

CSE 702: SEMINAR ON PROGRAMMING MASSIVELY PARALLEL SYSTEMS

Learning and Implementing Odd-even
transposition sort in OpenMP

 University at Buffalo
School of Engineering and Applied Sciences



PREPARED BY:

ASIF IMRAN (UB PERSON NUMBER: 50249959)

PREPARED FOR:

PROF. DR. RUSS MILLER

UB DISTINGUISHED PROFESSOR,

DEPARTMENT OF COMPUTER SCIENCE AND

ENGINEERING, UNIVERSITY AT BUFFALO (SUNY)



Overview

- Overview of Odd Even Transposition Sort
- Discussions of Goals and Assumptions
- Obtained results
- PE's behavior
- Discussion of results
- References



Background

- Aim to enhance knowledge gained on parallel programming
- Implemented the project on OpenMP, different model than MPI, therefore obtaining detailed knowledge on different aspects of parallel programming
- Harness massively parallel computing machines at CCR
- Used OpenMP directive on machines with 2, 4, 8 and 16 cores respectively.



Consider Bubble Sort

- Bubble sort is a $O(N^2)$ sorting algorithm.
- It is simple to understand and implement.

So why discuss it?

Understandable

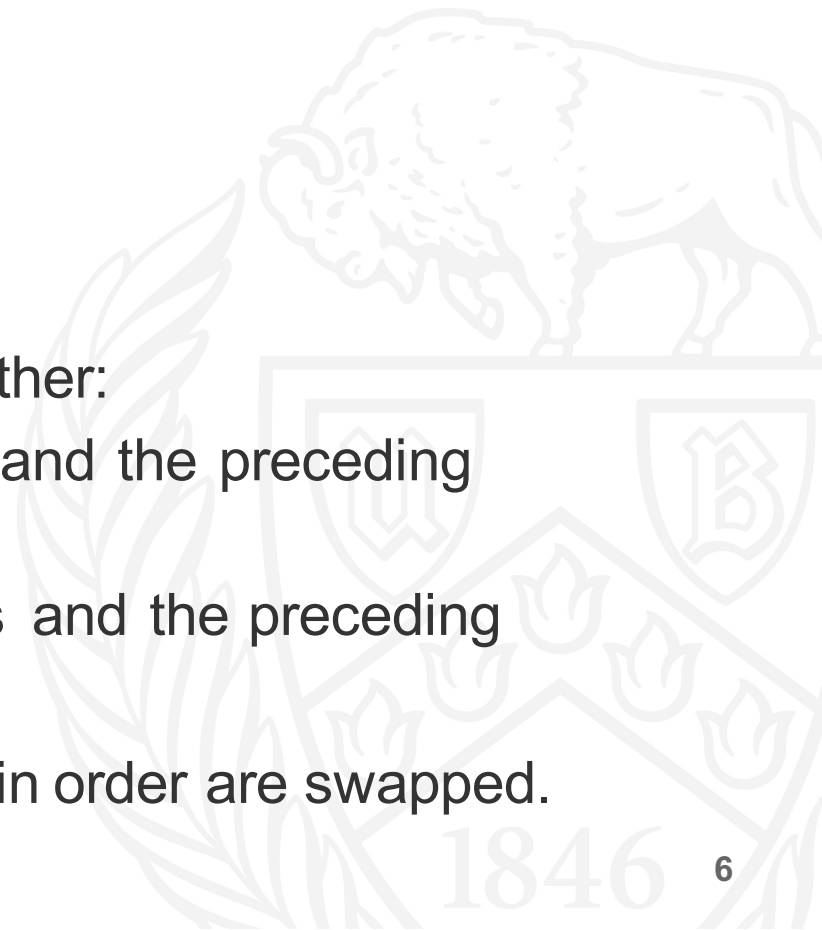
Implementable

Can be parallelized



Odd Even Transposition sort

- Parallelizable version of Bubble sort
- Requires N passes through the array.
- Each pass through the array analyzes either:
 - Every pair of odd indexed elements and the preceding element, or
 - Every pair of even indexed elements and the preceding element.
- Within each pass, elements that are not in order are swapped.




