High Performance Graphics (2012), pp. 1–1
C. Dachsbacher, J. Munkberg, and J. Pantaleoni (Editors)

An Effective Task Scheduling Scheme for Multi-core Tile based Rendering GPU

Won-Jong Lee[1] and Seok-Yoon Jung[1] and Shi-Hwa Lee[1]

[1]SAIT, Samsung Electronics

1. Introduction

We present a task-scheduling scheme for a multi-core tile-based rendering GPU. Modern GPUs have employed unified shaders, which allow more flexible use of the graphics rendering hardware. This trend has recently been realized in mobile GPUs. For the best GPU performance, the utilization of unified shader cores and fixed-function hardware should be maximized while maintaining load balancing among multiple cores. In order to satisfy this requirement, we propose an effective task-scheduling scheme, which has the following two key features: 1) a shader-interleaving scheme that can hide latencies by exclusively executing different kernels, and 2) a task slot-based dynamic load-balancing scheme that evenly distributes workloads to the multiple cores. Experimental results show that a GPU adopting our scheme achieves significant performance gains of up to 2.8x compared to GPUs with non-unified shaders.

2. Multi-core Unified Shader GPU Architecture

Our multi-core unified shader architecture is an improvement over our previous work, the SRP (Samsung Reconfigurable Processor)-based GPU [LWK*]. Multiple unified shaders can be easily implemented as the SRP is reconfigurable. Each core includes vertex and pixel processing hardware in order to handle both types of data. Specifically, these units are batch management units (BMUs), primitive assembly (PA) units, tile dispatch units (TDUs), fragment generators (FGs), and raster operators (ROPs). The tile-binning unit (TBU) is individually equipped for parallel processing. The task scheduler (TS) is a newly designed unit responsible for dynamic load balancing and shader interleaving.

3. Latency Hiding and Dynamic Load Balancing

We developed a shader-interleaving scheme for hiding latency. The GPU drives the fixed-function hardware and the shader cores to execute in a fully pipelined manner. Unfortunately, this can cause pipeline bubbles because of inconsistent latencies between the hardware and the shader stage. If the kernel running on the shader core finishes before the hardware stage, the shader core remains idle until the next stream arrives. In order to hide these latencies, we propose a shader interleaving that allows different kinds of kernels to run exclusively. When a certain kernel (e.g., the vertex shader) becomes idle, the other kernel (e.g., the pixel shader) is immediately assigned to the shader core and executed without any delays. This can greatly improve overall throughput, which in turn leads to faster execution.

Dynamic load balancing can be achieved with the TS, which is designed for creating, scheduling, and assigning tasks to the shader cores. The TS first accepts graphics commands from hosts, generates the tasks on a per unit (e.g., a drawcall or a tile) basis, and then schedules them to the idle cores. Task slots indicating the status of each core are defined and used to select a target core in the TS. The TS also takes into consideration the fact that different tasks (e.g., vertex, pixel) should be interleaved as much as possible for hiding latency.

4. Experimental Results

The prototype GPU, including the proposed scheduler, was verified and evaluated by cycle-accurate simulation, RTL simulation, and FPGA targeting. The GPU was synthesized for a Xilinx Vertex6 FPGA board running at 25 Mhz. Rightware's 3DMark Mobile ES and GLbenchmark were used as test benchmarks. Various benchmarks having different workload characteristics were selected for testing unified shader efficiency. We compared the performance of our GPU (4 unified shaders) with that of a GPU with non-unified shaders (1 vertex and 3 pixel shaders) [LWK*] and having the same number of cores. Our GPU, which adopts a scheduling scheme, is 1.2 to 2.8 times faster, thanks to the effective combination of latency hiding and dynamic load balancing. Finally, it is expected that our GPU, including the proposed scheduler, will be a core intellectual property for future application processors.

References

[LWK*] LEE W.-J., WOO S.-O., KWON K.-T., SON S.-J., MIN K.-J., JANG G.-J., LEE C.-H., JUNG S.-Y., PARK C.-M., LEE S.-H.: A scalable gpu architecture based on dynamically reconfigurable embedded processor. In *HPG 11, Poster*. 1