

Experimental run times of sorting algorithms

In this course, we have seen in detail several sorting algorithms: insertion sort, merge sort, heap sort, quicksort, counting sort, radix sort and bucket sort (we also mentioned selection sort, bubble sort and tree sort). These algorithms have illustrated different design techniques (iteration, divide-and-conquer, randomization, etc.), and they have different asymptotic running times as a function of the array size n and possibly of other parameters (e.g. the maximum number of different integers k in counting sort), as well as best, worst and average cases.

In several of the labs, a bonus question has been to measure the running times of sorting algorithms in a given computer. This project expands on that. It consists of repeating the following for different algorithms (counting sort, in this example):

1. Coding the counting sort algorithm in C or C++.
2. Running your counting sort code with input arrays of different size n , for different k , and when the input array is in ascending order, descending order, or unsorted. Record the run time T in each case.
3. Plotting a curve $T(n)$ as a function of n for each case (each value of k , and for ascending order, descending order and unsorted).

The objective of this exercise is to test whether the formulas predicted by the theory (e.g. “quicksort runs in $\Theta(n^2)$ in the worst case”) match well with the measured run time in an actual computer, and if they do not match well, to try to understand why and give an explanation or a guess. (You may be surprised about how well, or how badly, the theory applies.) The more algorithms you study and the more insightful your comments about their measured and theoretical run time, the higher your grade.

What you have to submit

Submit the following (as a single file `cse100.tar.gz`, by email to the TA):

1. A report `cse100.pdf` of your results including, for each algorithm you tested:
 - The theoretical run time for the best, worst and average case.
 - A plot of the runtime T as a function of n (and possibly other parameters) for ascending order, descending order and unsorted. See an example in fig. 1.
 - A brief discussion of whether the curves match with the theory and an explanation (or guess) of why they do not, if they disagree. Consider the questions: how do the curves compare with the asymptotic running times in the best, worst and average case for the algorithms? What happens for very small or very large n ? How do different parameters affect the result (e.g. different k for counting sort)? When does counting sort beat quicksort? Etc.

Also, give the basic specs for the computer that you used (processor, clock frequency, memory size, number of processors, operating system). Note: all run times must be measured with the same computer and (approximately) the same workload.

2. Your C/C++ code for the algorithms (each one in a separate subdirectory, e.g. `CountingSort/`).
3. A single Matlab file `runtimes.mat` containing your measured run times for all the algorithms you tested, from which the plots were generated.
4. If you want, repeat several times the same identical experiment (e.g. counting sort with $n = 1000$ and $k = 10$ with a sorted array) to obtain error bars (since no two runs will take the exact same time) and include them in the plot.

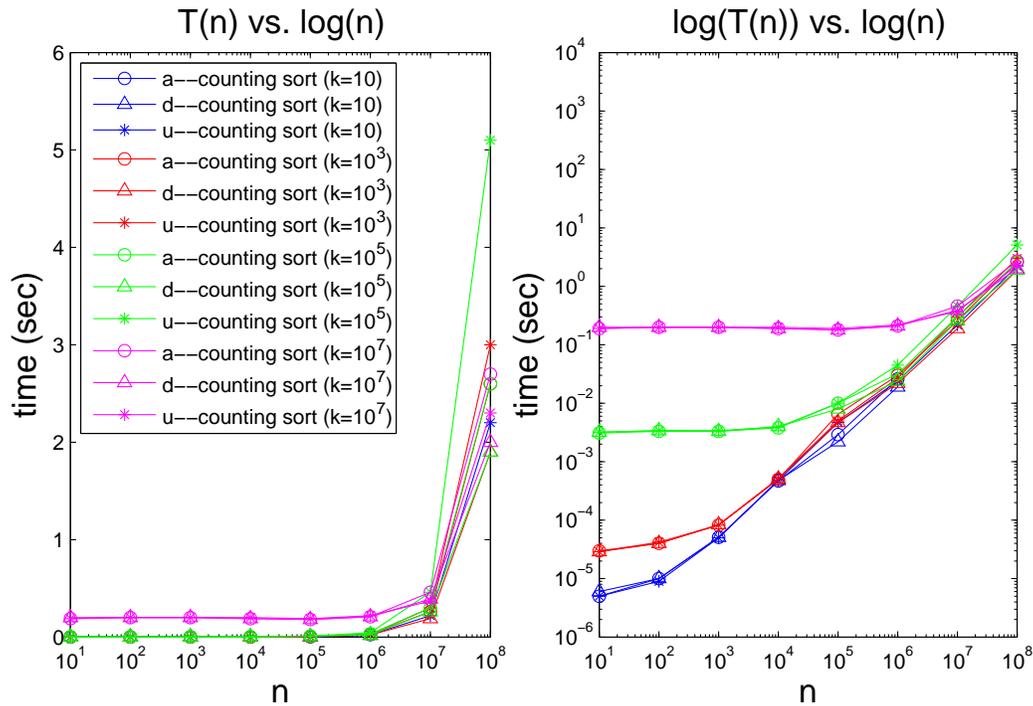


Figure 1: Experimental run times for counting sort for $n \in \{10, 10^2, 10^3, 10^4, 10^5, 10^6, 10^7, 10^8\}$, $k \in \{10, 10^3, 10^5, 10^7\}$ and a random input array of integers that is sorted in ascending order ('a'), descending order ('d') or unsorted ('u'). *Left*: semilog graph. *Right*: log-log graph.

Practical details

For all practical questions, see the TA.

- Generate all your test cases with the C function (provided by us) `rndarray(s,k,n,seed)`, which creates a random array of integers, as follows:
 - `s` (char): array sorted in ascending ('a') or descending order ('d'), or unsorted ('u').
 - `k` (int): the array values will be in $0..k$.
 - `n` (int): the size of the array to be generated.
 - `seed` (int): the seed to initialize the pseudorandom number generator. Note: fix this to your student id# for all your experiments.
- Use the C `gettimeofday()` function to measure elapsed time, by bracketing your sorting function with `gettimeofday()` calls. We have provided the template C program `engapps00:/home/cqin/bin/cse100-project/SortTime.c` for this purpose. You just need to merge your sorting function into `SortTime.c` and build the executable.
- Use the Matlab function `engapps00:/home/cqin/bin/cse100-project/PlotRuntime.m` to plot the runtime curves (modify it if you want). For example, if you have 6 curves each with 8 values of n , all you need to do is type your running times into a matrix T of 8×6 (use NaN for missing entries) and run `PlotRuntime.m`. The figure will appear on screen and on a file `runtime.eps`.
- You can automatize your experiments and save time by running your code with a Unix shellscript that loops over, say, $n = 10^2, 10^3, \dots$ and $s = \{ 'a', 'd', 'u' \}$. See an example in `engapps00:/home/cqin/bin/cse100-project/sortscript.sh`.
- We suggest you try $n \in \{10^2, 10^3, 10^4, 10^5, 10^6, 10^7, 10^8\}$, for each of $s = \{ 'a', 'd', 'u' \}$; for counting sort, $k \in \{10, 10^3, 10^5, 10^7\}$; for radix sort, $r \in \{1, 5, 10, 20\}$ (the number of bits in each digit).