Pictorial depiction

- Even positions

$(a[0], a[1]), (a[2], a[3]), (a[4], a[5]), \dots,$

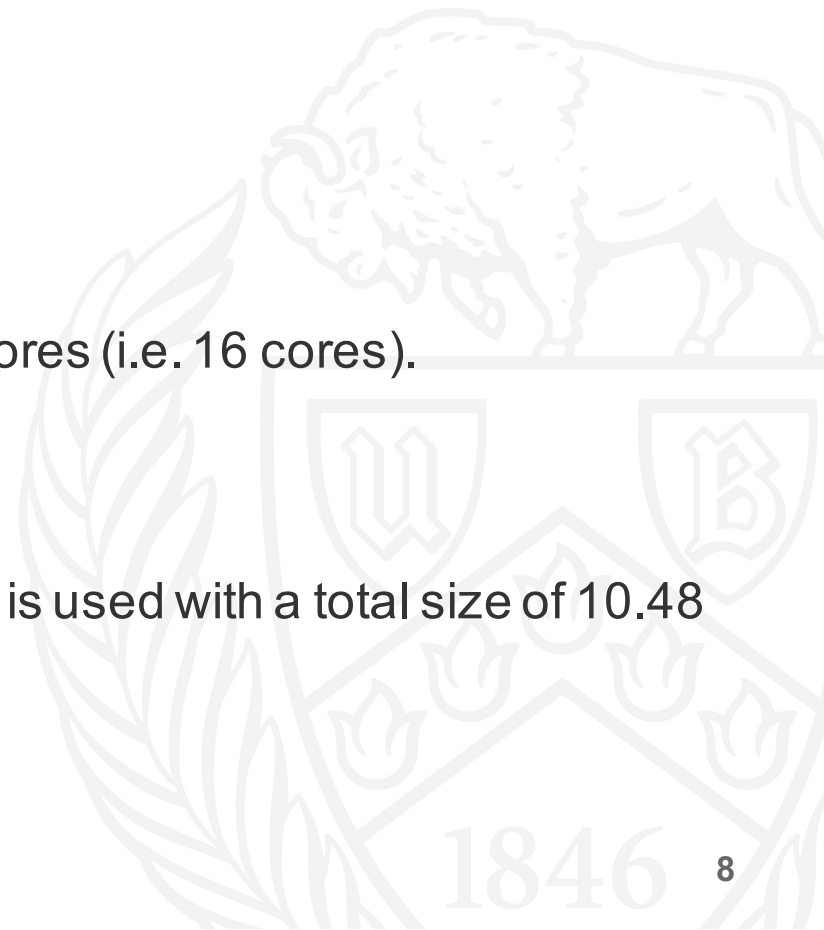
- Odd positions

$(a[1], a[2]), (a[3], a[4]), (a[5], a[6]), \dots$



Goals of this project

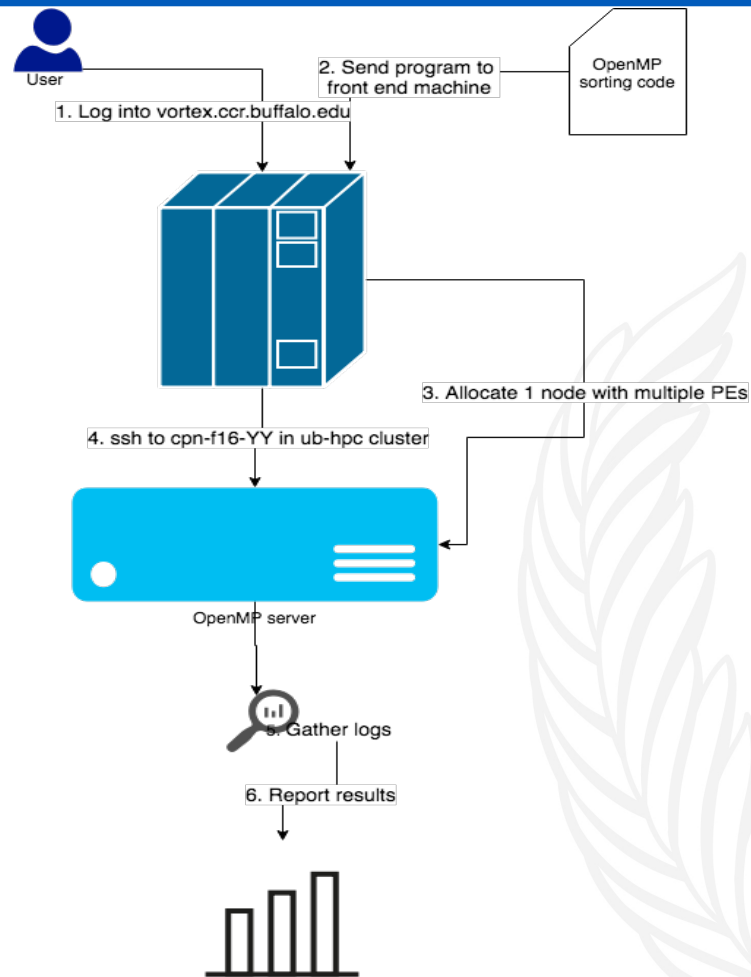
- Run OpenMP code
- One node should have multiple threads
- Aim is to allocate nodes with largest number of cores (i.e. 16 cores).
- Sorted dataset of varying sizes
- Used standard dataset
- There are 2621440 integers in the dataset which is used with a total size of 10.48 MB
- Take integer blocks of varying sizes



