



Review report of a final thesis

Reviewer: Ing. Tomáš Oberhuber, Ph.D.
Student: Emil Eyvazov
Thesis title: Gauss-Jordan Solver of Linear Equation Systems on GPU
Branch / specialization: Computer Science
Created on: 2 June 2021

Evaluation criteria

1. Fulfillment of the assignment

- [1] assignment fulfilled
- [2] assignment fulfilled with minor objections
- [3] assignment fulfilled with major objections
- ▶ **[4] assignment not fulfilled**

In my opinion, the author did not fulfilled the assignment. I tried to run the program made by the author and I obtained several error message. Obviously the results obtained by algorithms for CPU and GPU were completely different. From a source code, one can see that the algorithms modifies just a matrix of the linear system but not its right hand side (the array `d_values_out` is not handled well, the CPU algorithm does not work with any right hand side at all). Because of this, no linear system is solved and one cannot say if the implemented code works correctly. This is the main reason why I cannot consider the assignment to be fulfilled. It also seems to me that the author did not understand the principles of programming GPUs using CUDA. The algorithm executes CUDA kernels with only one or two CUDA blocks which leads to insufficient paralellism which is required for efficient utilization of GPU. Instead of it, the author uses CUDA streams. However, maximal number of concurrently running CUDA streams is limited to 32 only. At the end, tha author reports unrealistic speed-up. The author does not specify on what hardware such speed-up was achieved.

2. Main written part

50/100 (E)

The written part is well structured. However, the quality of english langauge and grammar is rather low. There are statements which does not make sense like "The goal of the project is to implement System of Linear Equations" (p. 1 - Introduction) or "the system are used..." or "sentences" without verbs like "SLE with n unknowns and m equations." References are strangely organized into sections and the author refers the sections instead of the references which is very confusing. There are misleading statements like "every CUDA core has L1 cache". Derivation of time complexity on page 11 is not correct.

3. Non-written part, attachments

30/100 (F)

The source code written by the author is rather hard to understand. It seems to me that the author is not a matured programmer yet. As I have mentioned already, the program does not solve any system of linear equations, it just modifies the matrix of the linear system and not its right hand side. It just does not make sense this way. Without solving any specific linear system one can never say if the algorithm works correctly. Concerning the computational study, the author does not specify what hardware he used and so it is impossible to validate the obtained speed-up.

4. Evaluation of results, publication outputs and awards

20/100 (F)

As I have mentioned already, the program does not solve any linear system and so one cannot validate its correctness. The author does not deal with pivoting which is very important modification of the Gaussian elimination method. Without this modification, the basic algorithm is numerically unstable. I do not see any contribution in the work of the author.

The overall evaluation

20/100 (F)

The program made by the author reports error messages indicating that the algorithm does not work well. The author does not solve any linear system because his algorithm does not modify the right-hand side of the linear system. Thus the algorithm as it is presented does not make any sense and it is questionable how the author was testing the correctness of his algorithm. The algorithm, as it is described by the author, uses just one or two CUDA blocks for one CUDA kernel which must lead to insufficient parallelism and inefficient use of GPU. At the end, the author reports unrealistic speed-up on unspecified hardware. There are number of grammatical and typographical errors in the text and several incorrect or misleading statements in the written part.

Questions for the defense

1. How did the author test the correctness of results obtained by his algorithm?
2. Why does the author set the CUDA block size to 1024 threads, did he try other numbers? Why did he choose 2D indexing of threads? Did he try 1D indexing?
3. Why does the author execute only one or two CUDA blocks in one CUDA grid?
4. How is the right-hand side of the linear system treated in the program?
5. On what hardware did the author measure the speed-up of GPU version of the algorithm?

Instructions

Fulfillment of the assignment

Assess whether the submitted FT defines the objectives sufficiently and in line with the assignment; whether the objectives are formulated correctly and fulfilled sufficiently. In the comment, specify the points of the assignment that have not been met, assess the severity, impact, and, if appropriate, also the cause of the deficiencies. If the assignment differs substantially from the standards for the FT or if the student has developed the FT beyond the assignment, describe the way it got reflected on the quality of the assignment's fulfilment and the way it affected your final evaluation.

Main written part

Evaluate whether the extent of the FT is adequate to its content and scope: are all the parts of the FT contentful and necessary? Next, consider whether the submitted FT is actually correct – are there factual errors or inaccuracies?

Evaluate the logical structure of the FT, the thematic flow between chapters and whether the text is comprehensible to the reader. Assess whether the formal notations in the FT are used correctly. Assess the typographic and language aspects of the FT, follow the Dean's Directive No. 26/2017, Art. 3.

Evaluate whether the relevant sources are properly used, quoted and cited. Verify that all quotes are properly distinguished from the results achieved in the FT, thus, that the citation ethics has not been violated and that the citations are complete and in accordance with citation practices and standards. Finally, evaluate whether the software and other copyrighted works have been used in accordance with their license terms.

Non-written part, attachments

Depending on the nature of the FT, comment on the non-written part of the thesis. For example: SW work – the overall quality of the program. Is the technology used (from the development to deployment) suitable and adequate? HW – functional sample. Evaluate the technology and tools used. Research and experimental work – repeatability of the experiment.

Evaluation of results, publication outputs and awards

Depending on the nature of the thesis, estimate whether the thesis results could be deployed in practice; alternatively, evaluate whether the results of the FT extend the already published/known results or whether they bring in completely new findings.

The overall evaluation

Summarize which of the aspects of the FT affected your grading process the most. The overall grade does not need to be an arithmetic mean (or other value) calculated from the evaluation in the previous criteria. Generally, a well-fulfilled assignment is assessed by grade A.