

Intro to GPGPU

General Purpose GPU programming

Massimo Coppola

13/05/2015

GPU Computing

- The need for efficient specialized processing of 3D meshes promoted the adoption of the SIMD programming model
- How the model evolved
- What are GPUs good at?
 - Large data sets
 - Arithmetic intensity = High compute/IO ratio
 - Minimal control flow or recursion
 - High locality

The birth of Graphic Processing Units

- The graphics pipeline
 - General methodology to produce graphic output on raster devices like computer displays
 - Start from elementary data (vertexes) and transform them into pixels
 - State of the art evolved over the years, to possibly very complex structures
 - Cfr. OpenGL 1.1 state machine
 - We only survey the basic principles
- Graphics pipeline, or its stages, can have both SW and HW implementation
- Tradeoff between flexibility and performance

The objects

- Vertex : a point in a coordinate system
- Primitive : graphic object comprising one or more vertexes, possibly other parameters
- Pixel : image element in a raster display

- Coordinate systems for Vertexes, Primitive, Pixel usually do not coincide
- They have typically different dimensionality
 - E.g. render 3D space on a 2D display

- Widespread use of homogeneous coordinates
 - Represent points in 2D spaces with 3 coordinates, and points in 3D spaces with 4-dimension coordinates
 - Allow representing linear affine transformations and projections as linear operators → implemented as matrix multiplication
 - Common, very efficient execution of graphic transformations

Elementary Graphics Pipeline

1. Vertex generation
 2. Vertex processing
 3. Primitive generation
 4. Primitive processing
 5. Pixel generation (Rasterization)
 6. Pixel Processing
 7. Pixel writing
- Some steps are more deeply customizable
 - Some steps are efficiently realized in HW

Example

1. Vertex generation
 - retrieve/generate coordinates, apply geometric transformation
2. Vertex processing
 - Apply/attach visualization parameters to vertexes, apply per-object transformations
3. Primitive generation
 - Group connected vertexes and turn them into squares, spheres, surfaces, lines ...
4. Primitive processing
 - Apply shading models, colors, textures custom transformation to primitives
5. Pixel generation (Rasterization)
 - Slice primitives according to the output device resolution and features
 - Compute/interpolate texture pixels from texture memory matching with primitive coordinates, to define each pixel characteristics in the slices
6. Pixel Processing
 - Process pixels according to lighting models, (anti) aliasing and other postprocessing techniques
7. Pixel writing
 - Framebuffer operation, appropriate memory format (e.g. alpha channel)

Evolution and transformation of GPUs

- From 1985 (e.g. Commodore Amiga) to 1990 (S3 chips and followers) and beyond, 2D and then 3D accelerated units spread in the personal computer market
 - Early experiences at Xerox PARC in 1975
 - Driven mainly by the game market
 - Less by Windowing systems, professional graphic use
- More and more specific stages in the pipeline implemented in HW on a chip of the graphic device
- In the end, all stages of a 3D pipeline implemented in HW
- **Load balancing among the stages and flexibility become issues for all-HW implementation**

Load balance in the pipeline

- More pixel than raster elements (slices of primitives)
- More raster elements than vertexes
- Expected primitive distribution, surface hiding and other masking effects can affect this balance




Figure 14. Characteristic pixel and vertex shader workload variation over time


Image from NVIDIA GeForce 8800 architecture documentation, 2006

Push toward unification

- A Fixed number of vertex units and pixel units leads to poor resource use on different workloads
- Fixed, HW-cabled functionalities are easily reproduced in SW
 - no generic CPU functionalities needed
- Special units gradually replaced by unified units alike to stream processors, with limited programming capabilities
- Allocation of code to stream units initially done by specialized SW = graphic drivers

First “programmable” GPUS

- Replace the graphic pipeline in the HW
 - e.g. texture caching and sampling
- Maintain some special purpose units in HW
 - Custom RAM bus
 - No read/write conflicts
 - Small caches
 - High on-chip ALU/memory ratio
 - Single precision, non IEEE floating point
- Architecture optimized for streaming



Example from GeForce 8800 docs

GPGPU

- General Purpose Graphic Unit Programming
- More and more graphic cores, and increasing core computing power
- People started to tap into the graphic unit via OpenGL primitives
 - Exploit the computational semantics of specific graphic operation to achieve access to the HW
 - Tasks fit for stream processing : physics, image manipulation, large data with few dependencies
- GPGPU research area was born
 - Physical simulation coupled with rendering
 - Textures and vertexes (read-only) are input streams
 - Need to write results !
 - Copy framebuffer (write-only) to texture after computation
 - Skip last pipeline stages and save results to texture memory (*stream output* in DirectX10)

