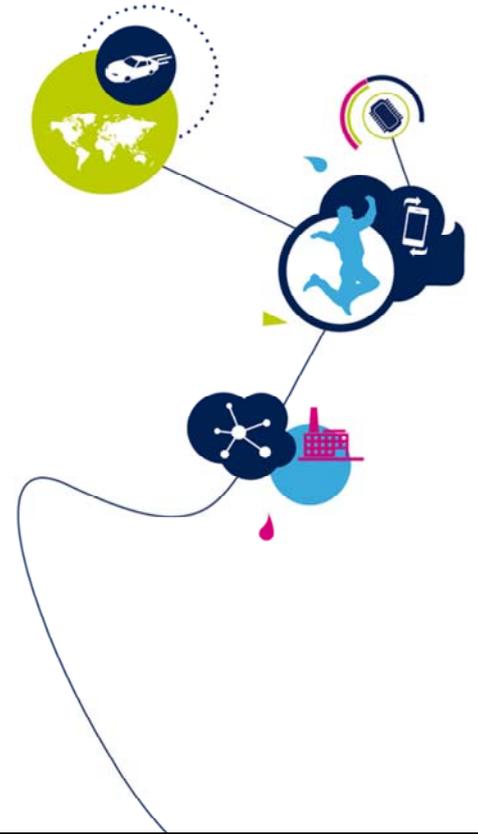


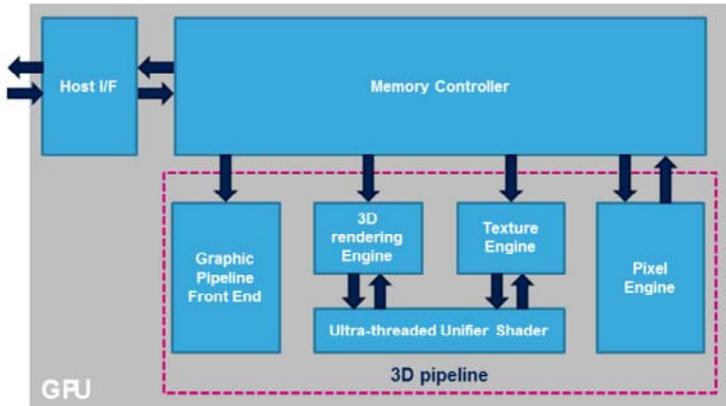
STM32MP1 – GPU

Graphics Processing Unit
Revision 0.1



Hello, and welcome to this presentation of the Graphics Processing Unit (GPU) embedded in the STM32MP1 microprocessor family. It covers the features of this processor dedicated to graphics creation, which is widely used for the rendering of Human Man Interface (HMI) and multimedia content.

Best-in-class Graphic Acceleration for HMI



- The GPU is build around a 3D pipeline with a unified shader
 - Specific 3D rendering engine and texture engine
 - Dedicated memory controller for transfers from/to the system and/or external memories

Application benefits

- 3D Graphics for HMI application
- State of the art graphic effects
- Standard API support (OpenGLES)



The STM32MP1 GPU is a powerful 3D Graphic processing unit. Its 3D pipeline is build around a Ultra-Threaded Unified Shader with a specific 3D rendering engine and a texture engine to efficiently render 3D scenes.

A dedicated memory controller optimizes all the transfers from and to the system or external memories.

Flexible operating modes to reduce CPU load

- OpenGL ES 2.0 / 1.1 compliance, including extensions; OpenVG 1.1
- IEEE 32-bit floating-point pipeline
- Ultra-threaded, unified vertex and fragment (pixel) shaders
- Low bandwidth at both high and low data rates & Low CPU load
- Up to 12 programmable elements per vertex
- Texture dependent operation with high-performance
- Alpha blending
- Depth and stencil compare
- Support for 8 fragment shader simultaneous textures
- Support for 4 vertex shader simultaneous textures
- Point sampling, bi-linear sampling, tri-linear filtering, and cubic textures
- 8k x 8k texture size and 8k x 8k rendering target
- 4 Vertex DMA streams



