

QUICKSORT USING OPENMP

By Mohd Ehtesham Shareef



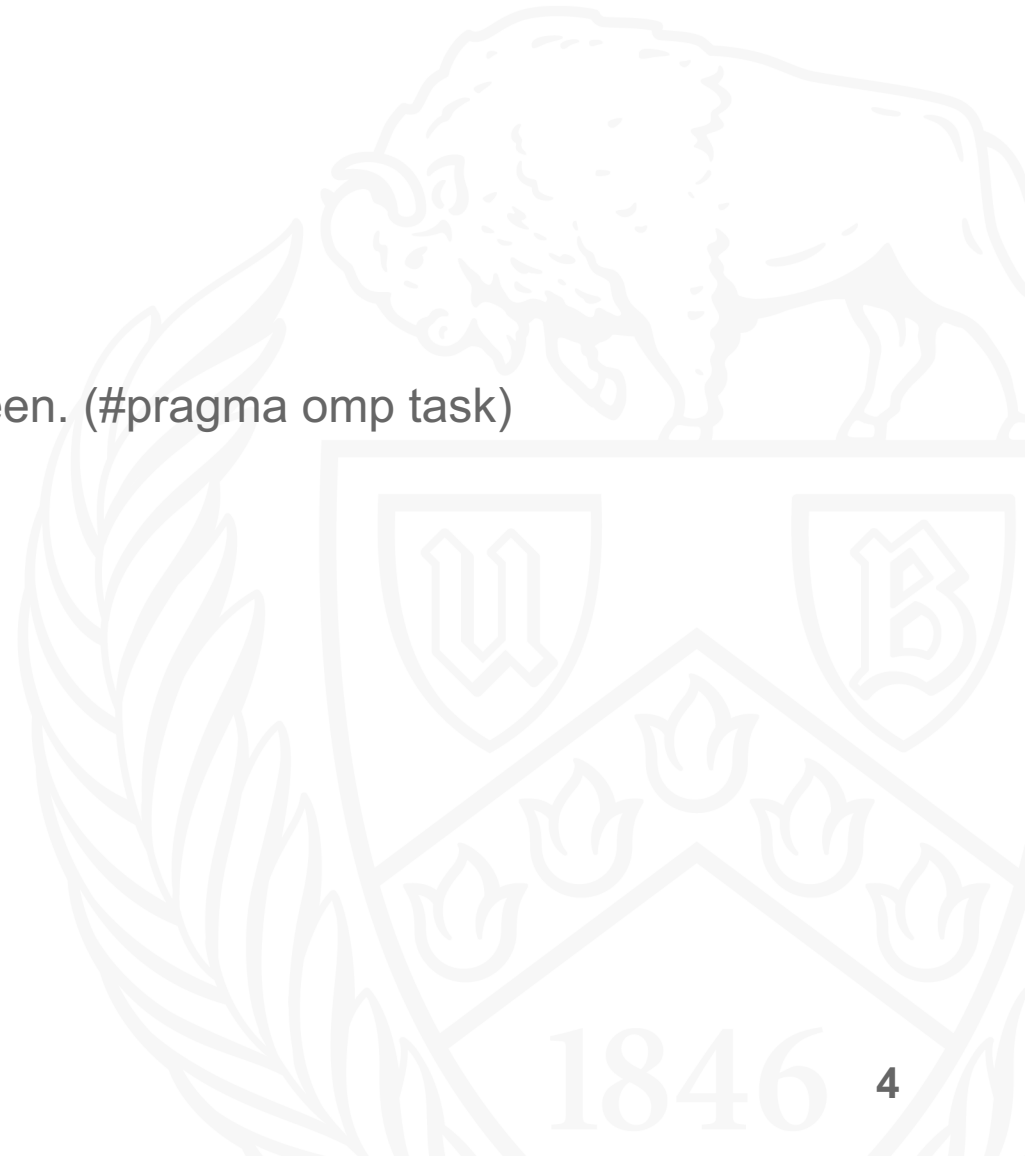
Sequential Quicksort

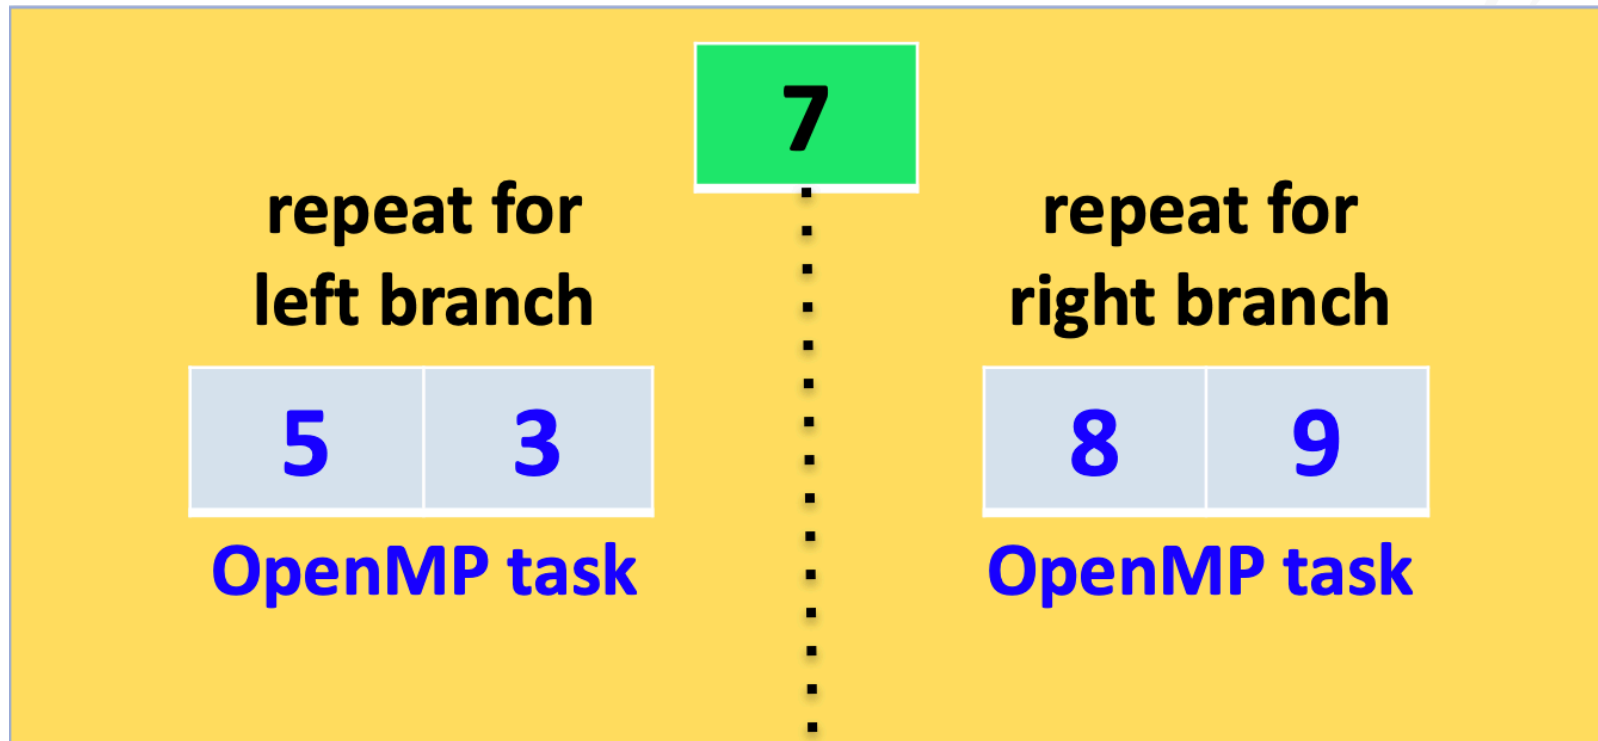
- Select median as pivot from the sample data set picked from the actual data set.
- Divide the list into two sub lists: a “low list” containing numbers smaller than the pivot, and a “high list” containing numbers larger than the pivot
- The low list and high list recursively repeat the procedure to sort themselves.
- The final sorted result is the concatenation of the sorted low list, the pivot, and the sorted high list.

