

Algorithms in Secondary Memory

Project Assignment

2018-2019

Assignment

In this assignment you are asked to implement an external-memory merge-sort algorithm (such as the one described in Section 15.4.1 of TCB), and examine its performance under different parameters. As a warm-up exercise you will need to explore several different ways to read data from, and write data to secondary memory. The overall goal of the assignment is to get real-world experience with the performance of external-memory algorithms.

Follow the outline of tasks below. You need to document your findings in a report that needs to be handed in together with your implementation. Code may be written in either C++ or Java.

1 Project outline

1.1 Reading and writing: streams

Your merge sort implementation will need to be able to sort disk files consisting of 32-bit integers. It will hence need to read data from, and write data to disk. A first step, therefore, is to develop *stream* classes that can sequentially read and write a file consisting of 32-bit integers. You will need two types of streams: *input streams* and *output streams*. The input stream should support the following operations: `open` (open an existing file for reading), `read_next`, (read the next element from the stream), and `end_of_stream` (a boolean operation that returns true if the end of stream has been reached). The output stream should support the following operations: `create` (create a new file), `write` (write an element to the stream), and `close` (close the stream).

Use the following four distinct I/O mechanisms to obtain four different implementations of the input and output streams.

1. Read and write one element at a time using the `read` and `write` system calls. These calls are available in C/C++ through the `io.h` header (if you're programming on windows in Visual C++) or through the `unistd.h` header (on Unix/Linux). If you are doing the project in Java these system calls can be mimicked. You can mimic `read` by calling `readInt()` on a `java.io.DataInputStream` that is wrapped directly around a `java.io.FileInputStream`, as illustrated by the following code snippet:

```
InputStream is = new FileInputStream( new File("input.data" ) );
DataInputStream ds = new DataInputStream(is);
ds.readInt();
```

The `write` system call can be mimicked similarly.

2. Read and write using `fread` and `fwrite` functions from the C `stdio` library which implements its own buffering mechanism for these functions. If you are doing the project in Java these functions can be mimicked. You can mimic `fread` by calling `readInt()` on a `java.io.DataInputStream` that is wrapped around a `java.io.BufferedInputStream` that itself is wrapped around a `java.io.FileInputStream`, as illustrated by the following code snippet:

```
InputStream is = FileInputStream( new File("input.data" ) );
BufferedInputStream bis = new BufferedInputStream( is );
DataInputStream ds = new DataInputStream( bis );
ds.readInt();
```

The `fwrite` function can be mimicked similarly.

3. Read and write are implemented as in step (1), except that now you equip your streams with a buffer of size B in internal memory. Whenever the buffer becomes empty/full the next B elements are read/written from/to the file.
4. Read and write is performed by mapping and unmapping a B element portion of the file into internal memory through *memory mapping*. Whenever you need to read/write outside of the mapped portion, the next B element portion of the file is mapped.

Memory mapping is available in C/C++ through the `mmap/munmap` functions of the `<sys/mman.h>` header (in Linux) and the `CreateFileMapping` function of `<windows.h>` (in Windows). Alternatively, the Boost C++ library¹ provides platform-independent interfaces to memory-mapped files.² In Java, you should use the `map` method of the `java.nio.channels.FileChannel` class.

For this particular implementation you are required to do a little research on what *memory mapping* actually is; you are expected to explain this concept in your report.

Experiment with each stream implementation and determine which implementation is the most performant. The goal is to determine which works the best, and to clearly document the process of this discovery in the report. What are the limitations and advantages of each implementation? In particular:

- Try opening k streams (on distinct files) and read/write one element to/from each stream, N times. Try for large N and $k = 1, 2, \dots, \text{MAX}$ where MAX is the minimum of (1) the maximum number of streams supported by your operating system/programming language and (2) 30.
- For implementations (3) and (4), experiment with different values of B (including very large ones). Identify the optimal values.

Don't forget to describe, in your report, the environment on which your experiments have been performed (machine type, hard disk type, operating system, total memory available, programming language used, and libraries used).

¹<http://www.boost.org/>

²It actually provides 2 such interfaces, one in the `iostreams` sublibrary, and one in the `interprocess` sublibrary.

1.2 Multi-way merge

Implement a d -way merging algorithm that, given d sorted input streams of 32-bit integers, creates a single output stream containing the elements from the input stream in sorted order. The merging should use a priority queue (e.g., a heap) to obtain the next element to be output at all times.