New GPUs


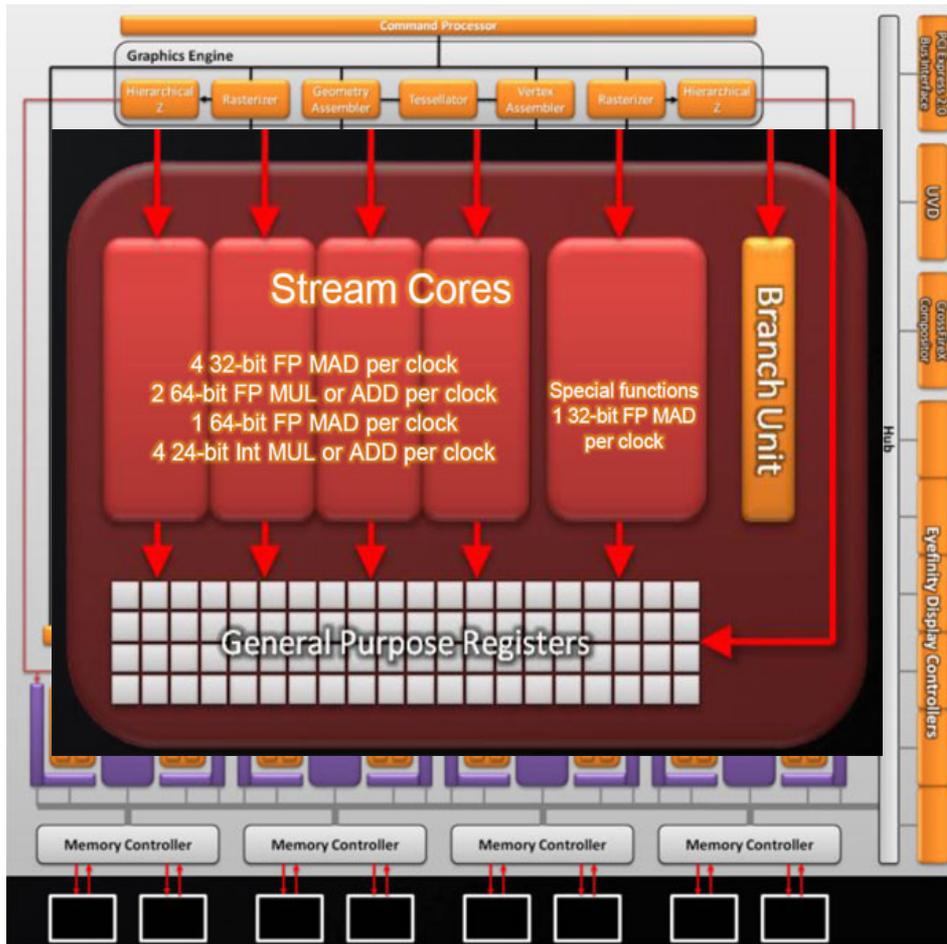
- GPU producers understood the market value
 - GPU became more programmable
 - General programming issues accounted for
 - Double precision IEEE f.p. arithmetic
 - More efficient branches in GPU code
- Architecture is still optimized for streaming
 - The model exposed is very much SIMD like
 - No support for reading/writing the same memory area
 - No or limited support for communication among code instances
 - to avoid synchronization and pipeline stall detection logic
- GPUs are optimized for long computation run with reduced dependencies
- CPUs for general access patterns and concurrency

GPU optimizations

- Very large RAM bus
- Multiple transfer per cycle
 - rising/falling clock edge
- Low latency for sequential access
- High ALU density
 - Many ALU controlled by the same control unit
- Grouped as thread processors
 - All the core in a same thread block share same SIMD model Code
 - Share code and program flow, cores can just skip
 - Sometimes available: shared set of registers and caches
- Different threads blocks are truly independent


