not Update Data Buffers Mid-Frame	8
6. Do Use Texture Compression	9
7. Do Use Mipmapping	10
8. Do Not Use Discard	11
9. Do Not Force Unnecessary Synchronisation	12
10. Do Move Calculations 'Up the Chain'	13
Other Considerations	14
Contact Details	17

The Golden Rules

This document covers key principles developers should follow in order to avoid critical performance flaws in their graphics applications. These recommendations come from the combined experience of the PowerVR Developer Technology Support team and the developers they work with, profiling and optimising their applications and games during development.

1. Do Understand Your Target Device

Seek to learn as much information about your target platforms as possible in order to understand different architectures.

2. Do Profile Your Application

Identify the bottlenecks in your application and determine whether there are opportunities for improvement.

3. Do Not Use Alpha Blend Unnecessarily

Be sure Alpha Blending is used only when required to make the most of deferred architectures and to save bandwidth.

4. Do Perform Clear

Perform a clear on a framebuffer's contents to avoid fetching the previous frame's data on tile-based graphics architectures, which reduces memory bandwidth.

5. Do Not Update Data Buffers Mid-Frame

Avoid touching any buffer when a frame is mid-flight to reduce stalls and temporary buffer stores.

6. Do Use Texture Compression

Reduce the memory footprint and bandwidth cost of the texture assets.

7. Do Use Mipmapping

This increases texture cache efficiency, which reduces bandwidth and increases performance.

8. Do Not Use Discard

Avoid forcing depth-test processing in the texture stage as this will decrease performance in the early depth rejection architectures.

9. Do Not Force Unnecessary Synchronisation

Avoid API functionality that could stall the graphics pipeline and do not access any hardware buffer directly.

10. Do Move Calculations 'Up the Chain'

Reduce the overall number of calculations by moving them earlier in the pipeline, where there are fewer instances to process.

1. Do Understand Your Target Device

Seek to learn as much information about your target platforms as possible in order to understand different graphics architectures, so that you use the device in the most efficient manner possible.

Manufacturers' websites for devices are a good place to look for specifications and they may also provide other helpful developer community resources. The PowerVR Graphics SDK provides public architecture and performance recommendation documents for reference:

- PowerVR Hardware Architecture Overview for Developers
- PowerVR Series5 Architecture Guide for Developers
- PowerVR Performance Recommendations
- PowerVR Low Level GLSL Optimisation
- PowerVR Instruction Set Reference

Further PowerVR architecture documentation is available from us under a non-disclosure agreement.

Even after the graphics architecture is thoroughly understood it is important to remember that other factors such as variations in CPU processing power, memory bandwidth and thermal load will also impact an applications performance.

2. Do Profile Your Application

Identify the bottlenecks in your application and determine whether there are opportunities for improvement.

It is important to understand where performance is bottlenecked before attempting to optimise an application. This ensures effort isn't wasted or visual quality isn't sacrificed for minimal gains. If an optimisation is inappropriately applied to an area that isn't bottlenecking performance, there may be no performance improvements and in some cases, an incorrectly applied optimisation may lead to worse performance.

From the PowerVR Developer Technology team's experience, we have derived the following list of common bottlenecks generally found in applications that have not been optimised, as ordered from most to least common:

- CPU Usage
- Bandwidth Usage
- CPU/Graphics Core Synchronisation
- Fragment Shader Instructions
- Geometry Upload
- Texture Upload
- Vertex Shader Instructions
- Geometry Complexity

Profiling tools are vital in this process for developers to understand what is happening in their application, the hardware it is running on, and how and where bottlenecks are occurring. The PowerVR SDK includes the profiling tools PVRTrace and PVRTune to aid development on platforms powered by PowerVR hardware.

3. Do Not Use Alpha Blend Unnecessarily

Be sure Alpha Blending is used only when required to make the most of deferred architectures and to save bandwidth.

Disable alpha blending wherever possible and if transparent objects are required keep the number of transparent objects to a minimum. The reasoning behind this is that deferred renderers such as PowerVR graphics cores calculate the visibility of fragments before the corresponding fragment shader is invoked to process it. This prevents invisible fragments in the output image being processed unnecessarily.