1.3 External multi-way merge-sort

Implement an external memory multi-way merge-sort algorithm for sorting 32-bit integers. The program should take as input an input file and the following parameters:

- M — the size of the main memory available during the first sort phase, in number of 32-bit integers;
- d — the number of streams to merge in one pass in the later sort phases;

and sort as follows. Let N be the size of the input file, measured in number of 32-bit integers.

1. Read the (single) input file and split it into $\lceil N/M \rceil$ streams, each of size at most M . This splitting itself should take care not to read more than M integers from the input file into memory at any given time.

Each stream that is created should be sorted in internal memory using heapsort, (internal-memory) mergesort, or a sorting method that has comparable $O(n \log n)$ worst-case performance before being written to disk. In particular, you can reuse the sort algorithm provided with the language you use if it meets the performance requirement.

2. Store the references to the $\lceil N/M \rceil$ streams in a queue (if necessary in external memory). Again, take care to only store the references to the streams in the queue. The content of the streams themselves should not be loaded into memory until the stream is treated by the following step.
3. Repeatedly until only one stream remains, merge the d first streams in the queue, and put the resulting stream at the end of the queue (if only $x < d$ streams remain in the queue, merge those).

Experiment with your merge-sort implementation using the best stream implementation from 1.1. Hereto, you should generate (unsorted) input files that contain randomly generated 32-bit integers. Try different values of N , M , and d . Identify what are good choices of these parameters for various input sizes (including very large ones). In case that the input file is small enough to fit in memory, compare your merge-sort implementation with a main-memory sort (e.g., heapsort, mergesort). Document everything you discover in your report.

2 Tips

- Describe in your report how you chose the values of the parameters N , B , k , M , d , that you ran your experiment with. In particular, motivate why these values will be able to tell you something about the real behavior of runtime systems. (Don't forget that a file with $N = 10^6$ 32-bit integers is only 4 MB in size, which is very small by today's standards.)

- Run your programs a suitable number of times for each input size and use the average running time. This should level out fluctuations due to other processes.
- There are quite a number of Unix commands to time the execution of the program. Among these are `time` and `timex`. In some shells (e.g. `csh`), there are also built-in versions of `time`. There are also corresponding unix routines called `times`, `getrusage`, and `gettimeofday`. Their availability may differ from machine to machine. The Windows API also provides `QueryPerformanceCounter` and `QueryFrequencyCounter` that can also be used to time the execution. Both the Boost library and the C++11 STL also provides a `chrono` class that provides timing in a platform independent way.
- Alternatively there are a number of *benchmarking* frameworks available for both C++ and Java that allow you to specify the code that needs to be benchmarked, and that takes care of running this benchmark a number of times and reporting running time averages, standard deviations, etc.
- When comparing different algorithms (or the same algorithm with different parameters), make sure to run them on the *same input file* in order to get meaningful results.

3 Report

Your report should document the progress of your project, your findings, as well as try to relate your experimental observations to what we have discussed during the theory lectures. (In particular: the cost formulas from Lecture 6.) Here is a suggested table of contents.

1. Introduction and Environment.

Briefly summarize the project goals. Document the machine environment on which the experiments were run; the programming language used; the external libraries used; the manner in which test data was generated; the size of the test data.

2. Observations on streams

(a) Expected behavior.

For each of the 4 implementations, briefly explain the implementation in your own words. Give details on how you programmed the implementation if necessary. Explain what kind of performance behavior you expect (before running any experiment) from the implementation. Hereto, define, for each of the 4 implementations, a cost formula that estimates the total number of I/Os that needs to be done, in function of N , k , and B .

(b) Experimental observations.

Discuss what you observe during your experiments. Use graphs and plots to compare the implementations. In each experiment, vary only one parameter (e.g., N) while keeping the other parameters (e.g., k) fixed.

(c) Discussion of expected behavior vs experimental observations.

When these behaviors are not the same, postulate a reasonable explanation for the difference. Give general conclusions and identify the stream implementation to use for the external multi-way merge-sort algorithm.

3. Observations on multi-way merge sort.

(a) Expected behavior.

Explain the implementation in your own words. Give details on how you programmed the implementation if necessary. Explain what kind of performance behavior you expect (before running any experiment) from the implementation. Hereto, define a cost function that estimates the total number of I/Os that need to be done, in function of N , d , and M

(b) Experimental observations.

