2.4 A 210mW Graphics LSI Implementing Full 3D Pipeline with 264Mtexels/s Texturing for Mobile Multimedia Applications

Ramchan Woo, Sungdae Choi, Ju-Ho Sohn, Seong-Jun Song, Young-Don Bae, Chi-Weon Yoon, Byeong-Gyu Nam, Jeong-Ho Woo, Sung-Eun Kim, In-Cheol Park, Sungwon Shin[1], Kyung-Dong Yoo[1], Jin-Yong Chung[1], Hoi-Jun Yoo

KAIST, Daejeon, Republic of Korea
[1]Hynix Semiconductor, Icheon, Republic of Korea

As mobile applications are moving to realtime multimedia, more functions are integrated into handheld devices [1-4]. Unlike the realtime decoding of MPEG-4 video [1,4], the performances of hardware-accelerated 3D solutions designed for mobile platforms are still below the market demands showing only limited shading operations [2,3]. Since the realization of various 3D functions requires huge computing power and corresponding memory bandwidth, previous LSIs integrate DRAM using the embedded memory logic technology although it is cost-inefficient due to process complexity. In this work, a graphics LSI using the pure DRAM technology is implemented to integrate both the logic and memory at low cost. Its circuits and architecture are optimized so that the full 3D pipeline is realized with less than 210mW at the drawing speed of 264Mtexels/s bilinear MIPMAP texturing and antialiasing, applicable to handheld devices.

As shown in Fig. 2.4.1, the graphics LSI consists of a 32b RISC, 3D rendering engine (3DRE), 29Mb DRAM, bandwidth equalizer (BEQ) and programmable power optimizer (PPO). The ARM-9 compatible RISC with 4kB I/D caches operates to 132MHz [6]. The RISC includes a 32 x 32b MAC in its datapath to accelerate the 3D geometry operations so that it can calculate as many as 1.04Mvertices/s transformation running a customized fixed-point graphics library, a 43% improvement over the conventional ARM9 processor [5]. SlimShader performs the main rendering operations such as texturing, shading, blending, and depth comparison. Memory Programmer (MP) enables the special effects such as antialiasing, motion blur and fog to be programmable. Integrated 29Mb DRAM with partial wordline scheme [1-3] provides sufficient bandwidth and capacity required for 3D rendering operations. Dedicated hardware engines and 1.6GB/s bandwidth through 416b-wide DRAM lowers the operation frequency of 3DRE even to 33MHz. To compensate the difference of the processing speed and data width between the RISC and 3DRE, BEQ buffers the vertex data with 1kB Dual-Ported SRAM (DP-SRAM). BEQ partially activates the banks of DP-SRAM according to the required buffer size saving 20% power of DP-SRAM. For DSP applications, DP-SRAM can also be used by the RISC as a scratch-pad RAM. PPO reduces the power consumption of the chip by varying four different clock domains.

3DRE is shown in Fig. 2.4.2. The triangle setup engine (TSE), which contains single-cycle parallel dividers, distributes polygons to 2 pixel processors (PP). It enhances the overall performance of 3D pipeline by accelerating setup operations which took ~7,000 RISC cycles in the previous work [1-3]. A depth-first clock-gating (DFCG) scheme is applied to SlimShader for low power. DFCG can prevent the unnecessary shading and texturing by gating off the clock in the following datapath according to the results of the depth comparison. For realtime special effects, MP post-processes the rendered pixels of the previous frame transferring them to the display controller, while SlimShader renders the current-frame pixels on the back-buffers. 3DRE can accelerate the drawing of points, lines and rectangles for 2D graphics as well.
To realize low-power texturing, 24Mb texture memory is integrated with two texture engines (TE). Per-pixel dividers are implemented in each TE to support the perspective-correct addressing which can remove the artifacts in large polygons. Since TEs perform the bilinear MIPMAP filtering to draw more realistic images [7], 8 texel requests are generated at every cycle. Fetching 8 texels directly from 8 TMs may consume large amount of power due to the concurrent data transitions in many capacitive I/Os and the activation power of TMs themselves. Therefore, the number of memory requests is reduced by adopting Address Alignment Logic (AAL) as illustrated in Fig. 2.4.3. Since two PPs process adjacent pixels, the nature of MIPMAP filtering allows Spatial Aligner to find and remove the overlapped texels. Then Temporal Aligner compares the current texture addresses with previous ones and leaves only the different addresses. Since AAL reduces the average number of requests to less than 2.5, texels can be fetched from only 4 TMs at every cycle. AAL saves power by reducing the number of activated TMs while doubling the performance compared with single TE architecture.

PPO provides four clocks synchronized with PLL as shown in Fig. 2.4.4. Each clock can be selectively gated and its frequency is scalable by the software to adjust the frame rate during runtime. RISCclk and BEQclk run at the full speed of the RISC core, and REclk and MEMclk operate at the quarter frequency – 132/33MHz (RISCclk/REclk) for FAST mode, 66/16.5MHz for NORMAL, and 33/8.25MHz for SLOW. Figure 2.4.5 shows the composition of the power consumption for various applications. The implemented graphics LSI consumes 210mW in continuous calculation of texture-mapped 3D graphics applications at FAST mode. The embedded DRAM drastically reduces the power consumption since the external I/Os for 3D rendering are eliminated, and an additional 22% reduction is obtained by AAL and DFCG. Non-textured 3D applications and MPEG-4 video decoding consume 145mW and 85mW, respectively. Textured 3D rendering consumes 110mW at NORMAL, and 65mW at SLOW mode, respectively

The Graphics LSI is implemented using a typical 0.16μm DRAM process and its die area takes 121mm^2. The chip contains 1M logic transistors, 29Mb DRAM, 72kb SRAM and PLL. The logic components, SRAM and analog blocks are drawn with the design rule of peripheral transistors of the DRAM. They meet the performance requirement of mobile applications with little leakage current although they show relatively large gate delay and routing area compared with the pure logic process. Figure 2.4.6 shows the chip micrograph and Fig. 2.4.7 summarizes its features.