Goals of this project (cont)

- We took two groups of data containing sub-groups amongst them:
 - Small Group: 0.00625, 0.0125 and 0.025 million
 - Large Group: 0.05, 0.1, 0.125, 0.25, 0.50, 1.00, 1.5 and 2.00 million
- The reason for using two subgroups was to demonstrate how OpenMP performs to sort various groups of data.
- Provide Runtime graphs when data size and nodes are both doubled
- Provide Runtime graphs when data size is constant and nodes are doubled

Process flow



SLURM output

```
[asifimra@vortex2:~]$ export | grep SLURM
declare -x SLURM_CLUSTER_NAME="ub-hpc"
declare -x SLURM_JOBID="9947720"
declare -x SLURM_JOB_CPUS_PER_NODE="16"
declare -x SLURM_JOB_ID="9947720"
declare -x SLURM_JOB_NAME="bash"
declare -x SLURM_JOB_NODELIST="cpn-f16-13"
declare -x SLURM_JOB_NUM_NODES="1"
declare -x SLURM_JOB_PARTITION="general-compute"
declare -x SLURM_MEM_PER_CPU="2800"
declare -x SLURM_NNODES="1"
declare -x SLURM_NODELIST="cpn-f16-13"
declare -x SLURM_NODE_ALIASES="(null)"
declare -x SLURM_NPROCS="16"
declare -x SLURM_NTASKS="16"
```





asif — asifimra@cpn-f16-13:~ — ssh -X asifimra@vortex.ccr.buffalo.edu — 80...