8	5	7	3	9
8	5	7	3	9
8	5	7	3	9
8	5	9	3	7
8	5	9	3	7
5	8	9	3	7
5	3	9	8	7
5	3	9	8	7
5	3	7	8	9



- Divide and Conquer – Parallelism can be introduced
- Each worker thread is spawned whenever a task construct is seen. (`#pragma omp task`)



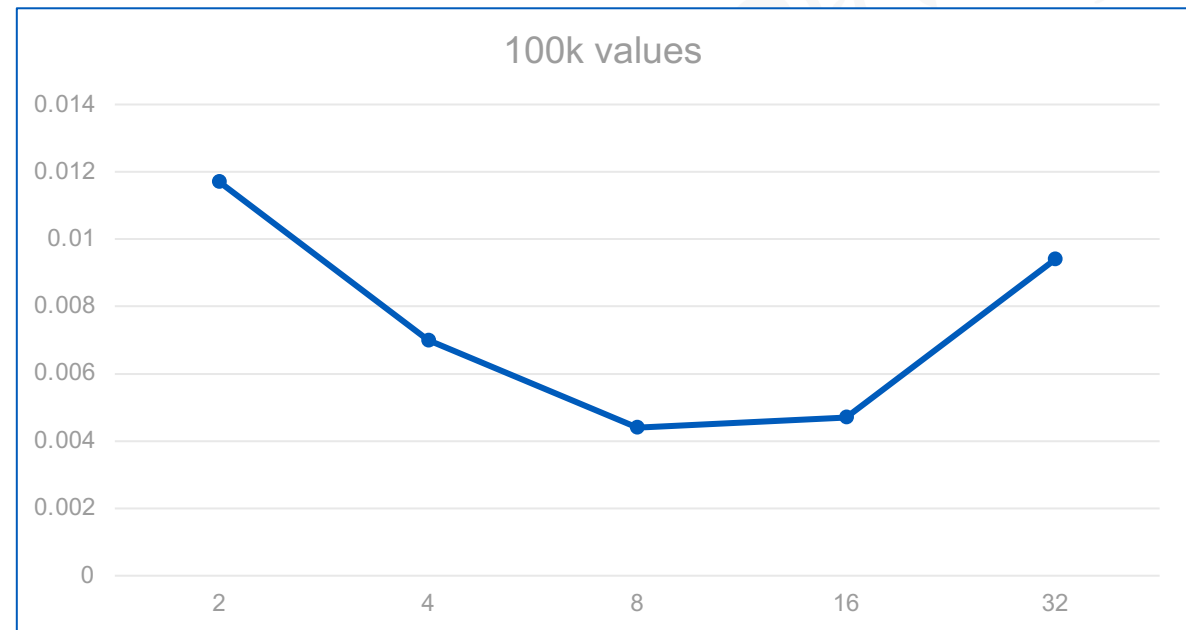


Results



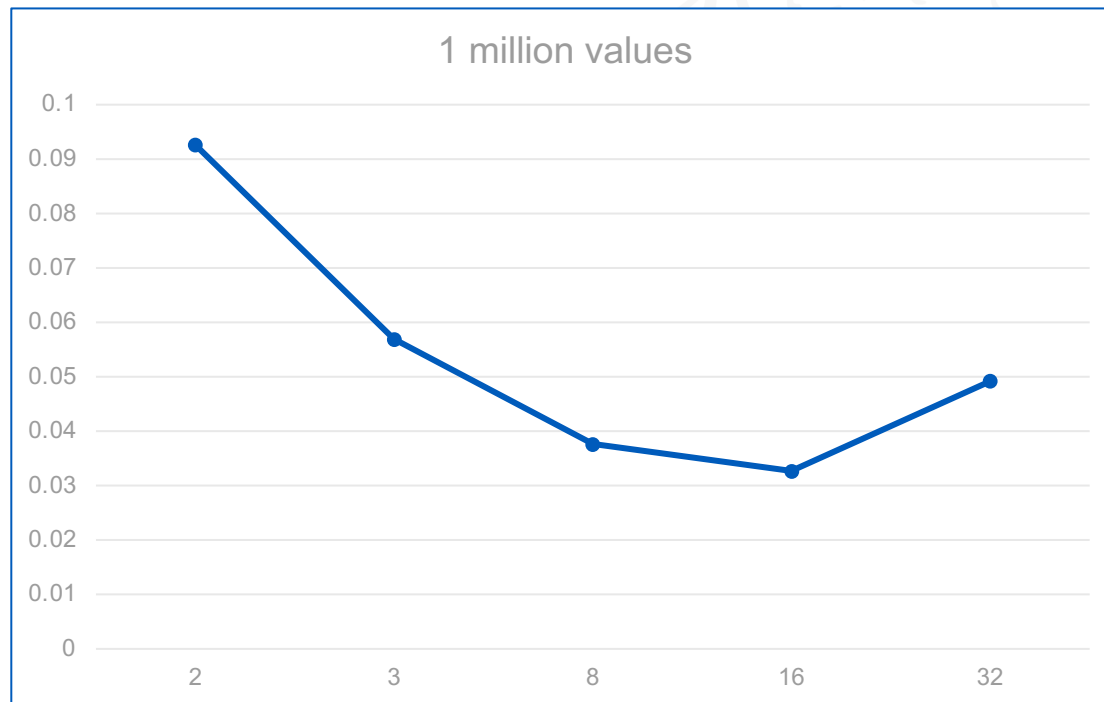
100k values

Threads	Time
2	0.0117
4	0.0070
8	0.0044
16	0.0047
32	0.0094



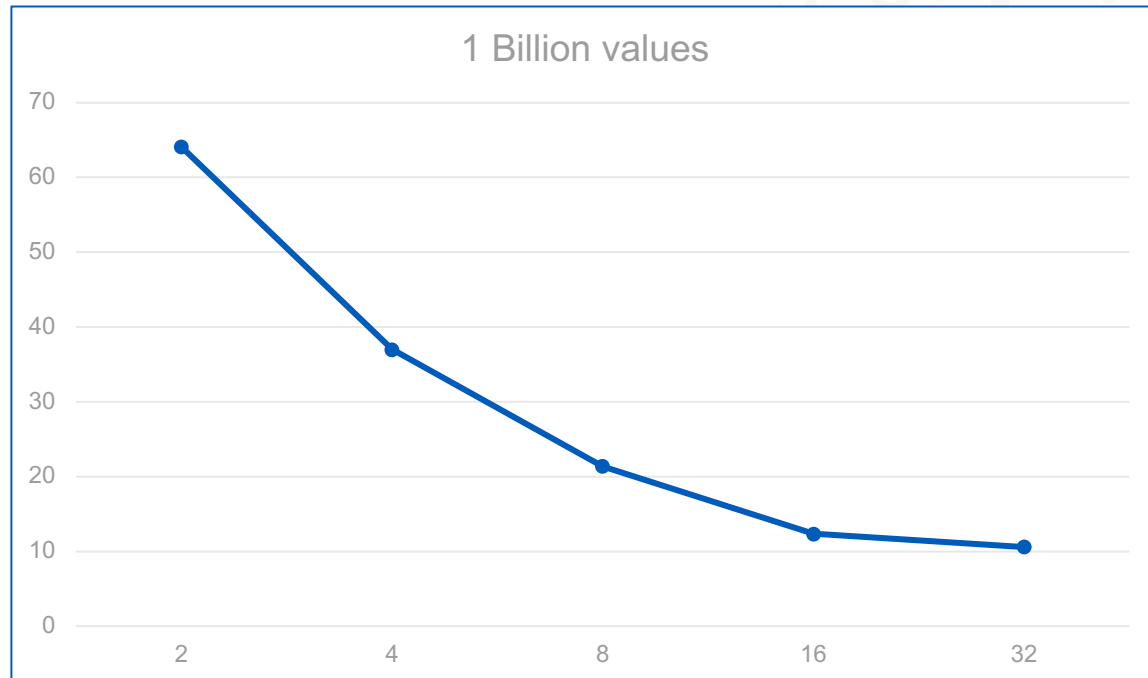