References

[1] C.W. Yoon, et al., "An 80/20MHz 160mW Multimedia Processor Integrated With Embedded DRAM, MPEG-4 Accelerator, and 3D Rendering Engine for Mobile Applications," *ISSCC Digest of Technical Papers*, pp. 142-143, 2001
[2] Y.H. Park, et al., "A 7.1GB/s Low-Power 3D Rendering Engine in 2D Array-Embedded Memory Logic CMOS," *ISSCC Digest of Technical Papers*, pp. 242-243, 2000
[3] R. Woo, et al., "A 120mW 3D Rendering Engine with 6Mb Embedded DRAM and 3.2Gbyte/s Runtime Reconfigurable Bus for PDA-Chip," *J Solid State Circuits*, pp. 1352-1355, Oct. 2002
[4] M. Ohashi, et al., "A 27MHz 11.1mW MPEG-4 Video Decoder LSI for Mobile Application," *ISSCC Digest of Technical Papers*, pp. 366-367, 2002
[5] J.H. Sohn, et al., "Optimization of Portable System Architecture for Real-Time 3D Graphics," *ISCAS Proceedings*, pp.769-772, 2002
[6] Y.D. Bae, et al., "A Single-Chip Programmable Platform Based on a Multithreaded Processor and Configurable Logic Clusters," *ISSCC Digest of Technical Papers*, pp. 336-337, 2002
[7] L. Williams, "Pyramidal Parametrics," *SIGGRAPH*, pp. 1-11, 1983

 0-7803-7707-9/03/$17.00

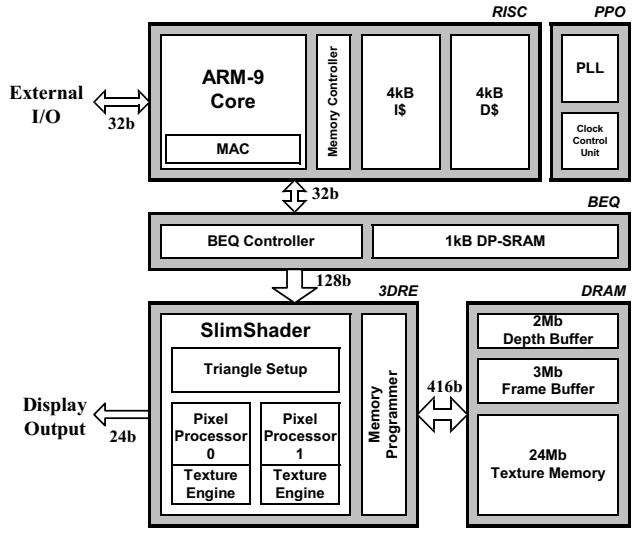

Figure 2.4.1 : Block diagram of Graphics LSI.

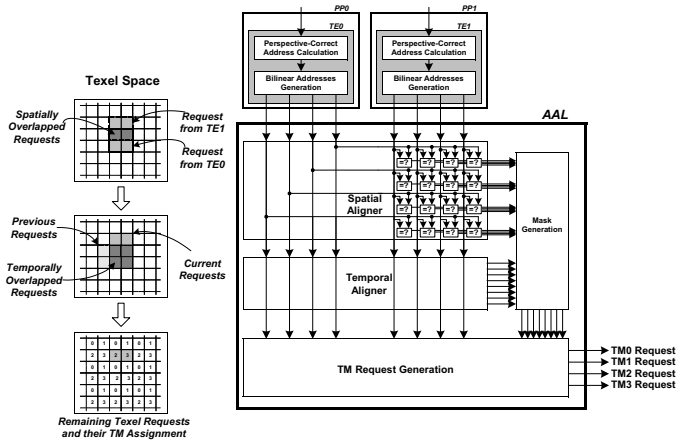

Figure 2.4.3 : Address Alignment Logic (AAL) and its operation.

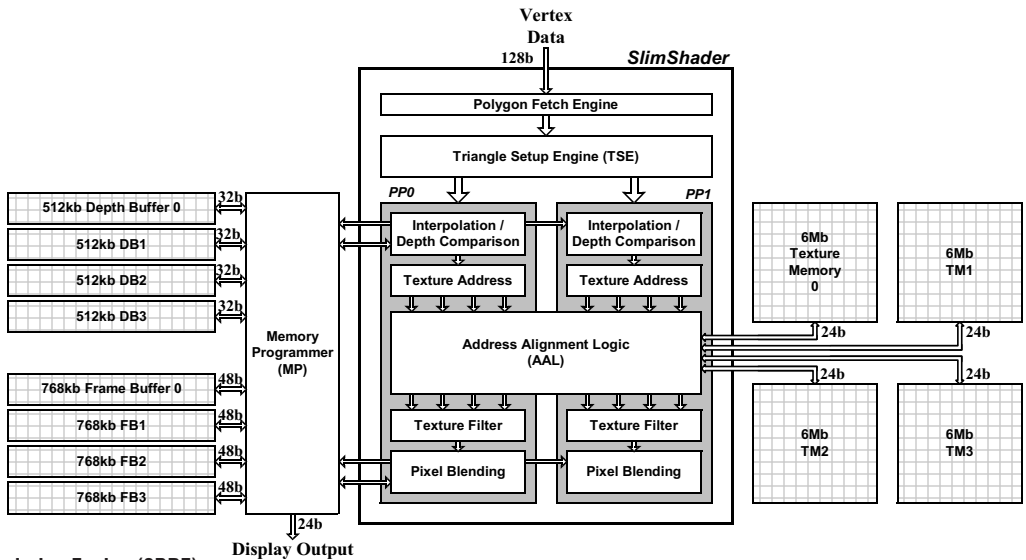

Figure 2.4.2 : 3D Rendering Engine (3DRE).

