

Parallel Programming Concepts

WS 2013 / 2014

Assignment 4 (Submission deadline: Jan 5th 2014, 23:59 CET)

Updated description (Dec 16th 2013)

General Rules

The assignment solutions have to be submitted at:

<https://www.dcl.hpi.uni-potsdam.de/submit/>

Our automated submission system is intended to give you feedback about the validity of your file upload. A submission is considered as accepted if the following rules are fulfilled:

- You did not miss the deadline.
- Your file upload can be decompressed with a zip / tar decompression tool.
- Your submitted solution contains only the source code files and a Makefile for Linux 2.6 64-bit. Please leave out any Git / Mercurial repository clones or SVN / CVS meta-information.
- Your solution can be compiled using the “make” command, without entering a separate sub-directory after decompression.
- Your program runs without expecting any kind of keyboard input or GUI interaction.
- Our assignment-specific validation script accepts your program output / generated files.

If something is wrong, you will be informed via email (console output, error code). Re-uploads of corrected solutions are possible until the deadline.

50% must be solved correctly in order to pass the assignment. Documentation should be done inside the source code.

Students can submit solutions either **alone or as team of max 2 persons**.

Assignment 4

This assignment covers GPU Computing as an example for accelerated computing. The tasks can be implemented by using OpenCL or CUDA. Since OpenCL allows you to run your code on CPUs as well, it may be more suitable for local testing. The technical details of the test machine GPU support can be checked in the submission system.

Two out of four tasks must be solved correctly in order to pass the assignment. Documentation should be done inside the source code.

Task 4.1: Parallel Sum with OpenCL / CUDA

Implement a program that sums up a range of numbers in parallel. The general algorithmic problem is called “parallel reduction”.

Input

Your application has to be named “parsum” and accept two parameters: The the *start index* and the *end index* (64bit numbers) of the range to compute. For example, the command line

```
Example: ./parsum 1 10000000000
```

has to result in a parallel summation of the numbers 1,2,...,10.000.000.000.

Output

Your program has to produce an output file with the name “output.txt” in the same directory. This file has to contain only the computed sum.

Validation

The solution is considered correct if a true parallelized computation takes place (no Gauss please), and if the application produces correct results for all inputs. We will evaluate your solution with different summation ranges.

Task 4.2: Heat Map with OpenCL / CUDA

Implement a program that simulates heat distribution on a two-dimensional field. The simulation is executed in rounds. The field is divided into equal-sized blocks. Initially some of the blocks are cold (value=0), some other blocks are active hot spots (value=1). The heat from the hot spots then transfers to the neighbor blocks in each of the rounds, which changes their temperature value.

A round is computed as follows:

1. The value of the hot spot fields may be set to 1 again, depending on the live time of the hot spot during a given number of rounds.
2. The new value for each block per round is computed by getting the values of the eight direct neighbor blocks from the last round. The new block value is the average of these values and the own block value from the last round. Blocks on the edges of the field have neighbor blocks outside of the fields, which should be considered to have the value 0.

You have to develop a parallel application for this simulation in C / C++ using OpenCL or CUDA. The goal is to minimize the execution time of the complete simulation. Specific optimizations for the given test hardware are not allowed, since we may have to opportunity to run your code on some larger GPU hardware for the performance comparison.

Input

Your application has to be named “heatmap” and needs to accept five parameters:

- The *width* of the field in number of blocks.

- The *height* of the field in number of blocks.
- The *number of rounds* to be simulated.
- The name of a file (in the same directory) describing the *hotspots*.
- The name of a file (in the same directory) containing *coordinates*. This is the only optional parameter. If it is passed, only the values at the indicated coordinates (starting at (0, 0) in the upper left corner) are to be written to the output file.

```
Example:  ./heatmap 20 7 17 hotspots.csv
          ./heatmap 20 7 17 hotspots.csv coords.csv
```

The *hotspots* file has the following structure:

- The first line can be ignored.
- All following lines describe one hotspot per line. The first two values indicate the position in the heat field (x, y). The hot spot is active from a start round (inclusive), which is indicated by the third value, to an end round (*exclusive!*), that is indicated by the last value of the line.

```
Example content of hotspots.csv:
x,y,startround,endround
5,2,0,20
15,5,5,15
```

With such an input file, you have to run a simulation of 17 rounds on a 20x7 field with two hotspots. The first one will be located at the coordinates (5, 2) and will be active from the first round to last round (and beyond). The second hotspot will be located at the coordinates (15, 5) and will be active starting from round 5. Starting from round 15, it will no longer be active. The starting round is inclusive, the final round is exclusive. We start counting at 0. So the first hotspot will be active at round 0,1,2...18,19; the second hotspot will be active at round 5,6,7,...13,14.

```
Example content of coords.csv:
x,y
5,2
10,5
```

With such a coordinate file, you only have to provide the values at the coordinates (5, 2) and (10, 5) as part of the output file.

Output

The program must terminate with exit code 0 and has to produce an output file with the name "output.txt" in the same directory.

If your program was called without a coordinate file, then this file represents the resulting field after simulation termination. The values in the field are encoded in the following way:

- A block with a value larger than 0.9 has to be represented as "X".

- All other values must be increased by 0.09. From the resulting value, the first digit after the decimal point is added to the output picture.

Example content of "output.txt" without coordinate file

```
11112221111111111100
1112343211111111110
11124X4221111111111
11124442111111222111
1112222211111222211
1111121111111223211
01111111111111222111
```

If your program was called with a coordinate file, then this file simply represents the list of exact values requested through the coordinate file.

Example content of "output.txt" with coordinate file

```
1.0
0.03056341073335933
```

The student achieving the lowest average runtime will be announced in the lecture.

Task 4.3: Decrypt with OpenCL / CUDA

Develop an OpenCL-based or CUDA-based command line tool that performs a brute-force dictionary attack on Unix crypt(3) passwords. An example password file to be attacked is available at:

<http://www.dcl.hpi.uni-potsdam.de/teaching/parProg/taskCryptPw.txt>

Each line of the password file contains the username and the encrypted password, separated by the character ":". Your program can use the example dictionary file available at:

<http://www.dcl.hpi.uni-potsdam.de/teaching/parProg/taskCryptDict.txt>

One of the users has a password exactly matching one dictionary entry. A second user has a password build from one of the dictionary entries plus a single number digit (0-9) attached, e.g. "Abakus5".

Please note that the first two characters of the encrypted password string in *taskCryptpw.txt* are the salt string used in the original encryption process. A correct solution therefore splits the encrypted password string into salt and encryption payload, calls some *crypt(3)* implementation with the salt and all of the dictionary entries, and checks if one of the crypt results matches with an entry from the user list.

Input

Your program has to be named “decrypt” and has to take two arguments, the name of the *password file* as the first and the name of the *dictionary file* as the second command line argument.

Example: `./decrypt ../../taskCryptPw.txt ../../taskCryptDict.txt`

Output

The program must terminate with exit code 0 and produce an output file with the name “output.txt” in the same directory. This file has to contain nothing but the users whose passwords could be decrypted with the dictionary. Each line of the result file has to be a combination of username and decrypted password, separated by semicolon:

```
User01;pass
User02;Abakus5
```

Submit a compressed archive with the sources as a solution. Beside the source code, the archive can also contain a file named “taskCryptSolution.txt” with the cracked users for the example data. In this case the validation step will tell you if you found the right ones. Please do not let the validation machine perform the cracking of the example data, since this may take several hours.

Task 4.4: Beautiful Wallpapers with OpenCL

Your task is to use OpenCL to calculate a two-dimensional picture that is so beautiful, that we will change our desktop wallpapers in favor of it. (Noise, fractals, geometry, prime number,...)

Input

Your program has to be named “beauty” and has to take two arguments, the *width* and the *height* of the image that is produced.

Example: `./beauty 1920 1080`

Output

The program must terminate with exit code 0 and produce an output file with the name “output.bmp” in the same directory.

The most beautiful pictures will be shown in the lecture.