Hardware Model

ATI "Cypress" RV870



Computational Model

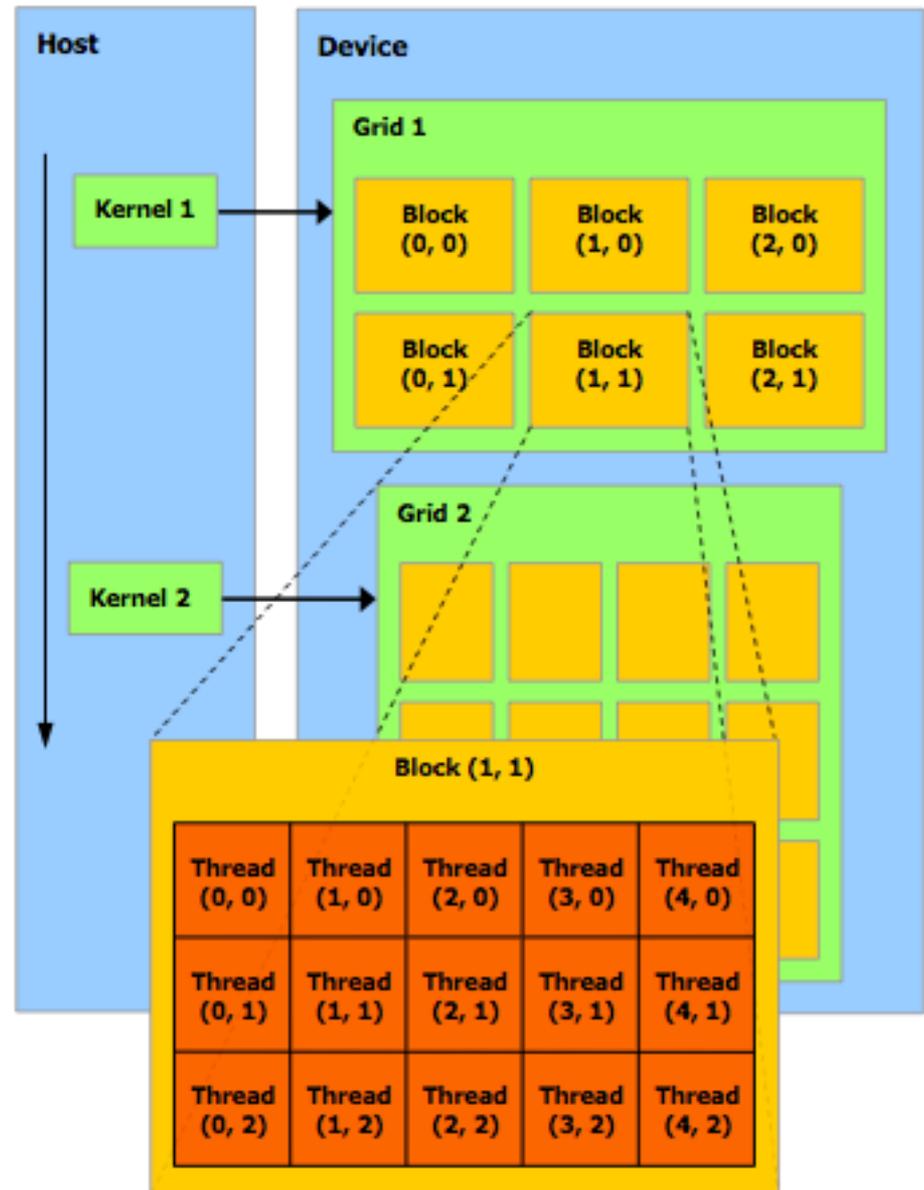
- Stream Computing
 - SIMD-like programming model
 - Multiple processing units
- Non-determinism
 - how data in streams gets processed by the cores is left to the board firmware
- The computation of each core is driven by a program, **kernel**
- The GPU infrastructure is responsible for assigning cores to kernels
 - each running instance of a kernel is called **thread**
 - each thread has an associated set of output locations in the GPU memory referred as the **domain of execution**.



Proprietary Programming Models


- Brook+ and CUDA
 - Provide sound language abstractions to define computational kernels
 - In a subset of standard sequential languages
 - each one assigned to one or more thread processors
 - Main issue is to define in which memory space each data/variable is actually kept

- Each kernel is mapped onto one or more thread blocks
- Each Block can execute several sub-computations
- Kernel instances (threads) in a thread block can be interleaved or parallel



The host issues a succession of kernel invocations to the device. Each kernel is executed as a batch of threads organized as a grid of thread blocks

- The kernel instance in each core can access several spaces
- Language qualifiers on functions and variables
- Concurrency issues
→ avoid dependencies



A thread has access to the device's DRAM and on-chip memory through a set of memory spaces of various scopes.

More general Programming Models

- OpenCL
 - New API more focused on computational exploitation of GPU
 - Will be part of this course
- RapidMind
 - Language based approach which focuses on portability
 - Same set of SIMD-like primitives could be compiled to
 - GPUs
 - Cell Multicore
 - X86 multicore CPUs
 - Interesting idea → acquired (by Intel) in 2009

GPU and CPU interaction

- The main limit of conventional GPU approach
- Interaction with the CPU bus is a bottleneck
 - PCI bus is fast, but slower than memory interface of the GPU
 - Data exchange rate and overhead is influenced by the driver/OS management and by the hardware capability (is DMA controlled by both ways?)
- To scale you need an ALU-intensive, regular problem and infrequent interaction with the CPU
- Or efficient asynchronous interaction with the CPU bus