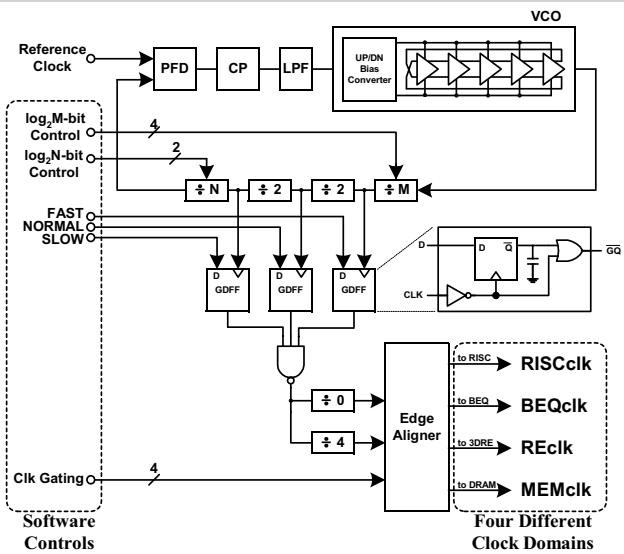

Figure 2.4.4 : Programmable Power Optimizer (PPO).

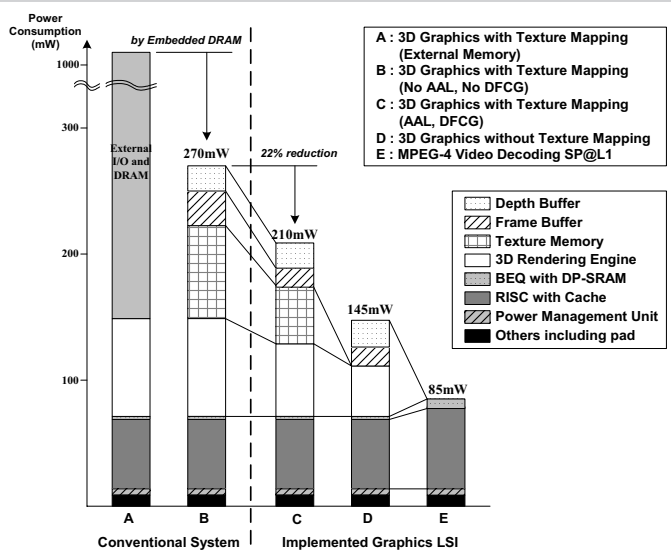

Figure 2.4.5 : Power consumption of Graphics LSI.

 0-7803-7707-9/03/$17.00

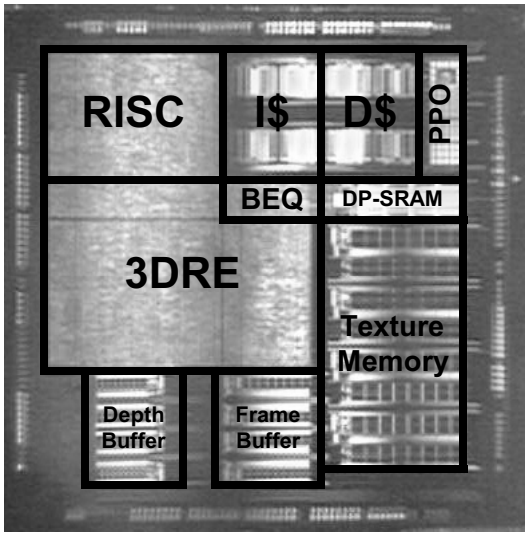

Figure 2.4.6 : Die micrograph.

Process Technology	0.16um CMOS DRAM with 1-W 3-Al
Power Supply	2.0V (DRAM Core), 2.5V (Logic), 3.3V (I/O)
Operating Frequency (RISC,BEQ/3DRE,DRAM)	FAST : 132MHz/33MHz NORMAL : 66MHz/16.5MHz SLOW : 33MHz/8.25MHz
Power Consumption	< 210mW
Transistor Counts	1M Logic 29Mbit DRAM 72kbit SRAM (9KByte)
Die Size	11mm x 11mm
Package	240pin PGA
Target Applications	Realtime 2D/3D Graphics Pipeline MPEG-4 SP@L1 Decoding MP3 Audio Decoding
3D Rendering Performance	66Mpixels/s, 264Mtexels/s Triangle Setup Engine Perspective-Correct Bilinear MIPMAP Texturing Gouraud Shading, Alpha Blending, Texture Blending Embedded 5Mb Double Depth/Frame Buffer Embedded 24Mb Texture Memory Antialiasing, Motion Blur, Fog, Special Effects

Figure 2.4.7 : Chip features and characteristics.