If alpha blending is enabled, then the hardware used to determine a fragment's visibility cannot be utilised. This is because the occluded (alpha-blended) fragment may impact the final rendered image. Due to this behaviour, enabling alpha blending eliminates the benefits of deferred rendering graphics architectures. This means the hardware is no longer able to make decisions about a fragment's visibility and drop it from the pipeline. This will likely result in overdraw which is where fragments are being processed that are not actually visible in the final image. Overdraw can negatively impact the application's performance, especially if the application is already limited by rendering.

4. Do Perform Clear

Perform a clear on a framebuffer's contents to avoid fetching the previous frame's data on tile-based graphics architectures, which reduces memory bandwidth.

System memory accesses use more bandwidth and power than any other graphics operation. Keeping memory accesses to a minimum will reduce the chances of your application being memory bandwidth bound and will also reduce the power consumption of your application.

Most applications need to generate a colour image at the end of the render, but have no need to preserve depth and stencil data between frames. Therefore if framebuffer attachments do not need to be preserved at the end of a render you should invalidate the appropriate framebuffer attachments to prevent them being written out to system memory.

Even fewer applications have a genuine need to upload the contents of the colour buffer's previous contents at the start of a new frame. Therefore if the contents previously written to framebuffer are not required, you can inform the driver not to load them from system memory to on-chip tile memory through a *clear* operation at the start of the render.

The net result of performing a clear and invalidating framebuffers will be a massive reduction system memory bandwidth usage and reduced power consumption.

In OpenGL ES you can perform a clear by calling the `glClear` function at the beginning of a render. Additionally you may use the `glDiscardFramebufferEXT` or `glInvalidateFramebuffer` functions to invalidate a framebuffer at the end of a render.

In Vulkan the API gives you explicit control over load and store operations on framebuffer attachments. When you are creating a framebuffer, you should set the load operation to either `VK_ATTACHMENT_LOAD_OP_DONT_CARE` or `VK_ATTACHMENT_LOAD_OP_CLEAR`. The store operation should preferably be set to `VK_ATTACHMENT_STORE_OP_DONT_CARE` unless you require the data to be preserved.

5. Do Not Update Data Buffers Mid-Frame

Avoid touching any buffer when a frame is mid-flight to reduce stalls and temporary buffer stores.

Modifying in-flight resources currently in use by the GPU such as vertex buffers and textures has a significant cost. As graphics processors tend to have at least one frame of latency to ensure that the hardware is always well occupied with work, altering a resource required by an outstanding render will usually result in one of the following actions being taken:

1. Stall in the buffer modifying API call until the outstanding render completes.
2. Allocate a temporary buffer for the new data so the buffer modifying API call can complete without stalling the CPU thread.




Figure 1. Using the CPU to modify the GPU buffers can be costly

As textures are generally accessed during fragment shading much later in the graphics pipeline than vertex attributes, the cost of a graphics driver stalling a texture modification is higher than modifying a vertex buffer. The driver may choose to avoid a stall entirely by creating temporary buffer stores (a.k.a. ghosting) which is good for performance, but it may not be desirable for applications that are already running out of buffer storage space.

