

10.5 An SoC with 1.3Gtexels/s 3D Graphics Full Pipeline Engine for Consumer Applications

Donghyun Kim¹, Kyusik Chung¹, Chang-Hyo Yu¹, Chun-Ho Kim¹, Inho Lee¹, Jaewan Bae¹, Young-Jun Kim², YoungJin Chung², Sungbeen Kim², Yong-Ha Park², NakHee Seong², Jin-Aeon Lee², Jaehong Park², Stephen Oh², Seh-Woong Jeong², Lee-Sup Kim¹

¹KAIST, Daejeon, Korea

²Samsung Electronics, Giheung, Korea

Recently, embedded 3D graphics hardware for portable multimedia products and home AV electronics has become popular. Several embedded 3D graphics engines have been reported [1,2], but better performance and functionality are required for higher display resolution and more realistic image quality in today's consumer applications. In this paper, an SoC for consumer electronics such as game consoles, DTVs, next generation DVD players is presented. It contains a 3D graphics full pipeline engine to provide rendering performance up to 33Mvertices/s and 1.3Gtexels/s. Since minimizing external memory accesses and hiding SDRAM access overheads are key factors to achieve high rendering performance, two techniques, *A-index* two-level texture cache and *depth filter*, are implemented to reduce the external memory accesses by 85.2% [4,5]. In addition, the bus arbiter and memory interface units (MIU) are designed to hide SDRAM access overheads efficiently.

The SoC provides three major features to achieve better visual effects. First, the SoC fully supports the OpenGL ES which is a standard API for embedded 3D graphics. Supporting the standard 3D graphics API makes it easy and affordable to offer a variety of advanced 3D graphics and games migrated from other platforms. Second, the SoC has the enhanced programmability to facilitate advanced 3D graphics algorithms, such as bidirectional reflectance distribution function, non-linear displacement, fisheye-lens effect, and so on. Third, the SoC contains a video composition subsystem to overlay 3D graphics images over a video layer or to use dynamic video streaming data as texture primitives.

Figure 10.5.1 shows a block diagram of the developed SoC. The major key features mentioned above are inherited from integrated 3D graphics IP and video-related IPs. The SoC also integrates an ARM11 RISC core useful as a stand-alone main processor and a video decoder. The floating-point co-processor in ARM11 RISC enables collision detection, artificial intelligence, and other CPU-level algorithms required in game applications. Each IP communicates with other IPs and MIUs through a dedicated 4-layer 32b bus which is backward compatible with AMBA AHB bus system. The order of SDRAM accesses is determined by the bus arbiter in a well-designed scheduling scheme concerning interleaved accesses on different banks and consecutive accesses in a row. clock generator that provides various clocks from three PLLs, controls power management of each IP by turning on/off a corresponding clock with configurable options.

The integrated 3D graphics IP consists of bus interface units and two functional units: geometry engine (GE) and rasterization engine (RE). Figure 10.5.2 depicts the block diagram of the GE. It consists of 2 floating-point programmable engines, vertex shading engine (VSE) and triangle setup engine (TSE). VSE has a 4-way 32b floating-point SIMD unit to accelerate transformation of 4-dimensional vectors. Four parallel multipliers and a cascaded 4-input adder evaluate an inner product of 4-dimensional vectors in one instruction. Special function unit (SFU) is

used to evaluate exponent, logarithm, division, and reciprocal-square-root for lighting calculations and special effects. Write-after-write (WAW) hazards by different write-back timing and the data dependency between instructions are checked and eliminated by a hazard controller. The architecture of TSE is similar to VSE, but it has a 3-way 32b floating-point SIMD unit and a dedicated pixel center adjustment (PCT) logic. Format converter (FC) scales floating-point data to appropriate range for fixed-point format used in RE with minimized precision errors.

As shown in Fig. 10.5.3, there are 3 major functional blocks in RE: rasterizer, texturing engine (TE), and per-pixel processing engine (PPE). Rasterizer has 4 parallel pixel generators to produce a 2x2 pixel stamp per clock cycle. The texture coordinates are perspective-correctly evaluated in division-free methods [3]. TE processes texture mapping on 4 pixels simultaneously and supports 2-level multi-texturing as well as trilinear multum in parvo map (MIPMAP) texture mapping. PPE performs fog-effect, color blending, alpha test, depth test, and stencil test. The 1kB Z-line cache and the C-line cache assemble depth and color data of the consecutive pixels in the screen coordinate. They store 8 stamp-lines and a stamp-line consists of two 16-pixel-lines. When a stamp-line is replaced, upper and lower pixel-lines are concurrently processed by two bus masters for efficient bus interleaving.