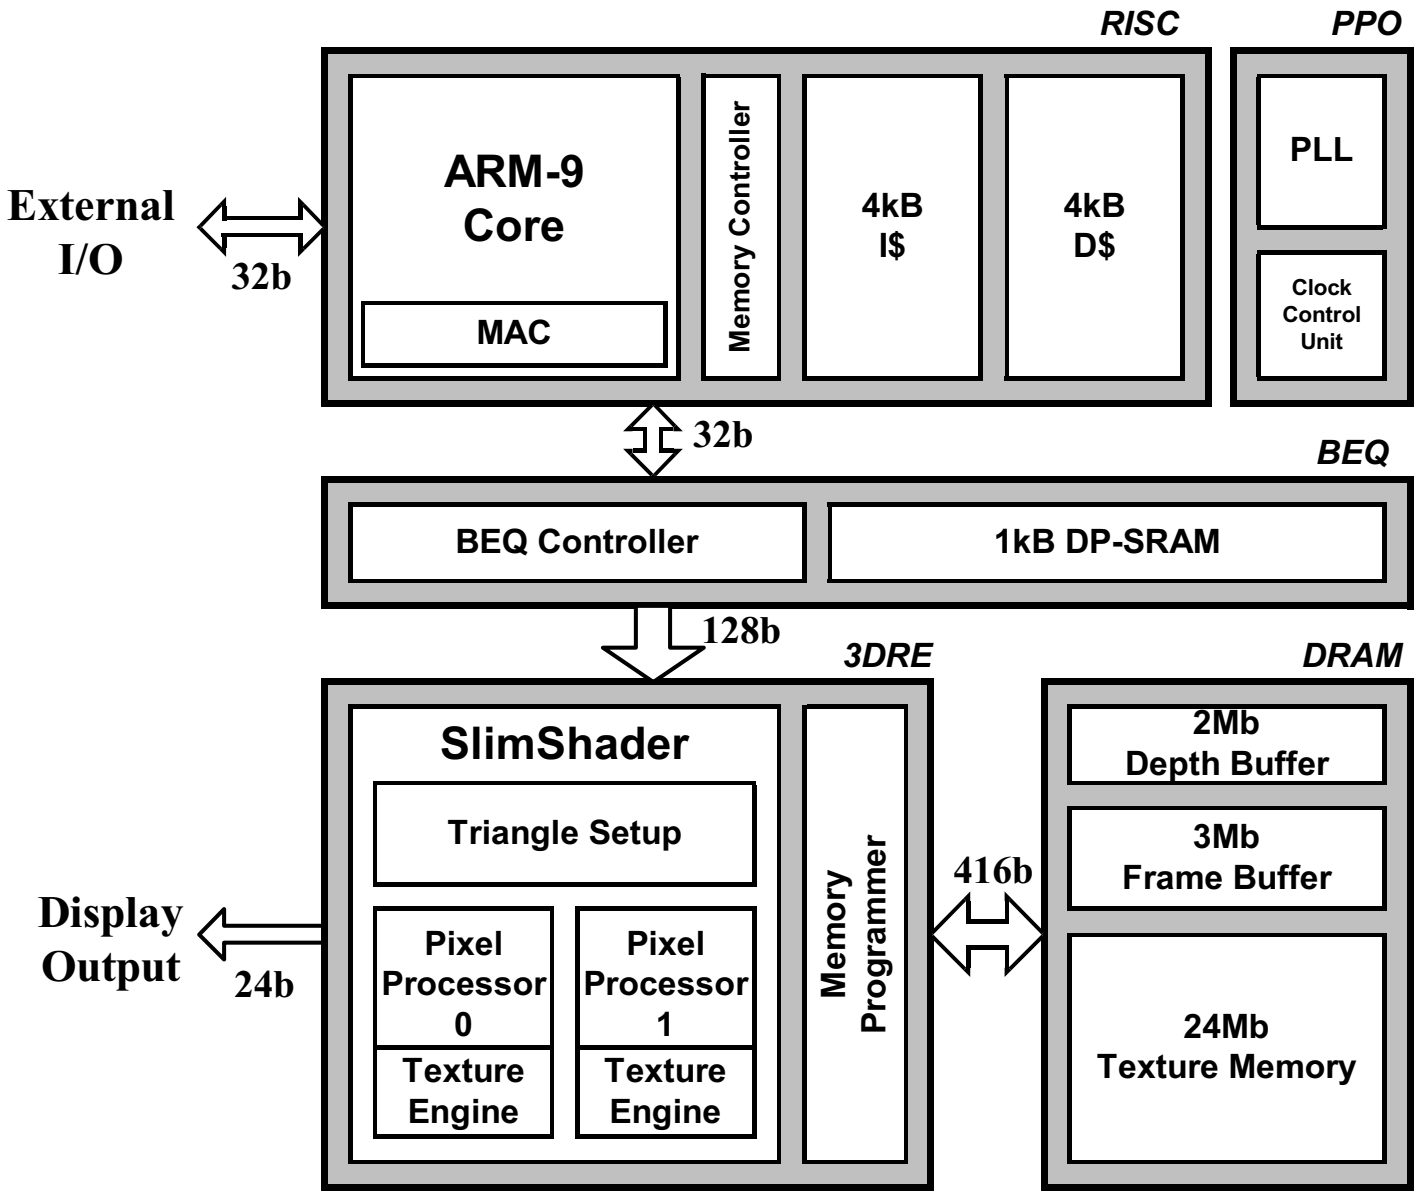

Figure 2.4.1 : Block diagram of Graphics LSI.

