GPU Programming in Computer Vision: Day 4

Date: Thursday, 14. September 2017

Please work in groups of 2–3 people. We will check your solutions tomorrow. 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 13: Denoising (8P, *Bonus*)

Compute a minimizer $\widehat{u}$ of the *denoising energy* with Huber-regularization,

$$\min_u \int_\Omega (u-f)^2 + \lambda\, h_\varepsilon\big(|\nabla u|\big)\; dx, \qquad \text{with} \quad h_\varepsilon(s) := \begin{Bmatrix} \frac{s^2}{2\varepsilon} & \text{if } s < \epsilon \\ s - \frac{\varepsilon}{2} & \text{else} \end{Bmatrix},$$

by solving the corresponding Euler-Lagrange equation:

$$\frac{\partial E}{\partial u} = 2(u-f) - \lambda \operatorname{div}\Big(\widehat{g}\big(|\nabla u|\big)\nabla u\Big) = 0, \qquad \widehat{g}(s) := h'_\varepsilon(s)/s = \frac{1}{\max(\varepsilon, s)}.$$

1. Add Gaussian noise to your input image, with some noise variance $\sigma > 0$. For this, right after loading the input image, use the `addNoise` function from `aux.h`. You can set the noise level to e.g. $\sigma = 0.1$.
2. Compute the minimizer u by implementing the *Jacobi method* in several steps:
   (a) Compute $g = \widehat{g}(|\nabla u|)$. Reuse your code from exercise 10.
   (b) Compute the update step for u using the discretization from the lecture.
   (c) Compute N iterations and visualize the result. Test with different values λ. Check experimentally how many interations you need to achieve convergence (i.e. until visually there are no significant changes anymore).
3. Compute the minimizer u by implementing the *red-black SOR method* in several steps:
   (a) Compute g as above.
   (b) Compute the update step for u.

   The red and black updates are essentially the same, so write *only one* kernel, which will compute either the red-update or the black-update depending on a parameter. Note that you now also have a parameter $0 \le \theta < 1$ for the SOR-extrapolation.

   Hint: The SOR update step is essentially the same as for Jacobi. The only change is that there is an additional θ-extrapolation, and that the update is performed only in certain pixels.

   (c) Compute N iterations and visualize the result. Test with different λ values, and also with different values of $0 \le \theta < 1$.
   (d) Check how many interations you need for convergence. Do you observe a speed up when using SOR, compared to the Jacobi method?

Exercise 14: Parallel Reduction & Energy Computation (4P)

Write code to compute the sum of a float array u with n elements:

$$s = \sum_{i=1}^{n} u_i$$

Do so by performing a parallel reduction. In particular, complete the following tasks:

1. Implement a parallel reduction algorithm yourself. Calculate the sum of a `float` array with $n = 10^5$ elements to test it.

2. Use the library function `cublasSasum`. Compare the performance of your implementation with the performance of the library function.

Exercise 15: Histograms (3P)

Compute the intensity histogram with 256 bins of the input image using atomic operations.

1. Represent the histogram as an integer array of length 256 and initialize it to zero. Write a kernel where each thread performs a global memory `atomicAdd` on the bin corresponding to the intensity of that pixel.

2. Use the `showHistogram256` function to visualize the histogram.

3. Think of a way to improve the performance of your kernel by using shared memory and shared memory atomics. Compare your solution to the naive version from 1.