

Lab 8

In this lab, you will be implementing some simple image processing commands in CUDA. Go to colab.research.google.com and upload the given project (the file with extension `.ipynb`). Follow the directions through the first two code blocks (changing your runtime and then running the code cells one at a time).

Hello World! The first really interesting code cell in the notebook is the third one. It runs the “Hello World” program discussed in class. Run this example to check that everything is working; it should print out a series of messages, one per thread.

Spend a bit of time playing with the program to better understand how it works. In particular, I suggest that you change the numbers in the `<<< >>>` in `main`. These control the numbers of blocks and threads per block respectively.

Removing Red. Once you’re comfortable with the “Hello World” program, it’s time to begin the main part of lab. In preparation for this, you need to upload some files that we use for this lab. These are available on Google Classroom (the entry for this lab). The files are the library code that manages images in the `.ppm` file format (`ppmFile.h` and `ppmFile.c`) and also the image that you want to process; the sample I provide is `640x426.ppm`.

Now we’re ready to work with images in the last code cell in the notebook. It starts with the first version of our image processing program. This version removes all the red from our sample image (`640x426.ppm`) and creates a new file `out.ppm`. (You may have to refresh the file panel to the left by hitting the button with circular arrow to see `out.ppm`.) Download this file and observe that it looks basically the same as our original image, but the colors are slightly shifted. (If you’re on a Mac, Preview will open the file. If you’re on a Windows system, upload the file to photopea.com to view it.) The image is stored as an array of RGB values, each color in a separate channel of the pixel.

Let’s look through the kernel to see how it works; it’s the function `kernel`. This kernel is called once for each pixel of the input image. First, the variables `i` and `j` are set to the x and y coordinates of the pixel for which this call is responsible. The kernel calls are grouped into a 2D grid of blocks and each block is itself 2 dimensional. The values of `blockIdx.x` and `blockIdx.y` give the coordinates of the block within the 2D grid of blocks. Each block has dimensions `blockDim.x` by `blockDim.y`. The values of `threadIdx.x` and `threadIdx.y` are the coordinates of the thread within its block.

The image itself is stored linearly as triples of integers for each pixel. The three integers representing a pixel are called the *red channel*, the *green channel*, and the *blue channel* respectively because they store the amount of red, green, and blue in their pixel. The given code simply changes the value in the red channel to 0, leaving the other channels unchanged. This has the effect of removing all the red from the image.

Grayscale. As a first task, let’s change this kernel to convert the image into grayscale (black and white). To do this, take the values of the red, green, and blue channels and average them (add them up and divide by 3). Set the value of all three channels to be this average. Run the resulting program and verify that `out.ppm` is now a grayscale image before proceeding.

Blurring. Next, you will modify the kernel to create a blur effect. This is also done by taking an average, but a different sort of average. Each channel is blurred separately— each gets the average of that channel’s value for nearby pixels. Specifically, the red value of a pixel should be set to the average of the original red values for the 3x3 group of pixels centered on p . Similarly, the green value of p should be the average of their green values and the blue value of p should be the average of their blue values. Pixels on the boundary of the image don’t have neighbors in all directions; just leave these pixels unchanged.

At the end of the period. Your code will disappear when you close Colab. In order to be able to come back to it (or turn it in), copy the code of your kernel into a text file. Then, if you want to return, start with the original notebook and follow the earlier steps to set it up again before pasting in your modified kernel function.