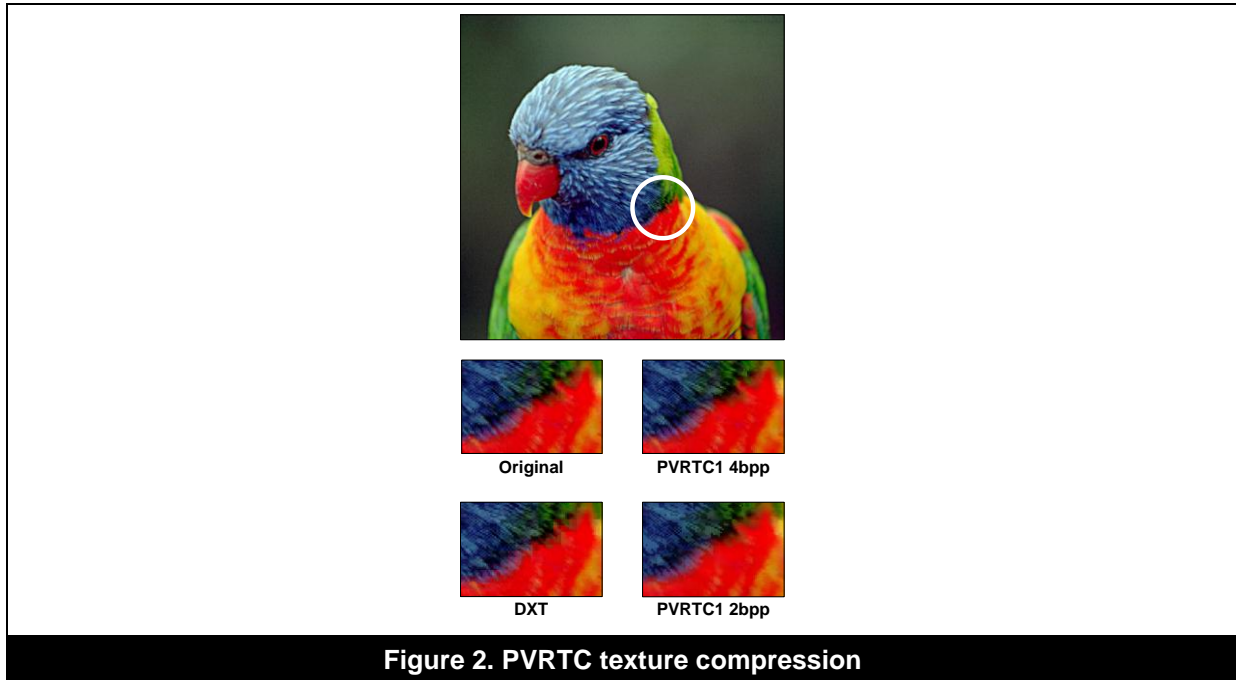
The stalling and ghosting behaviour of graphics processors varies between different GPUs and driver versions. For optimal performance, you should only modify vertex buffers and textures when absolutely necessary. If buffers must be modified, you should use application-side circular buffering so that the graphics processor can read from one buffer object while the application's CPU thread writes to another, thus preventing the stalling and ghosting behaviours.

If your application is using the Vulkan graphics API, then it is the responsibility of the application developer to synchronise with the graphics processor. The developer needs to put the appropriate mechanisms such as fences and semaphores in place to ensure that the application does not access a resource while the graphics processor is using it. This means that the developer has much more control over how and when resources are accessed, but comes at the cost of a more complex application as the driver will not safeguard against accessing data currently in use by the graphics processor.

6. Do Use Texture Compression

Reduce the memory footprint and bandwidth cost of the texture assets.

In some instances, it is worth considering the balance between texture size and texture compression. It may be possible to use a larger texture and a low-bitrate compression scheme and achieve a better balance of bandwidth savings and acceptable image quality.



Texture compression, not to be confused with image file compression, minimises the runtime memory footprint of your textures. This provides several performance benefits, but primarily reducing the amount of system memory bandwidth consumed sending data to the Graphics Core. PVRTC and PVRTCII are PowerVR specific compression technologies and will achieve best performance on the hardware, consuming as little as 2 bits per pixel. These textures are also very texture cache efficient as the lower pixel size allows a more pixels to fit in the limited amount of cache memory available to the texture units.

Depending on the PowerVR generation and graphics API you are targeting, additional compressed texture formats may be supported, such as ASTC.

7. Do Use Mipmapping

This increases texture cache efficiency, which reduces bandwidth and increases performance.

Mipmaps are smaller, pre-filtered variants of a texture image, representing different levels of detail of a texture. By using a minification filter mode that uses mipmaps, the graphics core can be set up to automatically calculate which level of detail comes closest to mapping the texels of a mipmap to pixels in the render target, and use the right mipmap for texturing.