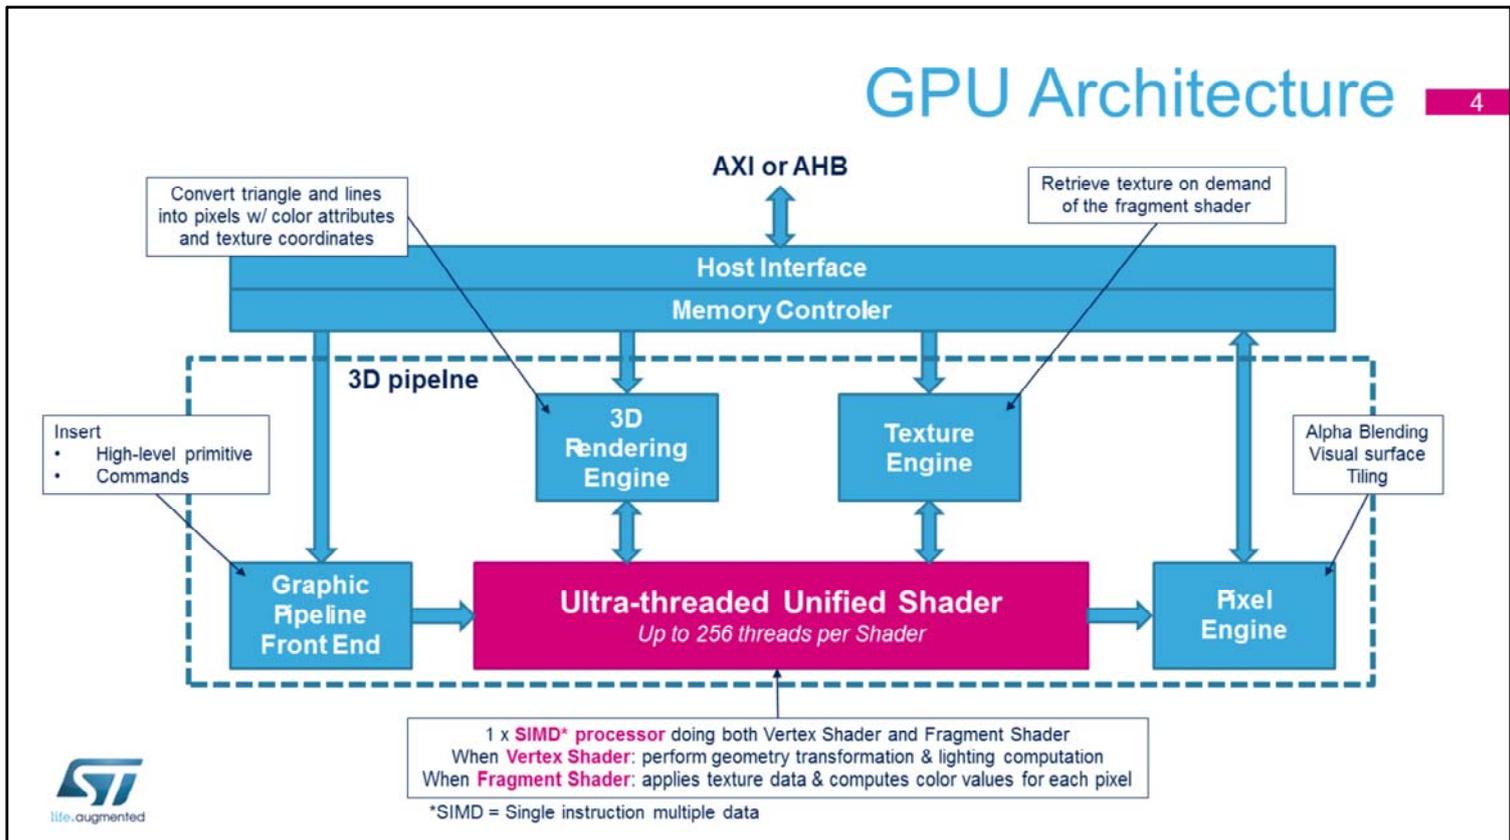
The STM32MP1 GPU is compliant with OpenGL ES 1.1 and 2.0 and OpenVG 1.1 standards.