RE has 2-level texture cache that employs *A-index*, an adaptive method to reduce the cache miss by selecting cache indices based on the data-access characteristics of TE [5]. *A-index* 2-level texture cache pulls down the total external memory accesses by 65.6%. *Depth filter* technique is also adopted to remove invisible pixels ahead of per-pixel pipeline [4]. It determines if a pixel is hidden by a certain mask plane having the history of appearances of a pixel in front of the mask plane. Since 3-plane *depth filter* needs 2b per pixel, it requires 2Mb of internal memory for 1024x1024 display resolution. To avoid using large internal memory, a fully associative 256B *depth filter* cache is used. As described in Fig. 10.5.4, the cache consists of 8 lines that stores *depth filter* values of 32 pixel-stamps (128 pixels). *Depth filter* reduces the number of pixels passed to the next stage up to 71.7% and saves 58.7% of total memory bandwidth required by the 3D graphics engine. As shown in Fig. 10.5.5, huge memory bandwidth up to 7.1GB/s is required to achieve 1.3Gtexels/s texture performance, but the required memory bandwidth is drastically reduced from 7.1GB/s to 1.05GB/s so that 3D graphics IP is capable of utilizing its full performance within 1.3GB/s bandwidth provided by SDRAM.

The implemented SoC integrates 17.9M transistors in 7.1x7.0mm² die size using 1.2V 0.13μm CMOS technology. ARM11 RISC processor runs at 333MHz while 3D graphics and video IPs run at 166MHz. Figure 10.5.6 summarizes features and characteristics of the SoC. The chip micrograph is shown in Fig. 10.5.7.

References:

- [1] R. Woo et al., "A 210mW Graphics LSI Implementing Full 3D Pipeline with 264Mtexels/s Texturing for Mobile Multimedia Applications," *ISSCC Dig. Tech. Papers*, pp. 44-45, Feb., 2003.
- [2] M. Imai et al., "A 109.5mW 1.2V 600M texels/s 3D Graphics Engine," *ISSCC Dig. Tech. Papers*, pp. 332-333, Feb., 2004.
- [3] D. Kim and L. S. Kim, "Division-Free Rasterizer for Perspective-Correct Texture Filtering," *IEEE ISCAS*, pp. 153-156, May, 2004.
- [4] C. H. Yu and L. S. Kim, "An Adaptive Spatial Filter for Early Depth Test," *IEEE ISCAS*, pp. 137-140, May, 2004.
- [5] C.H. Kim, L.S. Kim, "Adaptive Selection of an Index in a Texture Cache," *International Conference on Computer Design*, pp. 295-300, Oct., 2004.

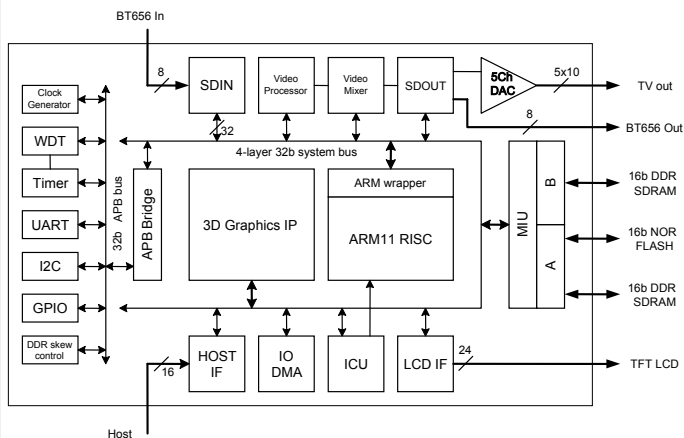


Figure 10.5.1: Block diagram of the SoC.

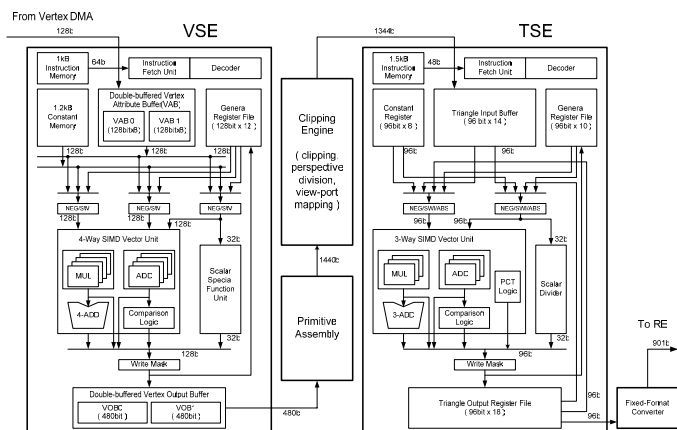


Figure 10.5.2: VSE and TSE in GE.

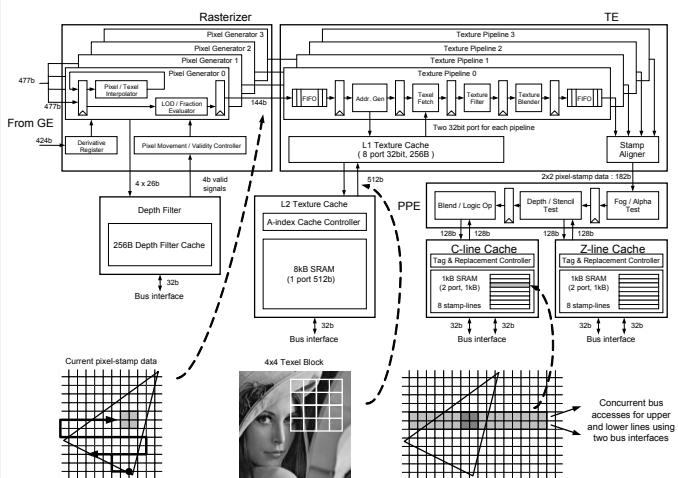


Figure 10.5.3: RE architecture.