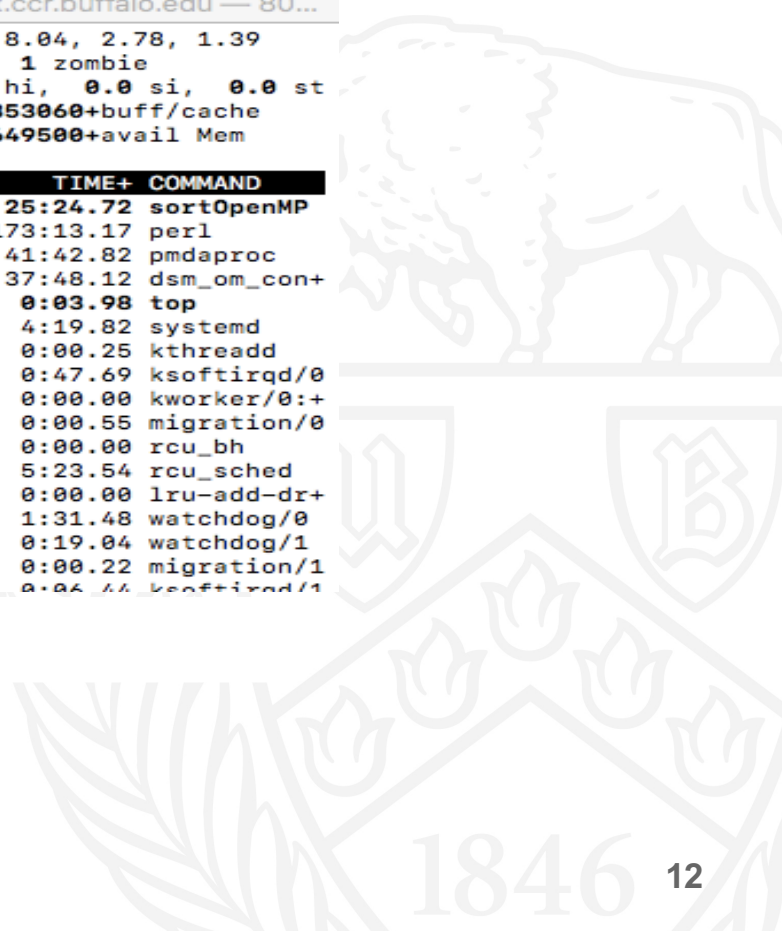
top - 23:08:37 up 18 days, 15:27, 2 users, load average: 8.04, 2.78, 1.39
Tasks: 279 total, 2 running, 276 sleeping, 0 stopped, 1 zombie
%Cpu(s): 99.6 us, 0.1 sy, 0.0 ni, 0.2 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 13181312+total, 15373364 free, 2909152 used, 11353060+buff/cache
KiB Swap: 13195673+total, 13184179+free, 114944 used. 12649500+avail Mem

Table with columns: PID, USER, PR, NI, VIRT, RES, SHR, S, %CPU, %MEM, TIME+, COMMAND. Row 17332 asifimra is highlighted with a blue box.

Table with columns: PID, USER, PR, NI, VIRT, RES, SHR, S, %CPU. Row 29309 asifimra.


Table with columns: PID, USER, PR, NI, VIRT, RES, SHR, S, %CPU. Row 37786 asifimra.

Table with columns: PID, USER, PR, NI, VIRT, RES, SHR, S, %CPU. Row 31088 asifimra.

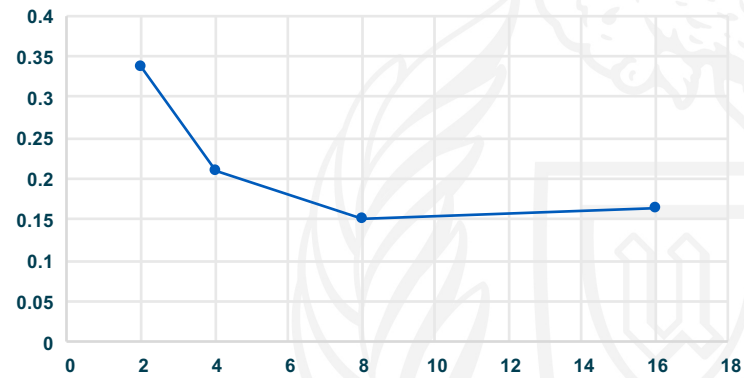


Obtained results

Running time for data size of 6250




Running time when data size is 12500




Obtained results [cont]

Run time when data size is 25000




Run time when data size is 50000




Obtained Results [cont]

Run time when data size is 100000




Run time for data size 125000




Obtained Results [cont]

Run time for datasize 250000




Run time when data size is 500000

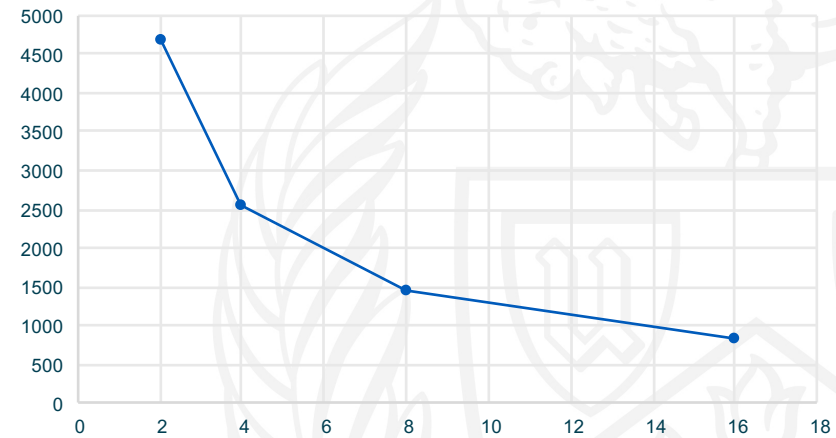


Obtained Results [cont]