Using mipmaps has two important advantages:

1. It increases graphics rendering performance by massively improving texture cache efficiency, especially in cases of strong minification – the texture data is more likely to fit inside tile memory.
2. It also improves image quality by reducing aliasing that is caused by the under sampling of textures that do not use mipmapping.

The single limitation of mipmapping is that it requires approximately a third more texture memory per image. Depending on the situation, this cost may be minor when compared to the benefits in terms of rendering speed and image quality.

There are some exceptions where mipmaps should not be used. For example, mipmapping should not be used where filtering cannot be applied sensibly, such as for textures that contain non-image data such as indices or depth textures. It should also be avoided for textures that are never minified, for example, UI elements where texels are always mapped one to one to pixels.

Ideally mipmaps should be created offline using a tool like PVRTexTool, which is available as part of the PowerVR Graphics SDK.

It is possible to generate mipmaps at runtime, which can be useful for updating the mipmaps for a render to texture target. In OpenGL ES this can be achieved using the function `glGenerateMipmap`. In Vulkan there is no such built in function and a developer must generate mipmaps manually.

Also keep in mind that generation of mipmaps online will not work with compressed textures such as PVRTC, which must have their mipmaps generated offline. A decision must be made as to which cost is the most appropriate: the storage cost of offline generation, or the runtime cost (and increased code complexity in the case of Vulkan) of generating mipmaps at runtime.

8. Do Not Use Discard

Avoid forcing depth-test processing in the texture stage as this will decrease performance in the early depth rejection architectures.

Applications should avoid the use of the discard operation in the fragment shader as using it will not improve performance. Most mobile graphics cores use a form of tile based deferred rendering (TBDR) and using discard negates some of the benefits of this type of architecture. If possible an application should prefer alpha blending over discarding.

Additionally applications should avoid alpha testing. When an alpha tested primitive is submitted, early depth testing, such as PowerVR's Hidden Surface Removal (HSR), can discard fragments that are occluded by other fragments closer to the camera. Unlike opaque primitives that would also perform depth writes at this pipeline stage, alpha tested primitives cannot write data to the depth buffer until the fragment shader has executed and fragment visibility is known. These deferred depth writes can impact performance, as subsequent primitives cannot be processed until the depth buffers are updated with the alpha tested primitive's values.




Figure 3. PowerVR alpha test/discard pipeline

For optimal performance, consider alpha blending instead of alpha test to avoid costly deferred depth write operations. To ensure HSR removes as much overdraw as possible, you should also submit your draws in the following order:

1. opaque
2. alpha-tested
3. blended

9. Do Not Force Unnecessary Synchronisation

Avoid API functionality that could stall the graphics pipeline, and do not access any hardware buffer directly.

Graphics applications achieve the best performance when the CPU and graphics core tasks run in parallel. Graphics cores also operate most efficiently when the vertex processing tasks of one frame are processed in parallel to the fragment colouring tasks of previous frames. When an application issues a command that causes the CPU to interrupt the graphics core, it can significantly reduce performance.

The most efficient way for the hardware to schedule tasks is vertex processing executing in parallel to fragment tasks. In order to achieve this, the application should aim to remove functions which cause synchronisation between the CPU and graphics core wherever possible. In OpenGL ES this includes synchronisation functions such as `glReadPixels`, `glFinish`, `eglClientWaitSync` and `glWaitSync`. In Vulkan the application developer has much finer control over synchronisation between resources as any synchronisation between the graphics processor and CPU is defined by the developer.

In addition one of the most common causes of poor application performance is applications accessing the contents of a framebuffer from the CPU. When such an operation is issued, the calling application's CPU thread must stall until the graphics core has finished rendering into the framebuffer attachment. Once the render is complete, the CPU can begin reading data from the attachment. During this time the graphics core will not have write access to that attachment which can cause the graphics core to stall subsequent renders to that framebuffer. Due to the severe cost these operations should only be used when absolutely necessary, for example to capture a screenshot of a game when a player requests one.

10. Do Move Calculations ‘Up the Chain’

Reduce the overall number of calculations by moving them earlier in the pipeline where there are fewer instances to process.