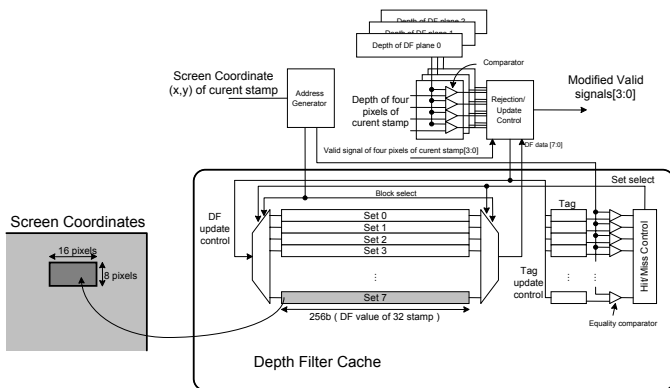


Figure 10.5.4: Depth Filter architecture.




Figure 10.5.5: Memory bandwidth reduction by the techniques used in 3-D graphics IP.

Process Technology	0.13μm CMOS 7-Metal
Supply Voltage	1.2V
Operating Frequency	166MHz* for IPs and 333MHz for RISC Core
Power Consumption	407mW for fully activated 3D Graphics IP
Transistor Counts	17.9M transistors (7.0M for 3D Graphics IP)
Die Size	7.1x7.0mm ²
Package	256 pin QFPA
3D Performance	33Mvertices/s for GE, 1.3Gtexels/s for RE
3D Function Feature	All OpenGL ES operations Programmable vertex and triangle processing 2-level multi-texturing
Video Interface	BT.656 video input / output TFT-LCD output Component output NTSC/PAL, S-Video Composite output

* The system clock is 117MHz in the worst condition which is 1.1V, 125°C, and the worst process variation.

Figure 10.5.6: Chip specifications and characteristics.

Continued on Page 592

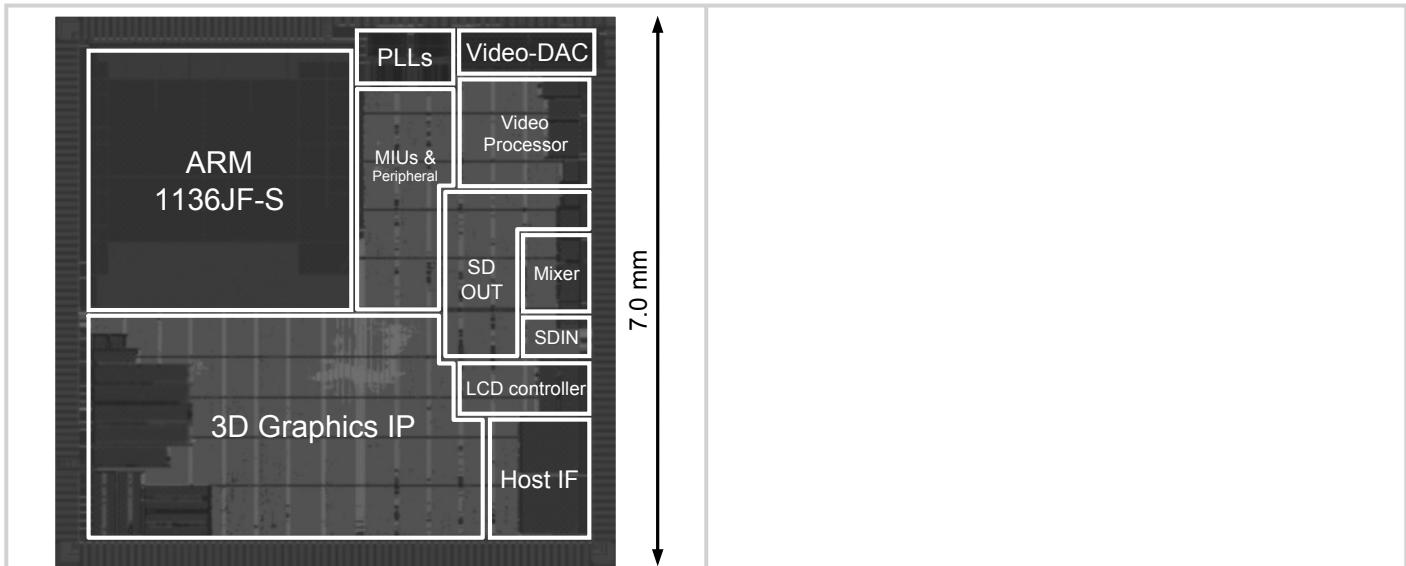


Figure 10.5.7: Chip micrograph.