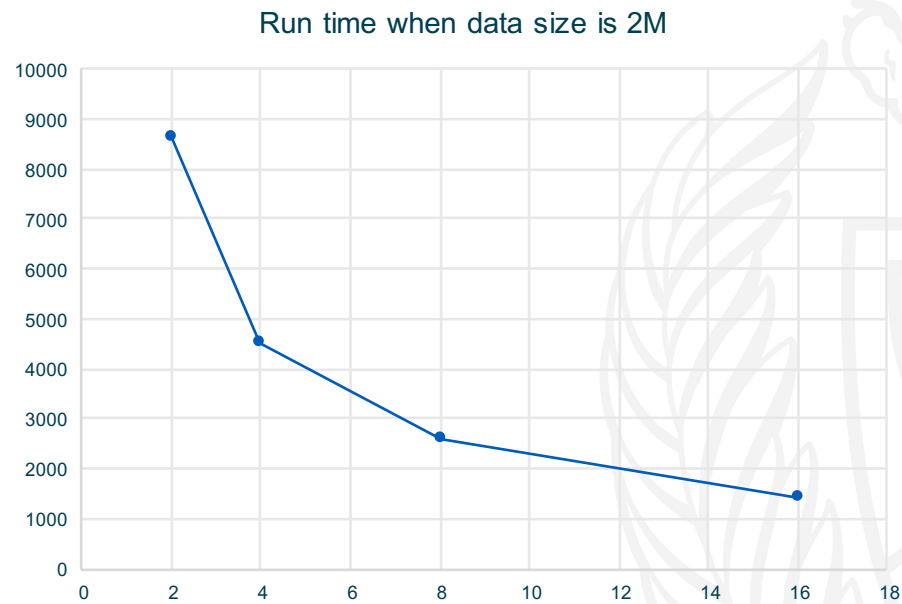
Run time for data size 1 million



Run time when data size is 1.5M




Results [cont]



Results (cont.) with 32 threads (i.e. PEs)


OpenMP runtime with 1.2M data



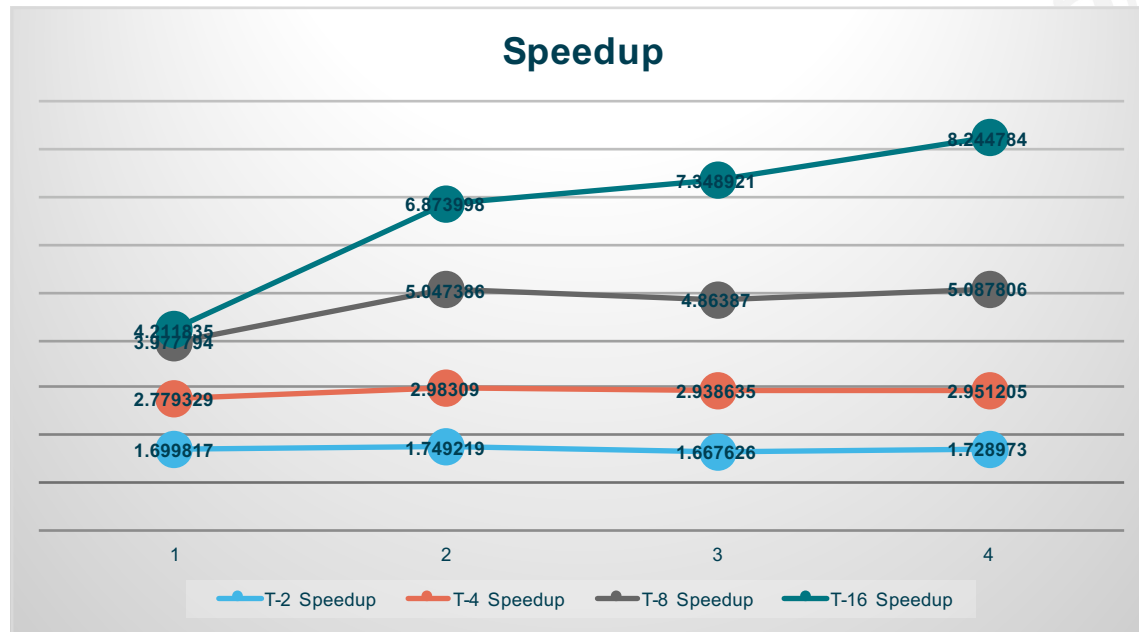
- For 32 cores:

- Difficult to obtain a single server with 32 cores by salloc command in CCR.
- Experiments were run in vortex1 front end of CCR which was equipped with 32 cores.
- It was conducted during off-peak hours (between 3:00 AM – 4:00 AM EST)

Double data size and PEs (threads)



Speedup



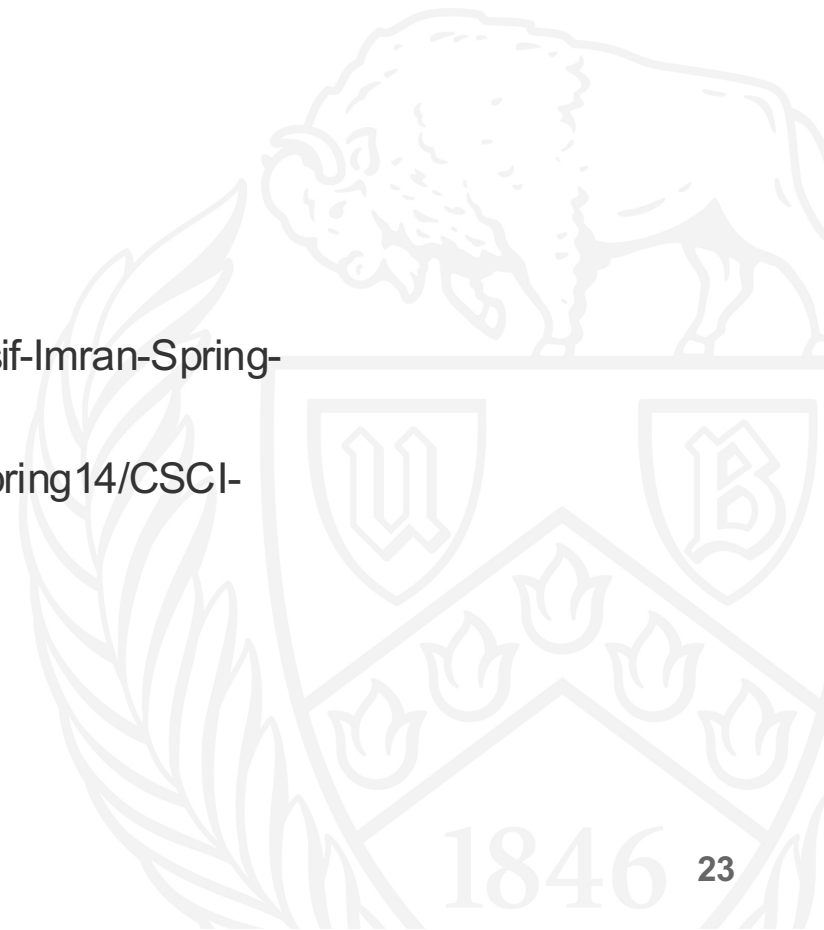
Outcome

- Presented results of the runtimes in a multitude of dimensions
- Discussions
 - OpenMP's runtime is desirable for considerably large datasets
 - Performance degrades when large number of PEs are used for data of size 25000 or lower
 - Desirable behavior when we double data and PEs

Made good use of free resources during thanksgiving

References

- Dr. Russ Miller's webpage:
<https://cse.buffalo.edu/faculty/miller/teaching.shtml>
- Parallel Odd Even Transposition sort:
<https://cse.buffalo.edu/faculty/miller/Courses/CSE633/Asif-Imran-Spring-2018.pdf>
- Parallel Computing Sorting <https://cs.nyu.edu/courses/spring14/CSCI-UA.0480-003/lecture11.pdf>



Thank you