By performing calculations earlier in the pipeline, the overall number of operations can be reduced and therefore workload can be substantially reduced. Generally in a scene there are far fewer vertices than fragments that need to be processed. Therefore processing per vertex instead of per fragment would greatly reduce the number of calculations. One use case, for example, could be to perform per vertex lighting instead of per pixel lighting.




Figure 4. Promote calculations up the chain to reduce the number of times they are executed

It is also possible to consider moving calculations off the graphics core altogether. Although the graphics core may be able to perform operations far more rapidly than the CPU can, it would be even faster for the CPU to perform an operation just once instead of allowing the operation to be performed for many vertices on the graphics core.

To take the concept even further, consider performing calculations offline by baking values into the scene, effectively replacing expensive run-time calculations with a simple lookup. For example, replacing real-time lighting with light maps for static objects in a scene, such as terrain, buildings and trees can be a particularly effective compromise. This substantially improves performance and in many cases provides higher quality lighting than would be possible to calculate at run-time.

Other Considerations

Do Group Per Material

Modifying the GL state machine incurs CPU overhead in the graphics driver, as changes need to be interpreted and converted into tasks that can be issued to the graphics core. To reduce this overhead, you must minimise the number of API calls and state changes made by your application.

For geometry data, combine as many meshes into a single draw call as possible. Here is an example use case:

Meshes for seats on a train have static position and orientation relative to one another and use the same render state. The seats and the train could all be combined into a single mesh. To draw the train interior, several draw calls have merged into a single call.

With the Hidden Surface Removal (HSR) feature on PowerVR hardware it is not necessary to submit geometry in depth order to reduce overdraw. By freeing applications from this restriction they can focus on sorting draws by render state, ensuring state changes are minimised.

Similar to geometry data, it is possible to combine several textures into a single bindable object by using texture atlases or texture arrays where available. Textures can then be applied per object with the appropriate shader uniforms.

As discussed in Avoid Read/Write of In-Use Buffer Objects, modifying buffer data may stall the graphics pipeline or increase the amount of memory allocated by the graphics driver. When batching draws together, it is important to consider the update frequency of buffers. For example, you should batch spatially coherent objects with static vertex data into one vertex buffer, and objects with dynamic data, such as soft body objects like cloth, into another.

Do Use Indexed List

Vertex buffers enable the graphics driver to cache vertex data attributes, such as texture coordinates for mapping 2D images to the mesh, and model/space position. For static objects which have vertex attributes that change infrequently if at all, vertex buffers improve performance as the cached data can be reused to render many frames.

In the example above, an index buffer is used in conjunction with a vertex buffer. Index buffers define the order in which elements of a vertex buffer should be accessed to represent the triangles in a mesh. Index buffers improve performance and reduce the storage space requirements of complex mesh data. This is because vertex attributes are written to the vertex buffer once, then referenced as many times as required to represent the triangles surrounding that vertex position.

PowerVR hardware is optimised for indexed triangle lists. For finely-tuned performance, vertex and index buffers should be sorted to improve cache efficiency when the data is accessed by the GPU. Our 3D scene exporter and converter tool, PVRGeoPOD, automatically applies sorting to mesh data when generating POD (PowerVR object data) files.

Do Utilise All Your CPU Cores

Modern mobile devices usually have more than a single CPU core. To achieve the best performance possible on modern CPU architectures it is crucial that your application utilises multi-threading wherever possible. For example you may have your graphics updates on the main thread while having physics updates running on a separate worker thread. Splitting large chunks of work such as physics, animations, file I/O and so on over multiple threads enables your application to utilise the CPU more efficiently and will usually result in a smoother end user experience. Additionally if your application is targeting the Vulkan graphics API, you may be able to split preparation of draw commands (building command buffers) over several threads.

Do Prefer Lower Data Precision

Variables in shaders declared with the mediump modifier are represented as 16 bit floating point (FP16) values. It is advised that an application uses FP16 wherever appropriate as it typically offers a significant performance improvement by theoretically double the floating point throughput over FP32 (highp). It should be considered wherever FP32 would normally be used, provided the precision is sufficient and the maximum and minimum values will not be overflowed, as visual artefacts may be introduced.

