263-3502-00L Data Processing on Modern Hardware **Assignment 6**
FS 2012 Out: 2012-11-29
Dr. Jens Teubner, Cagri Balkesen Discussion: 2012-12-13
Department of Computer Science, ETH Zurich

6 Query Execution on GPGPUs

In this assignment you are going to port the *column store* implementation introduced in assignment 2 onto the GPU using CUDA or OpenCL. We provide a code template [1] that contains the TPC-H data loading module and CUDA[1] skeleton code. Before you run cs_gpu do not forget to add lineitem.tbl from assignment 2 to the directory of cs_gpu.

The N row data file is loaded into the following column store representation in main memory.

```
/* column store */
uint32_t orderkey[N];
uint32_t partkey[N];
...
char     comment[45*N];
uint32_t shipdate[N];
```

In this exercise we now ask you to implement the following SQL query

```
SELECT sum(quantity*extendedprice)
  FROM lineitem
 WHERE suppkey<Z
```

where Z determines the selectivity of the query.

6.1 CUDA Setup

Important: You can solve this exercise even if you do not own a CUDA/OpenCL-capable GPU. You can run your code in device emulation mode (compiler option -deviceemu).

If you are working on Mac OS X 10.6 (Snow Leopard) OpenCL is already installed. If you have an nVidia GPGPU or want to run your program in simulation mode, you need to install CUDA on your machine. Note that the nVidia toolkit also supports OpenCL. Download the following packages from http://www.nvidia.com/object/cuda_get.html:

- Developer Drivers for your OS
- CUDA Toolkit (provides nvcc, cuda-gdb, and cudaprof)
- CUDA SDK code samples (documentation and examples).

Choose the appropriate download for your operating system (Linux, Windows, or MacOS X). If you have a CUDA-capable GPU[2] also install the latest CUDA driver from this URL. Install the packages by executing the downloaded files.
You may want to set the environment variables accordingly. For example, add the following lines to your .bashrc:

[1] If you prefer to program in OpenCL, translate the few CUDA functions in the code template to OpenCL.
[2] A list compatible GPUs is available at http://www.nvidia.com/object/cuda_learn_products.html.

```
# CUDA stuff
export PATH=/usr/local/cuda/bin:/usr/local/cuda/cudaprof/bin:$PATH
export LD_LIBRARY_PATH=/usr/local/cuda/lib64:\
   /usr/local/cuda/cudaprof/bin:$LD_LIBRARY_PATH
export MANPATH=/usr/local/cuda/man:$MANPATH
```

Test your installation by compiling the sources [1] we provide. Typing `make` should produce three executables (`cs_gpu`, `example-cuda` and `example-opencl`). Additional installation instructions can be found in [2].

6.2 Task

Implement the aforementioned query in CUDA or OpenCL. Start from the provided template. To help you get started we also provied a simple program that increments all values of an array on the GPU. One version (`example-cuda`) is written in CUDA and an equivalent version (`example-opencl`) is written in OpenCL. Especially, if you want to solve this exercise in OpenCL, you should look at `example-opencl` so that you can translate the provided template from CUDA to OpenCL accordingly.

Questions:

1. What is the optimal block size for your implementation? Use the *GPU Occupancy Calculator* Excel spreadsheet (see slides) and determine the optimal block size.

2. Which performance impact does the selectivity parameter `Z` have?

3. If you own a CUDA-capable GPU or OpenCL-capable device only: What speedup do you get with your parallelized implementation?

4. If you own a CUDA-capable GPU: What is the occupancy of the multiprocessors? Profile your application with `cudaprof`.

6.3 Tips

- For reference use the *NVIDIA CUDA Reference Manual v2.3* [4].

- A good introduction with additional information can be found in the *CUDA Programming Guide* [3].

- The `sum` computation is a parallel data reduction. Study the lecture slides. A complete implementation can be found in the `reduction` example of the NVIDIA SDK sample codes you installed earlier.

References

[1] TPC-H data loading modules and cuda skeletons. http://www.systems.ethz.ch/sites/default/files/file/dpmh_Fall2012/src-handout-06_tar.bz2.

[2] NVIDIA Inc. *NVIDIA CUDA Getting Started Guide for Linux – Installation and Verification on Linux System*, August 2010.

[3] NVIDIA Inc. *NVIDIA CUDA Programming Guide v3.2*, October 2010.

[4] NVIDIA Inc. *NVIDIA CUDA Reference Manual v3.2 Beta*, August 2010.