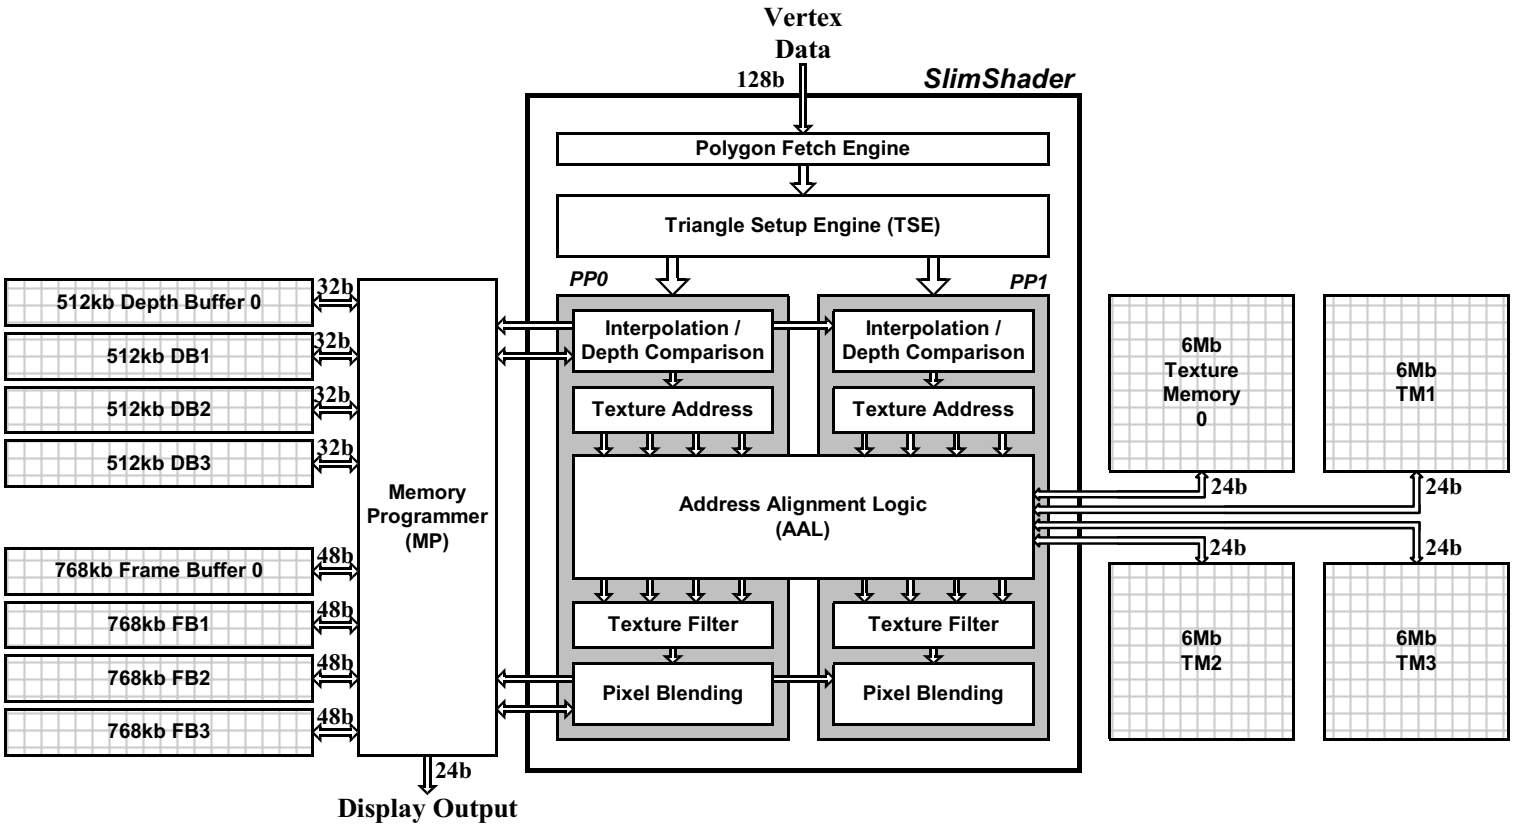

Figure 2.4.2 : 3D Rendering Engine (3DRE).

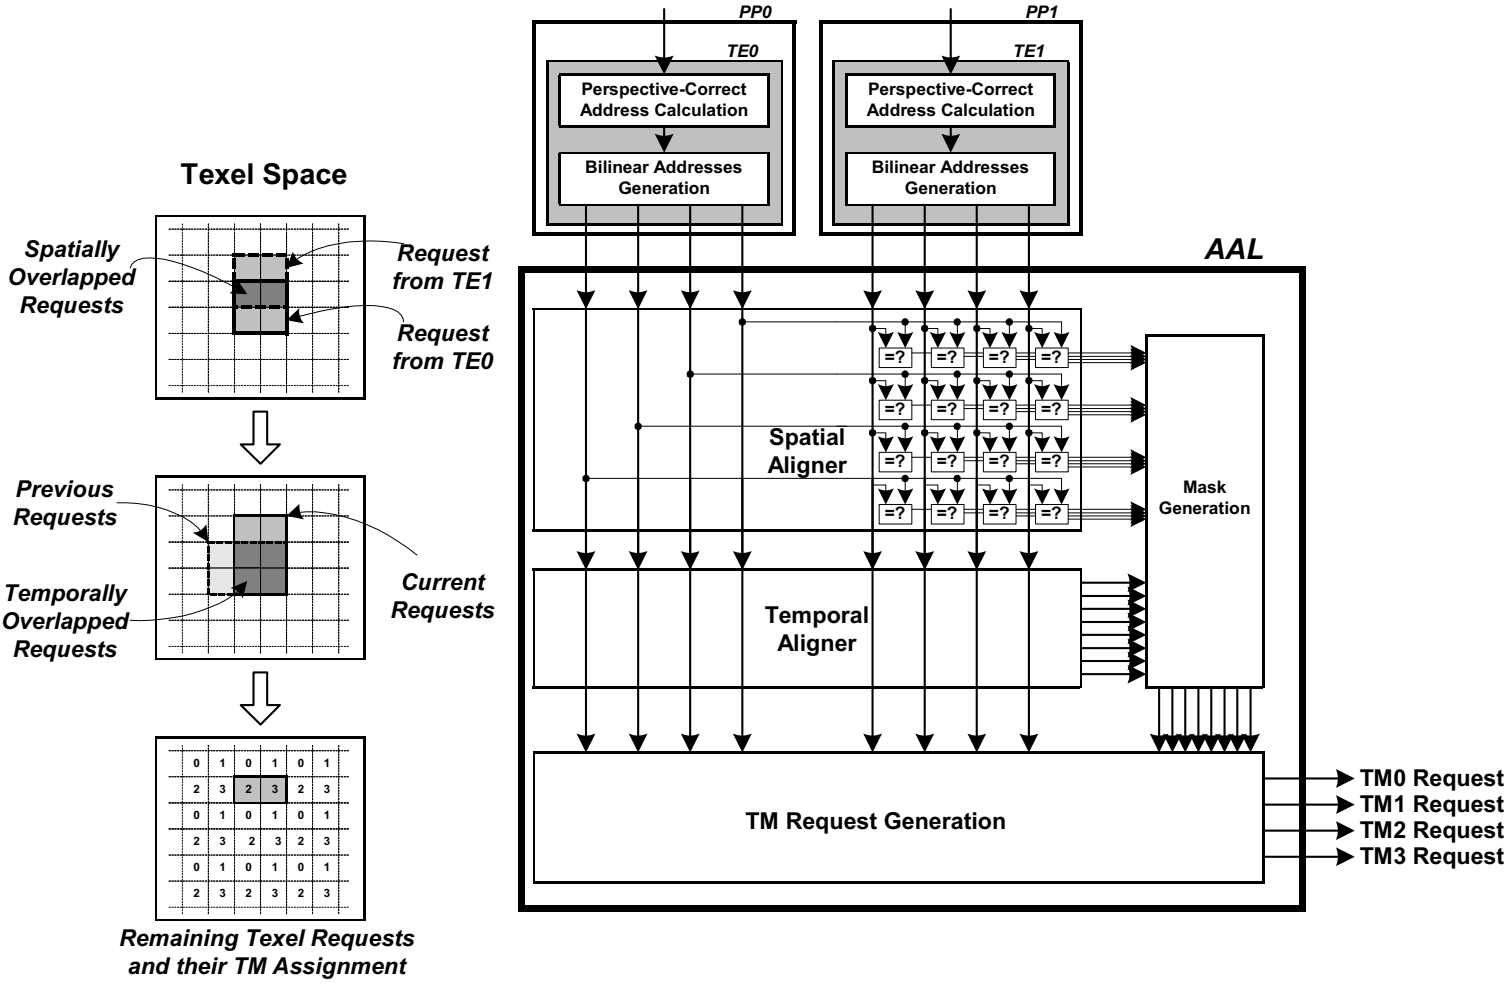

Figure 2.4.3 : Address Alignment Logic (AAL) and its operation.

