CUDA Exercises Lab 2 (26th Aug 2019)

Logistics

- Block A (Friday / 23rd Aug)
 - How to use CUDA
 - How to launch a CUDA kernel and array indexing
- Block B (Monday / 26th Aug)
 - How to index 2D indexes
 - How simple convolution works
 - (Bonus) How to use shared memory

Clarification

- Time difference computation should now be fixed
- Tegner supports nvprof
 - srun -n 1 nvprof ./a.out

Setting up...

- 1. kinit --forwardable your_username@NADA.KTH.SE
- 2. ssh -Y your_username@tegner.pdc.kth.se
- 3. cd /cfs/klemming/nobackup/[your_initial]/[your_username]
- 4. module load git
- 5. module load cuda/7.0
 - 1. FORTRAN: module load pgi
- 6. git clone https://github.com/PDC-support/cuda-lab-exercises.git
- 7. Go inside the code directories
 - 1. C: cd cuda-lab-exercises/lab 2/C
 - 2. FORTRAN: cd cuda-lab-exercise/lab_2/Fortran

□ CUDA Laboratory 2

Introduction to High-Performance Computing



[™] Introduction

In this second laboratory about GPU programming in CUDA, we are going to continue building your skills in order to develop more advanced GPU-accelerated applications. As a friendly reminder, the laboratory is divided in two different blocks and exercises:

(Again) Instructions also available in **README on Github page**





Compiling

nvcc -arch=sm_30 lab02_ex3_6.cu -o lab02_ex3_6.out

Compiling with Fortran



nvcc -arch=sm_30 lab02_ex3_6.cu -o lab02_ex3_6.out

Automatically creates binary

Compiling with Fortran

Main program and helper functions (timer) exists as C file, compiled and linked during compilation!



Running CUDA Program

salloc --nodes=1 -C Haswell --gres=gpu:K420:1 -t 00:05:00 -A edu19.summer --reservation=summer-2019-08-26

We are going to use the "thin nodes" on Tegner

Request for one K420 GPU on the node

Running CUDA Program

srun -n 1 ./lab02-ex3-6.out image/lab-02.bmp Just like before Path to an image



To view an image

display -resize 1280x720 images/lab02.bmp



One Pixel

В	G	R	В	G	R	В	G	R		
	G		B	G	R	B	G	R		
В	G	R	В	G	R	В	G	R		
B	G	R	B	G	R	B	G	R		
B	G	R	В	G	R	В	G	R		

Header

Colorimetric method (luminance-preserving) method

$$Y_{out} = 0.0722 \times B_{in} + 0.0722 \times B_{in}$$

$0.7152 \times G_{in} + 0.2126 \times R_{in}$



$Y_{out} = 0.0722 \times B_{in} + 0.7152 \times G_{in} + 0.2126 \times R_{in}$

- Understand cpu_greyscale()
- Implement gpu_greyscale()
 - Block and Grid in 2D is already setup with you
 - Assume each thread responsible for one output pixel
 - Boundary check!!!

display -resize 1280x720 images/lab02_result_1.bmp



- Apply a 3x3 convolution matrix on all the pixels of the image
 - Map each pixel as the center of the 3x3 matrix
 - Apply weights with surrounding pixels

$$\left(\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} * \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix} \right) [2,2] = (i * 1) + (h * 2) + (i * 1) + (i + (i$$

+(g*3)+(f*4)+(e*5)+(d*6)+(c*7)+(b*8)+(a*9).



_			t(0,0)	t(1,0)	t(2,0)	t(3,0)
A	B	С				
D	E	F	t(1,0)	•••		
G	Н					

- Understand cpu_applyFilter()
- Implement gpu_applyFilter()
 - The kernel is launched from kernel function from gpu_gaussian()
 - What is the modifier for functions launched by another kernel?

Understand how a convolution matrix is applied to a certain pixel

images/lab02_result_2_comp.jpg display images/lab02_result_2_comp.jpg

ImageMagick: lab01_result_2_comparison.jpg (on tegner-login-1.pdc.kth.se)



- montage -tile 2x1 -crop 320x180+512+512 -geometry 640x360 \ images/lab02_result_1.bmp images/lab02_result_2.bmp \

- Apply a Sobel filter to our smoothened image
 - Compute an approximation of the gradient of the image intensity
 - Creates an image where the edges are emphasized
 - A base of full edge detection algorithms such as Canny algorithm
 - Apply as two 3x3 convolution filter
 - Computes derivatives on horizontal and vertical directions
 - Resultant pixel value computed as a square root

$$\mathbf{G}_x = egin{bmatrix} +1 & 0 & -1 \ +2 & 0 & -2 \ +1 & 0 & -1 \end{bmatrix} * \mathbf{A} \hspace{1.5cm} ext{and} \hspace{1.5cm} \mathbf{G}_y = egin{bmatrix} +1 & +2 & +1 \ 0 & 0 & 0 \ -1 & -2 & -1 \end{bmatrix} * \mathbf{A}$$

 $\mathbf{G} = \sqrt{}$

$$\overline{{\mathbf{G}_x}^2+{\mathbf{G}_y}^2}$$

- Understand cpu_sobel()
- Implement gpu_sobel()

montage -border 0 -geometry 640x360 -tile 3x1 \
images/lab02.bmp images/lab02_result_1.bmp \
images/lab02_result_3.bmp images/lab02_result_3_comp.jpg
display images/lab02_result_3_comp.jpg



CPU and GPU Timing

Elapsed CPU: 52ms / Elapsed GPU: 16ms Elapsed CPU: 270ms / Elapsed GPU: 19ms Elapsed CPU: 570ms / Elapsed GPU: 20ms

Step #1 Completed - Result stored in "images/lab02_result_1.bmp". Step #2 Completed - Result stored in "images/lab02_result_2.bmp". Step #3 Completed - Result stored in "images/lab02_result_3.bmp".

When testing your GPU functions, comment out the respective CPU functions, otherwise you might be reading the CPU generated file whereas the GPU kernel failed!

Check README.md from lab Github for more information

Have fun!