

Summer Semester 2024

Assignment on Massively Parallel Algorithms - Sheet 3

Due Date May 19, 2024

Exercise 1 (Histogram, 3 Credits)

In class, you have learned the histogram algorithm (which uses atomic operations).

- a) What is the worst-case input? (in the sense that the GPU algorithm will take the longest time)
- b) What is the best-case input?
- c) In the best-case input, what is the probability that any two threads access the same memory location? Consider 1024 bins and 64 threads and 1 warp (that means all 64 threads operate in lockstep). It could help to think of the reverse problem of no collision.

Exercise 2 (CUDA: Matrix Vector Multiplication, 5 Credits)

Start from the framework `matrixVectorMul`. There, matrix `A` is stored in row major order.

- a) Implement a matrix vector multiplication kernel for the above matrix stored in row major order. (2 Credits)
- b) Implement the conversion of the above matrix in column major order before copying its elements to the GPU (create a new variable instead of modifying `h_A` for the tests to work) and the modify the above matrix vector multiplication kernel to handle matrices stored in column major order. (2 Credits) You can either comment out the old code in the kernel for 2a) or copy the kernel, modify it for 2b and comment out the old 2a kernel so that there's no shadowing of definitions.
- c) Compare run times between the above two implementations (row- vs column major order) for different matrix sizes and provide the arguments for the differences/similarities between run times for these two implementations. (1 Credit)