Prof. G. Zachmann
A. Srinivas

University of Bremen
School of Computer Science
CGVR Group
April 25, 2014

Summer Semester 2014

Assignment on Massively Parallel Algorithms - Sheet 2

Due Date 07. 05. 2014

Hint: You can use one of the examples on the lecture homepage or from the Cuda SDK (included in the Cuda installation package) to test if Cuda works at all on your computer. For example, you can run the example `1_Utilities/deviceQuery`.

Exercise 1 (CUDA basics: Memory, *3 Credits*)

In the framework `cudaMallocAndMemcpy`

a) Allocate two arrays `d_a` and `d_b` on the device of the same size as the array `h_a` on the host. You can use `sizeof(datatype)` to get the number of bytes for `datatype`.

b) Copy `h_a` on the host to `d_a` on the device.

c) Do a device to device copy from `d_a` to `d_b`.

d) Copy `d_b` on the device back to `h_a` on the host.

e) Free `d_a` and `d_b` on the device.

Exercise 2 (CUDA basics: Launching kernels, *3 Credits*)

Starting from the framework `myFirstKernel`

a) Allocate device memory for array `d_a` to hold the results of the kernel.
Overall `numBlocks`×`numThreadsPerBlock` threads will be launched, and each thread writes to one array element.

b) Configure and launch the kernel `myFirstKernel( int *d_a )` using a 1D grid of 1D thread blocks.

c) Have each thread set an element of `d_a` as follows:
```
idx = blockIdx.x*blockDim.x + threadIdx.x
d_a[idx] = (blockIdx.x - 6) * (100 - threadIdx.x)
```

d) Copy the result in `d_a` back to the host memory to array `h_a`.

e) Free the device array `d_a`

f) Cuda kernels cannot return a value. What could be the reason for this?