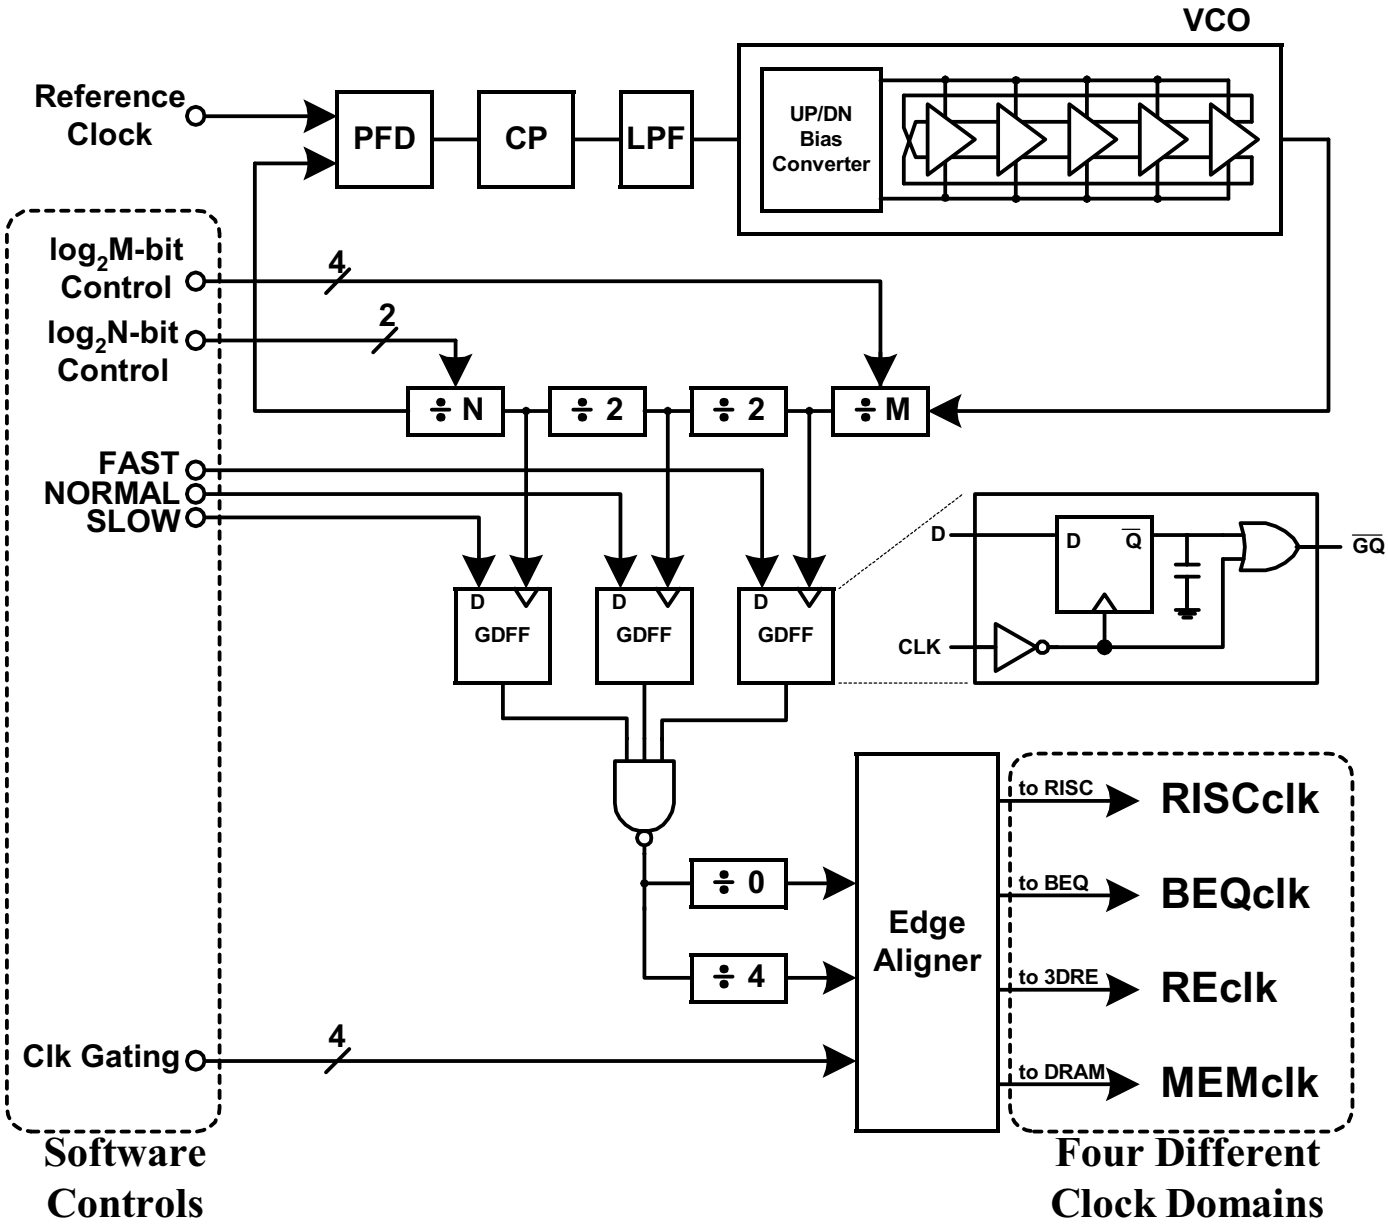

Figure 2.4.4 : Programmable Power Optimizer (PPO).