Do Use Level of Detail (LoD)

Level of detail is an important consideration for an application, the concept of 'good enough' should be employed here; developers should consider the usage of expensive graphics effects and high quality assets against the impact to performance.

Mipmapping is one form of LoD, which we have previously discussed in this document with Rule 7.

A second consideration for LoD is geometry complexity. An appropriate level of geometry complexity should be used for each object or portion of an object. The following are examples of a waste of compute and memory resources:

- using a large number of polygons for an object that will never cover more than a small area of the screen, like a distant background object.
- using polygons for detail that will never be seen due to camera angle, or culling (i.e. objects outside of the view frustum)
- using large numbers of primitives for objects that can be drawn with much fewer, with minimal to no loss in visual fidelity. As an example - using many hundreds of polygons to render a single quad.

Shader techniques such as bump mapping should be considered to minimise geometry complexity, but still maintain a high level of perceived detail. This is especially true for techniques such as reflection passes, where higher amounts of geometry may not be visible.

Do Not Use Depth Pre-pass

On graphics hardware that employs a deferred rendering architecture such as PowerVR, an application should not perform a depth pre-pass as there is no performance benefit. Performing this operation would be a waste of clock cycles and memory bandwidth. This is because the hardware will detect and remove occluded (opaque) geometry from the pipeline automatically during rasterization, before fragment processing begins.

Do Use On-Chip Memory Efficiently for Deferred Rendering

Graphics techniques such as deferred lighting are often implemented by attaching multiple colour render targets to a frame buffer object, rendering the required intermediate data, and then sampling from this data as textures. While flexible, this approach even when implemented optimally still consumes a large amount of system memory bandwidth, which comes at a premium on mobile devices.

Both OpenGL ES (3.x) and Vulkan graphics APIs provide a method to enable communication between fragment shader invocations which cover the same pixel location – through intermediate on-chip buffers. This buffer can only be read from and written to by shader invocations at the same pixel coordinate.

The GLES extension *shader_pixel_local_storage(2)* and Vulkan transient attachments enables applications to store the intermediate per-pixel data in on-chip tile memory. While each method has their own implementation details, they both provide similar functionality and both bring the same benefits. For example the "G-Buffer" attachments in a deferred lighting pass that are only needed once can be stored in tile memory and then completely discarded when drawing is complete.

Both of the API features described above are extremely beneficial for tile based renderers such as PowerVR graphics cores, as it allows an application to efficiently make use of on-chip tile memory. The intermediate framebuffer attachments are never allocated or written out to system memory - they only exist in on-chip tile memory. This is extremely beneficial for mobile and embedded systems where memory bandwidth is at a premium. Using these features correctly will result in a significant reduction in system memory bandwidth usage. Additionally most techniques (such as deferred lighting) that write intermediate data out to system memory and then sample from it at the same pixel location can be optimised using these API features.

Do Prefer Explicit APIs

Vulkan is a new generation graphics and compute API. It is highly efficient, streamlined and modern and designed to take advantage of current and future device architectures. Vulkan works on a wide variety of platforms such as desktop PCs, consoles, mobile devices and embedded devices.

Vulkan is designed from the ground up to take advantage of modern CPU architecture such as multi-core and multi-threaded systems, rendering work can be spread over many logical threads. The Vulkan "Gnome Horde" demo in the PowerVR SDK shows this aspect of the API very nicely.

Vulkan is designed to have minimal driver overhead, but this comes at the cost of a more complex programming paradigm – explicit. In Vulkan it is up to the application developer to handle low level details such as memory allocation for buffers and explicit synchronisation between resources. However once you have mastered the API, a Vulkan graphics application is likely to run much more efficiently and more predictably across various devices compared to legacy graphics APIs.

Contact Details

For further support, visit our forum:

<http://forum.imgtec.com>

Or file a ticket in our support system:

<https://pvrsupport.imgtec.com>

To learn more about our PowerVR Graphics SDK and Insider programme, please visit:

<http://www.powervrinsider.com>

For general enquiries, please visit our website:

<http://imgtec.com/corporate/contactus.asp>