Discuss what you observe during your experiments. Use graphs and plots to compare the implementations. In each experiment, vary only one parameter (e.g., N) while keeping the other parameters (e.g., M, d) fixed.

(c) Discussion of expected behavior vs experimental observations. Identify what good choices are of the parameters for varying input sizes.

4. Overall conclusion.

Summarize what you have learned in this project.

Careful! Do not simply copy-and-paste text from the internet to include in your report. If you copy-and-paste existing text/graphics without duly citing the source from which this was copied this constitutes fraud, which will result on 0/20 on the course and possibly an exclusion from other exams!

4 Modalities

The assignment has the following modalities:

1. This project assignment contributes 6/20 to the overall grade; the written exam contributes the remaining 14/20 points.
2. The assignment will be graded on (1) the implementation itself and (2) the report that you need to write describing both the implementation, the expected behavior, and the experimental work.
3. The assignment should be solved in groups of 3 (if the total number of students in the course is not divisible by 3, at most two groups of 2 students will be allowed). You are asked to register, per group, the names of the group members via the online poll available at

<https://docs.google.com/forms/d/e/1FAIpQLSetjHsgNHCInHcYSKFJ101QN5977mX6aNrHRubUwXDGkQ5vKg/viewform>

by October 26 at the latest. If you cannot find a partner, please indicate so by sending an email to prof. Vansummeren (svsummer@ulb.ac.be), who will hook you up with a partner.

4. This assignment is mandatory. If you do not make the assignment, you cannot pass the course in the first exam session.

5. You will have to create a GIT repository³ in the INFO-H-417/2018-2019-1 repository group at <http://wit-projects.ulb.ac.be/rhodecode/> to submit both your report and your code. The username and password to login to this system correspond to your ULB/VUB NetID. The repository will be named

`project-<student1>-<student2>-<student3>`

where `student` corresponds to your username and `<student1>-<student2>-<student3>` appear in sorted order. This repository must be made private. It is recommended that you create this repository *as soon as possible* to avoid last minute technical difficulties, and that you use it throughout the project to synchronize your changes.

Tip. If you attempt to push a large set of changes to a GIT repository with HTTP or HTTPS, you may get an error message such as `error: RPC failed; result=22, HTTP code = 411`. This is caused by a GIT configuration default which limits certain HTTP operations to 1 megabyte. To change this limit run within your local repository

```
git config http.postBuffer *bytes*
```

where `*bytes*` is the maximum number of bytes permitted.

6. Your solution should be pushed to the repository *no later than 6 January 2019*. You get a penalty of -1/6 points for each day that your solution is delayed. Only the latest commit will be considered as the solution. It is recommended, however, to finish & submit your project before the Christmas break.
7. As stated above, the project will need to use a *priority queue* and implement *heapsort* or equivalent main-memory sorting algorithm. You are only allowed to use existing code and existing libraries for these two aspects; all other code must be written by you for the project, unless explicitly stated otherwise in the assignment above. For example, you can use the `std::priority_queue` class from the C++ STL to implement priority queue (if you do the project in C++), or the Java `java.util.PriorityQueue` (if you do it in Java).

(Note that, if you decide to use a benchmarking framework, you are free to use that as a library in your code.)
8. Sharing of code or reports between groups is not allowed. (Groups may, however, verbally discuss ideas on how to tackle the project).
9. Plagiarism, in the sense of copy-pasting from existing reports or copy-pasting existing external memory merge-sort code is not allowed. In case plagiarism is detected, students risk being punished according to article 34 of the exam regulations shown in Figure 1.
10. As stated above, this project assignment is done in groups of 3. All members of the group receive the same grade on the project! Therefore, in case that a group member feels that another group member does not do his/her share of the agreed work, please contact prof. Vansummeren immediately.

³<http://git-scm.com/documentation>

Art.34 In case of fraud or plagiarism during an examination or during a test at an interim date during the academic year, or in relation with the preparation of written reports or papers, the course professor reports the case in writing prior to the jury deliberation to the relevant academic authority levels for disciplinary matters. A copy of that fraud report is addressed to the jury chairmen. The student can ask to be heard by a jury chairperson prior to the jury deliberation, in presence of the related course professor. Without prejudice to the disciplinary processes at the University Faculty level, in case of fraud the student points for the related course are brought down to 0/20. The jury further can:

- *decide to cancel the examination session;*
- *decide to refuse the student access to both examination sessions of that academic year.*

Figure 1: Excerpt of the Exam Regulations concerning fraud.