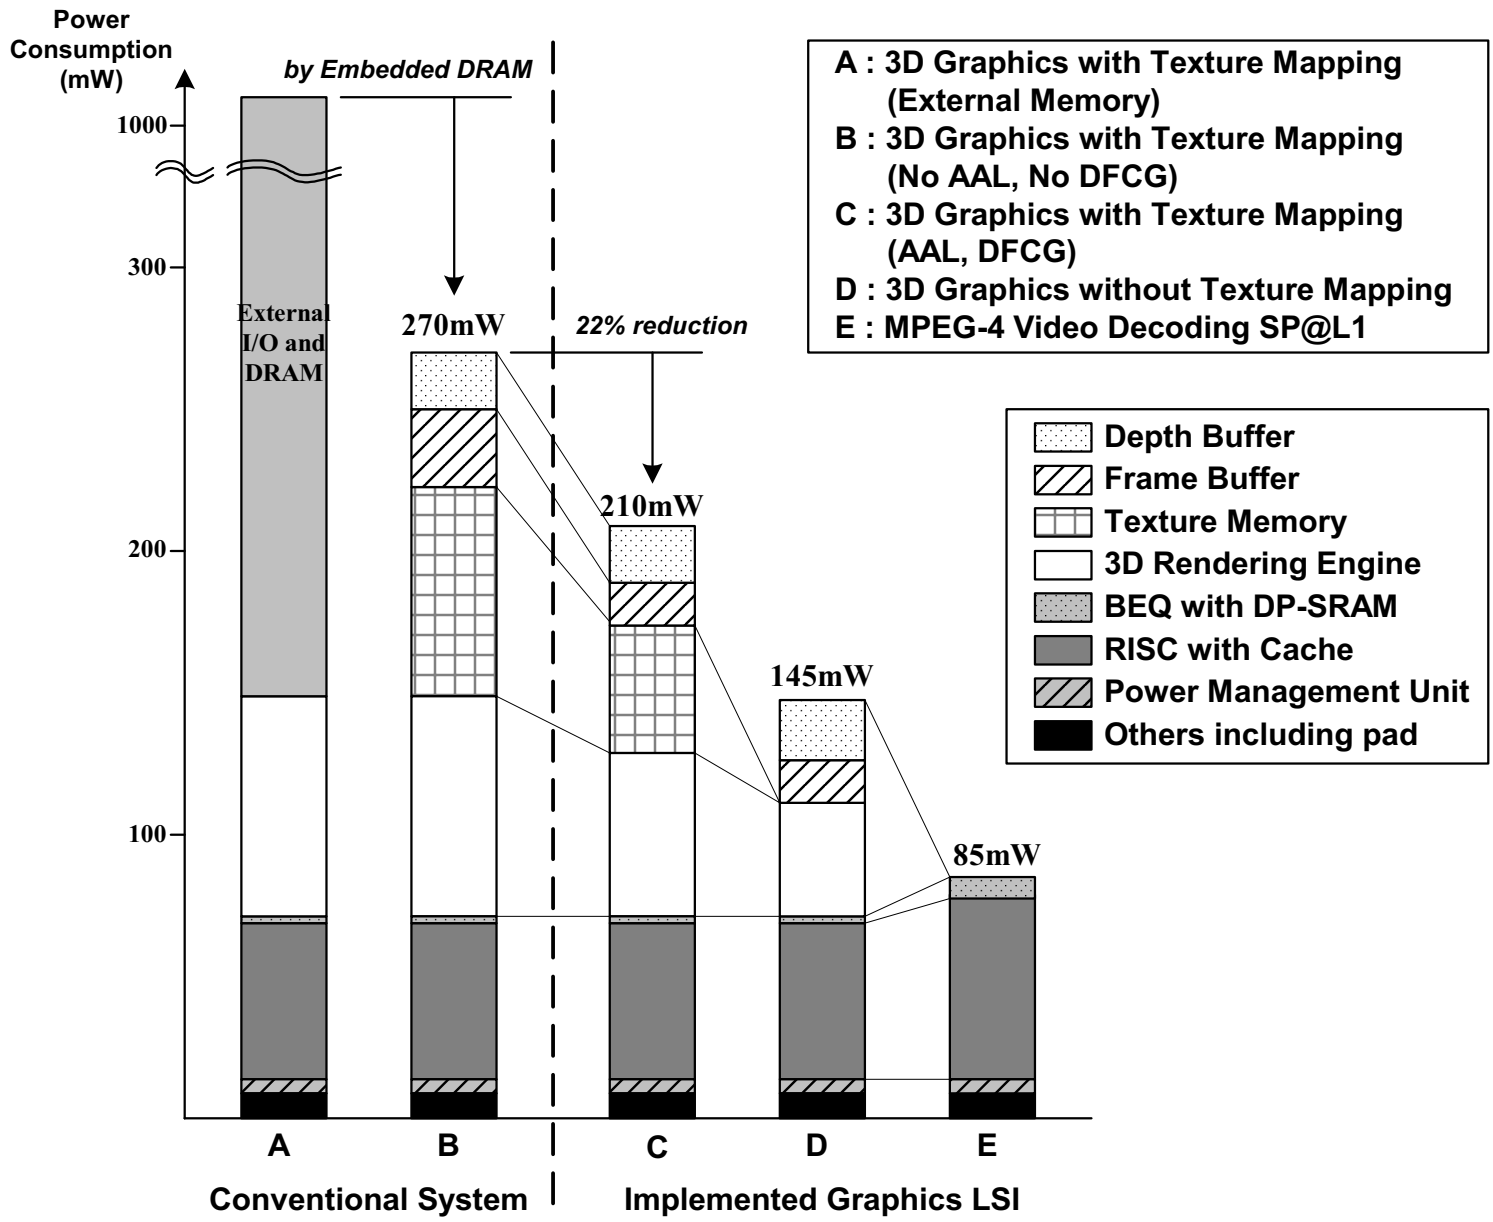

Figure 2.4.5 : Power consumption of Graphics LSI.