The GPU build around a powerful ultra-threaded unified sharder working with IEEE compliant 32-bit floating-point numbers and accelerates all the state of the art graphic functions lowering the CPU load.

Each vertex can have up to 12 programmable elements and the GPU supports simultaneous textures for 8 fragment shaders and 4 vertex shaders.

Several sampling techniques are supported like, point sampling, bilinear, trilinear or cubic.

The rendering target and texture size can go up to 8k per 8k pixels.



The GPU has key dedicated block to provide the best in class graphic acceleration:

- The graphic pipeline front end takes as an input the high-level graphic primitives and commands.
- They are transmitted to the Ultra-threaded Unified Shader, which is a single **instruction multiple data** (SIMD) processor supporting up to 256 threads.
- The unified shader acts as a vertex shader for geometry transformation and lighting computation or as a fragment shader to apply textures and compute the colors of each pixels.
- A specific 3D rendering engine converts triangles and lines into pixels with color attributes and texture coordinates.
- The texture engine ensures the retrieval of the textures on demand of the fragment shader.
- The pixel engine performs the final alpha blending and the tiling.

Interrupt event	Description
Gbbal interrupt	Set GPU event occurs.



The GPU has only one global interrupt. This interrupt is used by the low level drivers for graphic API implementation.

Mode	Description
Run	Active.
Sleep	Active. Peripheral interrupts cause the device to exit Sleep mode.
Stop + LP-Stop	Frozen. Peripheral registers content is kept.
FLV-Stop	Frozen. Peripheral registers content is kept.
Standby	Powered-down. The peripheral must be reinitialized after exiting Standby mode.



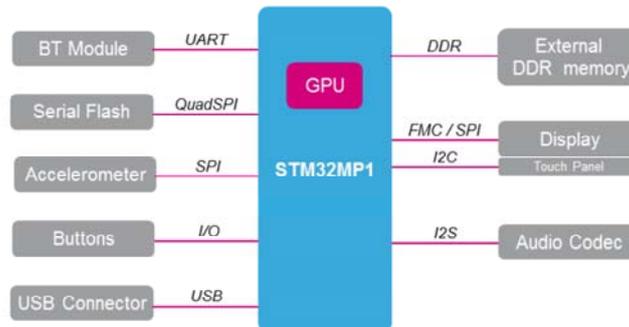
The GPU is active in Run and Sleep modes. A GPU interrupt can cause the device to exit Sleep mode. In Stop modes, the GPU is frozen and its registers content is kept. In Standby mode, the GPU is powered-down and it must be reinitialized afterwards.

.

Application examples

7

- Low-power embedded applications including connectivity and user interface:



- GPU can process all the graphic primitives for advanced graphical user interface rendering.



Embedded applications require low-power management functions together with a high-quality user interface. This can be achieved using the GPU to render advanced graphical user interfaces.

- Refer to these trainings related to this peripheral:
 - RCC (CPU clock control, CPU enable/reset)
 - Interrupts (CPU interrupt mapping)

You can refer to the training slides related to the RCC block and interrupts for additional information.

Comparison across STM32 series

9

GPU features	STM32 L4	STM32 F0	STM32 F1	STM32 F2	STM32 F3	STM32 F4	STM32 F7	STM32 H7	STM32 MP1
Number of instances	0	0	0	0	0	0	0	0	1
Maximum speed	N/A	533 MHz							
OpenGL ES 1.1	N/A	Y							
OpenGL ES 2.0	N/A	Y							
OpenVG 1.1	N/A	Y							
EGL 1.4	N/A	Y							
OpenGL 2.1	N/A	Y							



The GPU is available on STM32MP1 devices with OpenGL ES 1.1 & 2.0 OpenVG 1.1, EGL 1.4 and OpenGL 2.1 API support.