GPU Programming in Computer Vision: Day 2

Date: Tuesday, 9. September 2014

Please work in groups of 2–3 people. We will check your solutions tomorrow after the lecture. Please be prepared to present your solution and explain the code. The general code requirements from exercise sheet 1 still apply. The bonus exercises are not mandatory.

Exercise 7: Convolution (continued) (11P)

Output: same number of channels as input image. Input: general number of channels.

In exercise 6 of sheet 1 you have implemented the convolution $G_\sigma * u$ using *global memory* for u, and global memory for the kernel $k := G_\sigma$.

1. Finish exercise 6.

Implement the convolution $G_\sigma * u$ on the GPU, now using:

2. *Shared memory* for the image u, but still global memory for k. To solve this task, use shared memory in the following way:

   (a) To compute the convolution in an output pixel (x, y) one needs values of the input image in (x, y) and also in the neighboring pixels. For each kernel block, load into shared memory all input image values needed to compute the convolution in every pixel of this block. Note that overall there are *more* input image values to load than the size of the block, see Figure 1 on the next page. Set the size of the shared memory array accordingly.

   (b) When loading into shared memory, make the read accesses to global memory as coaleasced as possible (for instance, wherever possible, neighboring threads should access neighboring memory regions).

   (c) At image borders, use clamping.

   (d) Don't forget to use `__syncthreads()` after you've finished loading the data into shared memory.

   (e) The actual computation of the convolution may only use data from the shared memory (no global memory accesses allowed).

   (f) Use dynamic allocation of shared memory. Make sure you allocate exactly the right amount of shared memory for your kernel, and not more than needed.

   (g) For multi-channel images, apply the above procedure in a loop (within the kernel) separately for each channel. When you start processing each new channel, also synchronize *before* you begin writing to the shared memory.

3. *Texture memory* for the image u, but still global memory for k.

   Since we work with multi-channel images, but CUDA textures allow only one channel, define the texture as having width w and height $h \cdot n_c$. Here we use the fact that

the channels are arranged in memory one after another. Therefore we can view a multi-channel image u as an n_c-times larger grayscale image u_1, defined by $u(x, y, c) = u_1(x, y + h \cdot c)$.

4. Pick some memory variant for the input image u (global, shared, or texture), and use *constant memory* for the kernel k. To define the CUDA constant kernel array, assume a maximal kernel radius $r_{\max} = 20$.

5. For some fixed $\sigma > 0$, compare the run times for the different memory versions. Which combination is the fastest? How much faster is it compared to the CPU version?


Figure 1: Shared memory array for convolution

Exercise 8: Structure Tensor (5P)

Output: three grayscale images. Input: general number of channels.

For an input image u, the smoothed version is defined as $S := G_\sigma * u$. The *structure tensor* T of u is defined at each pixel (x, y) as the smoothing

$$T := G_\sigma * M$$

of the matrix

$$M := \nabla S \cdot \nabla S^\top = \begin{pmatrix} (\partial_x S)^2 & (\partial_x S)(\partial_y S) \\ (\partial_x S)(\partial_y S) & (\partial_y S)^2 \end{pmatrix},$$

where $\sigma > 0$ is a scale parameter. The entries of M are scalar products over the n_c channels:

$$\left((\partial_x S)^2\right)(x, y) := \sum_{c=1}^{n_c} (\partial_x S_c)(x, y)^2, \quad \left((\partial_y S)^2\right)(x, y) := \sum_{c=1}^{n_c} (\partial_y S_c)(x, y)^2,$$

and

$$\left((\partial_x S)(\partial_y S)\right)(x, y) := \sum_{c=1}^{n_c} (\partial_x S_c)(x, y) \cdot (\partial_y S_c)(x, y).$$

Compute the structure tensor. Reuse your kernels for the convolution, don't write new kernels for this. You can use just the global memory for everything for convenience.

Implement this in several steps:

1. Compute $S = G_\sigma * u$.

2. Compute $v^1 := \partial_x S$ and $v^2 := \partial_y S$ using the *more rotationally symmetric* derivative discretizations $\partial_x^r, \partial_y^r$ as given in the lecture. Note that v^1, v^2 and S each have n_c channels.

3. Compute the matrix M at each pixel. The output should consist of three grayscale images m_{11}, m_{12}, m_{22}, corresponding to the three independent components of M at each pixel.

4. Compute $T = G_\sigma * M$ by convolving the three grayscale images m_{11}, m_{12}, m_{22}.

5. Visualize the grayscale images m_{11}, m_{12} and m_{22}. For this, you will need to define three new output images in the code framework.
   Hint: You will need to scale up these images, otherwise they will appear too dark because m_{11}, m_{12}, m_{22} are usually very small. To multiply an OpenCV `cv::Mat` image `m` by a scalar factor `f`, use `m *= f;`

Exercise 9: Feature Detection (3P)

Output: color. Input: general number of channels.

Detect edges and corners in an image using the structure tensor T from the previous exercise:

1. Write a `__device__` function that computes the eigenvalues of a 2×2 matrix. [1]

2. Compute at each pixel (x, y) the two eigenvalues λ_1, λ_2 of the structure tensor. We use the convention that $\lambda_1 \leq \lambda_2$.

   - If a pixel is on a corner $(\lambda_2 \geq \lambda_1 \geq \alpha)$, mark it red.
   - If a pixel is on an edge $(\lambda_1 \leq \beta < \alpha \leq \lambda_2)$, mark it yellow.
   - Otherwise, make the pixel darker by multiplying its components with 0.5.

3. The scale parameter $\sigma > 0$ and feature parameters $\alpha > \beta > 0$ depend on the input image u. Experiment with different values. As a start, choose $\sigma = 0.5$, $\alpha = 10^{-2}$ and $\beta = 10^{-3}$, for the input image `gaudi.png`. Test also on live webcam images.

Exercise 10 (Bonus): Convolution using cuFFT (4P)

Output: same number of channels as input image. Input: general number of channels.

The convolution theorem states that the Fourier transform of a convolution (with cyclical boundary conditions) is the pointwise product of Fourier transforms. In this exercise, we will employ the cuFFT[2] library to efficiently compute the convolution $G_\sigma * u$ using Fast Fourier Transforms (FFTs):

$$G_\sigma * u = \mathcal{F}^{-1}\{\mathcal{F}\{G_\sigma\} \cdot \mathcal{F}\{u\}\}$$

Here $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the Fourier transform and its inverse, and $\cdot$ denotes the componen-twise (complex) product. Proceed in the following way:

---

[1] http://www.math.harvard.edu/archive/21b_fall_04/exhibits/2dmatrices/index.html
[2] https://developer.nvidia.com/cufft

1. Allocate memory for the complex images $\mathcal{F}\{G_\sigma\}$ and $\mathcal{F}\{u\}$. Since the result of the Fourier transform is complex, use the type `float2`. Note that the kernel has to have the same dimensions as the image, should be centered at $(-1, -1)$ and wrap around the borders.

2. Read up on the internet about the cuFFT library and look up the documentation for `cufftPlanMany`. Set up the FFT and inverse FFT plan using that function. Set the sizes and embed parameters to $\{h, w\}$, stride to 1, dist to $h \times w$, the type to either `CUFFT_R2C` or `CUFFT_C2R` and batch to the number of channels n_c.

3. Transform both the input image and the kernel into the Fourier domain using `cufftExecR2C`. Write a kernel which performs the pointwise complex multiplication. Normalize the result by dividing both the real and imaginary part by $w \times h$. Finally transform the result back using `cufftExecC2R`.

4. For which kernel sizes does the FFT-based convolution perform favorably to your implementation of Exercise 7?