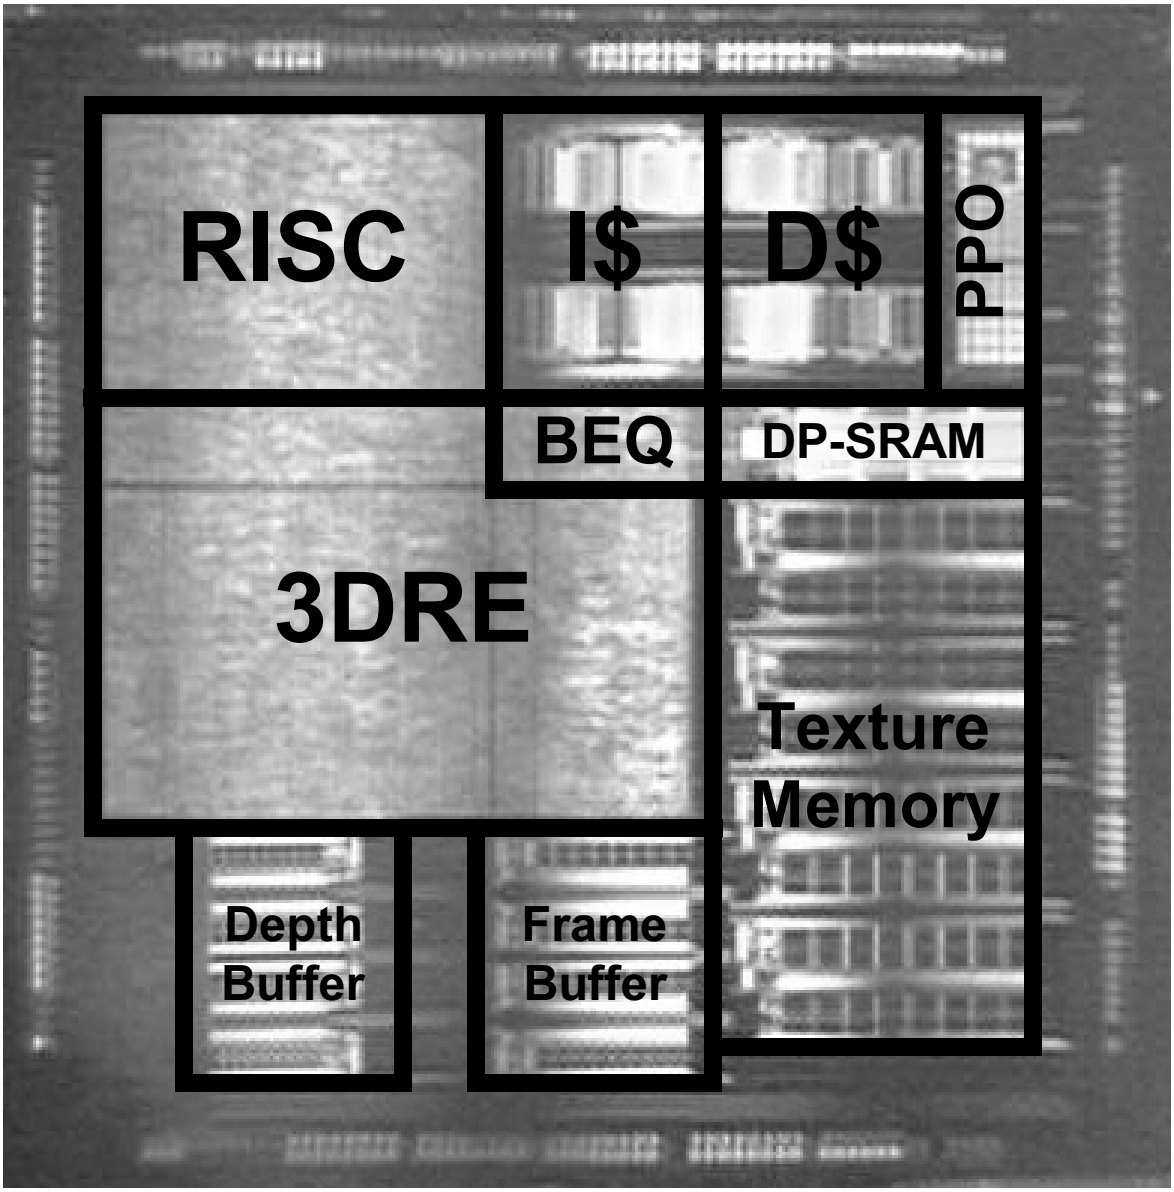

Figure 2.4.6 : Die micrograph.

Process Technology	0.16um CMOS DRAM with 1-W 3-Al
Power Supply	2.0V (DRAM Core), 2.5V (Logic), 3.3V (I/O)
Operating Frequency (RISC,BEQ/3DRE,DRAM)	FAST : 132MHz/33MHz NORMAL : 66MHz/16.5MHz SLOW : 33MHz/8.25MHz
Power Consumption	< 210mW
Transistor Counts	1M Logic 29Mbit DRAM 72kbit SRAM (9KByte)
Die Size	11mm x 11mm
Package	240pin PGA
Target Applications	Realtime 2D/3D Graphics Pipeline MPEG-4 SP@L1 Decoding MP3 Audio Decoding
3D Rendering Performance	66Mpixels/s, 264Mtexels/s Triangle Setup Engine Perspective-Correct Bilinear MIPMAP Texturing Gouraud Shading, Alpha Blending, Texture Blending Embedded 5Mb Double Depth/Frame Buffer Embedded 24Mb Texture Memory Antialiasing, Motion Blur, Fog, Special Effects

Figure 2.4.7 : Chip features and characteristics.