

# GPU Programming in Computer Vision: Day 1

---

Date: September 17, 2018

---

## Setup and Code Framework

Include the path to the nvcc Compiler:

Open the `.bashrc` script: `gedit ~/.bashrc`

Add at the end of the file: `export PATH=/usr/local/cuda-9.1/bin:$PATH`

Reload to apply the changes: `source ~/.bashrc`

In your home directory, execute:

```
git clone https://svncvpr.in.tum.de/git/cuda_ss18
```

The framework shows how to use OpenCV to load/save/display images, access the camera, measure the runtime, and process the command line parameters.

Create a directory to build: `mkdir build`

Change into the directory: `cd build`

Generate the Makefile: `cmake ..`

Compile: `make`

Run executable for exercise, e.g.: `./bin/ex3_gamma -i ../images/dog.png`

A complete framework for all exercises is already provided, in which you should insert the missing code for each exercise. Parts to be completed are indicated by `// TODO` comments.

## General Code Requirements for the Exercises

- Keep your code as general as possible. It must be applicable for images with an arbitrary number of channels  $n_c$  (if not stated otherwise).
- Always comment your code.
- Whenever new parameters are introduced, use the `cv::CommandLineParser` to read in these parameters from command line arguments.
- Always include code for measuring run times and test how much time your overall computation for the exercise takes.
- When finished, test on several still images. If you want, also test on live webcam stream.
- Always use the macro `CUDA_CHECK` after each CUDA call, e.g.  
`cudaMalloc(...); CUDA_CHECK;`
- Hint: Multi-channel images are layered: access `imgIn(x, y, channel c)` as  
`imgIn[x + (size_t)w*y + (size_t)w*h*c]`
- Always use a variable (of type `size_t`) for an index which you need *more than once*, e.g.  
`size_t ind = x + (size_t)w*y + (size_t)w*h*c;`
- Always cast to `size_t` in integer products when computing array indices or image sizes

## Exercise 1: Check CUDA and the installed GPU (1P)

1. Open a terminal and check whether CUDA is installed: `nvcc --version`. Which version is installed?
2. Go to the “CUDA samples” folder<sup>1</sup> and run `deviceQuery`. Find out the following:
  - (a) name of the installed GPU and its compute capability (“CUDA Capability”)
  - (b) number of multiprocessors and CUDA cores
  - (c) amount of global memory
  - (d) max. amount of registers and shared memory per block

## Exercise 2: First CUDA Kernels (3P)

Implement the following CUDA kernels:

1. In `basic/squareArray.cu`, complete the CUDA code for squaring an array on the GPU. Implement the square operation as a `__device__` function. Compile with `nvcc -o squareArray squareArray.cu`
2. In `basic/addArrays.cu`, complete the CUDA code for adding two arrays on the GPU. Implement the addition operation as a `__device__` function.
3. Now, compile both files with (similarly for `addArrays`):  
`nvcc -o squareArray squareArray.cu --ptxas-options=-v`  
How many registers are used by your kernels?

## Exercise 3: Gamma Correction (4P)

Perform gamma correction on the colors of the input image:  $u_c^{\text{out}}(x, y) = u_c(x, y)^\gamma$ ,  $\gamma > 0$  for each pixel  $(x, y) \in \Omega$  and for each channel  $c \in \{1, \dots, n_c\}$ .

1. Write the CPU version. Keep your code general, so that it can process grayscale ( $n_c = 1$ ) as well as color images ( $n_c = 3$ ). Test on several input images, with and without the `-b` parameter for grayscale. Then test on live webcam images (just don’t provide an input image).
2. Write the GPU version. Test on still images and on the webcam stream.
3. Compare the CPU and GPU run times on still images. Average the run times over `repeats`  $\geq 1$  repetitions and experiment with different values of `repeats`. For the GPU version, first measure all operations, and then only the kernel executions excluding `alloc/free/memcpy`. What do you observe?
4. Experiment with several different block sizes for the kernel launch, starting with `(32, 8, 1)`. Make sure that the overall number of threads per block is a multiple of 32. For which block size is the run time minimal?

---

<sup>1</sup>/work/sdks/cudacurrent/samples/1\_Uutilities/deviceQuery

## Exercise 4: Linear Operators

(4P)

Write code for computing the gradient of an image and the divergence of a vector field. Combine both kernels to compute the pixelwise norm of the Laplacian  $\Delta u = \text{div}(\nabla u)$ :

$$\|\Delta u(x, y)\|_2 = \sqrt{\sum_{c=1}^{n_c} \Delta u_c(x, y)^2}$$

Write only a GPU version. As usual, write your code for a general  $n_c$ . Implement this in several steps:

1. Write a kernel which computes the gradient  $v^1 := \partial_x^+ u$  and  $v^2 := \partial_y^+ u$  given an input image  $u$ . The images  $v^1$  and  $v^2$  have the same number of channels as  $u$ , and  $\partial_x^+$  and  $\partial_y^+$  are applied channelwise.
2. Write another kernel which computes the divergence  $w := \partial_x^- v_1 + \partial_y^- v_2$  of a given vector field  $v$ . The image  $w$  has the same number of channels as  $v_1$  and  $v_2$ . The operators  $\partial_x^-$  and  $\partial_y^-$  are applied channelwise.
3. Write a third kernel which calculates at each pixel  $(x, y)$  the  $\ell_2$ -norm across the color channels:

$$\|u(x, y)\|_2 = \sqrt{\sum_{c=1}^{n_c} u_c(x, y)^2}.$$

4. Finally combine all three kernels to compute the absolute value of the Laplacian. Visualize the result.

## Exercise 5: Convolution

(6P)

Implement the convolution  $G_\sigma * u$  of an input image  $u$  with a Gaussian kernel  $G_\sigma$ . Use GPU global memory for everything.

1. Compute the kernel  $k := G_\sigma$  on the CPU. Normalize so that the values sum up to 1. For a general variance  $\sigma > 0$  set the kernel window radius to  $r := \text{ceil}(3 \times \sigma)$  (i.e. round up).
2. Visualize the kernel using OpenCV. For visualization, define a copy  $k'$  which is equal to the kernel  $k$  but is scaled so that the maximum value is 1. Note that the kernel can be visualized as a grayscale image with width = height =  $2r + 1$ .
3. Compute the convolution  $k * u$  on the CPU. The convolution is done channelwise on  $u$ . When the convolution requires values of  $u$  in pixels outside of the image domain, use clamping. Visualize the result.
4. Copy the kernel  $k$  computed in step 1 from the CPU to the GPU memory. Compute the convolution  $k * u$  on the GPU. Use a single kernel execution to process all channels. Visualize the result.
5. Experiment with different values of  $\sigma$  on still images, compare the run times.
6. Test on webcam images.