1 million values

Threads	Time
2	0.0926
4	0.0564
8	0.0376
16	0.0327
32	0.0492



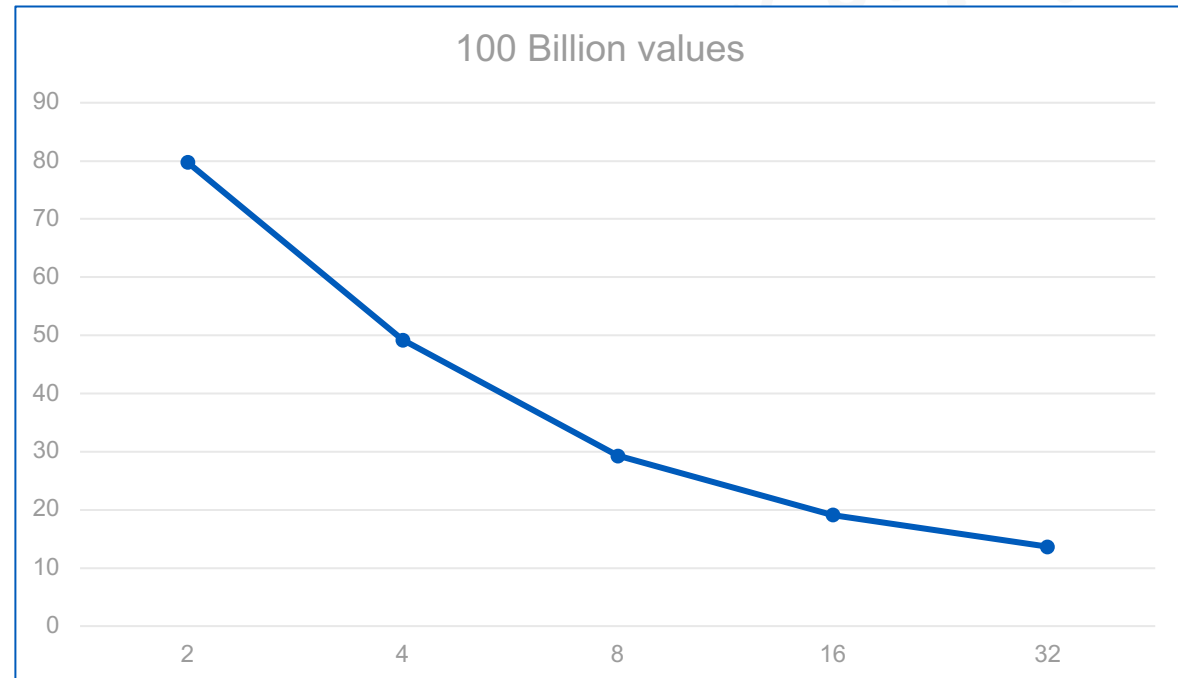
1 Billion values

Threads	Time (in s)
2	64.1
4	37.0
8	21.39
16	12.35
32	10.57



100 Billion values

Threads	Time (in s)
2	79.75
4	49.21
8	29.32
16	19.12
32	13.66



Inferences

- Performance degrades when using large number of threads when the size of the data is relatively small.
- Makes sense to increase the number of threads as the size of the data increases.
- Possible reason for performance degradation – overhead due to thread creation, thread scheduling.

References

- <https://cse.buffalo.edu/faculty/miller/teaching.shtml>
- <https://www.openmp.org/>
- <https://www.openmp.org/wp-content/uploads/sc16-openmp-booth-tasking-ruud.pdf>
- <https://stackoverflow.com/questions/16007640/openmp-parallel-quicksort>

