

# PARALLEL BREADTH-FIRST SEARCH USING MPI

CSE 708 Programming Massively Parallel Systems

Presenter: Sandeep Kunusoth (50465621)

Instructor: Dr. Russ Miller




# Contents

- Breadth First Search
- Applications of BFS
- Serial Implementation of BFS - Dry Run
- Issues with serial implementation and Need for Parallelization
- Parallel Implementation of BFS - Dry Run
- Advantages of Parallel over Serial Implementation
- Results
- References

# Breadth-First Search

- It is a graph traversal algorithm.
- Starts with a given start node and traverse the graph layer wise. We then move towards the next level neighbors.
- Drawback: Extra memory required. Generally Queue, to keep track of unexplored nodes.



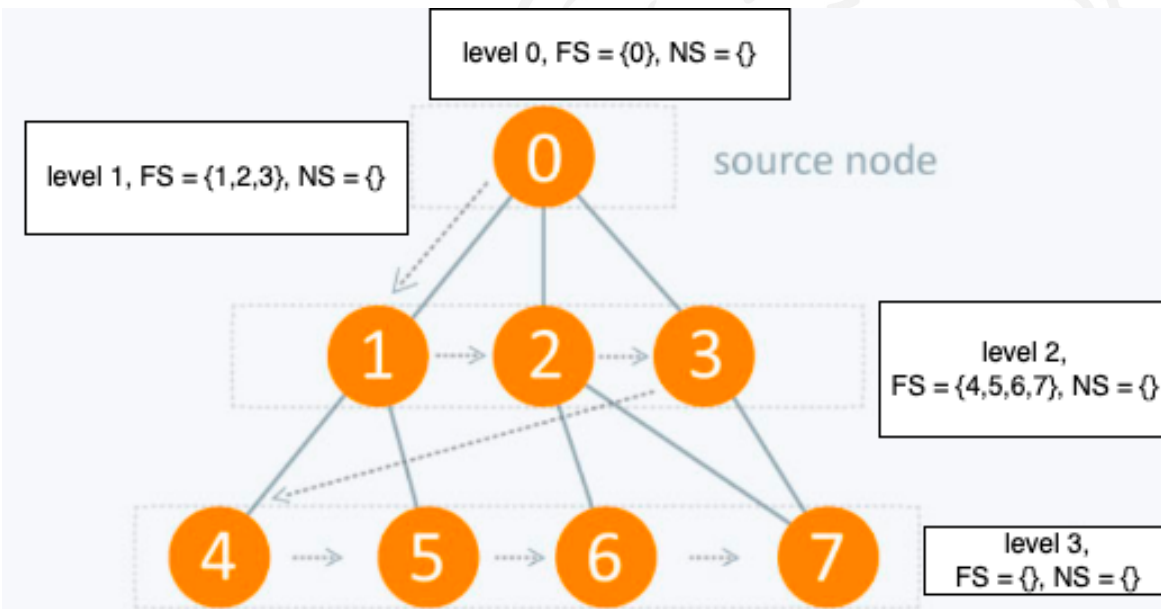
source: [hackerearth](https://leetcode.com/problems/breadth-first-search/)

# Applications of BFS:

- BFS can be used to find shortest path between 2 geographical locations on map as routing algorithms for navigation systems.
- BFS is used by search engines to index and crawl the web pages.
- Peer to Peer Networks like BitTorrent.
- BFS can be used in AI applications such as path finding, recommender systems.
- BFS can be used in game theory to find next best move in games like Chess etc.

# Serial BFS implementation

```
1  define serial_bfs(graph (V,E), source s):  
2      for all v in V do  
3          distance[v] = -1;  
4      distance[s] = 0; level = 0; FS = {s}; NS = {};  
5      while FS is not empty do  
6          level = level + 1;  
7          for u in FS do  
8              for each neighbour v of u do  
9                  if distance[v] = -1 then  
10                     push(v, NS);  
11                     distance[v] = level;  
12             FS = NS, NS = {};
```



order of traversal:

0 -> 1 -> 2 -> 3 -> 4 -> 5 -> 6 -> 7

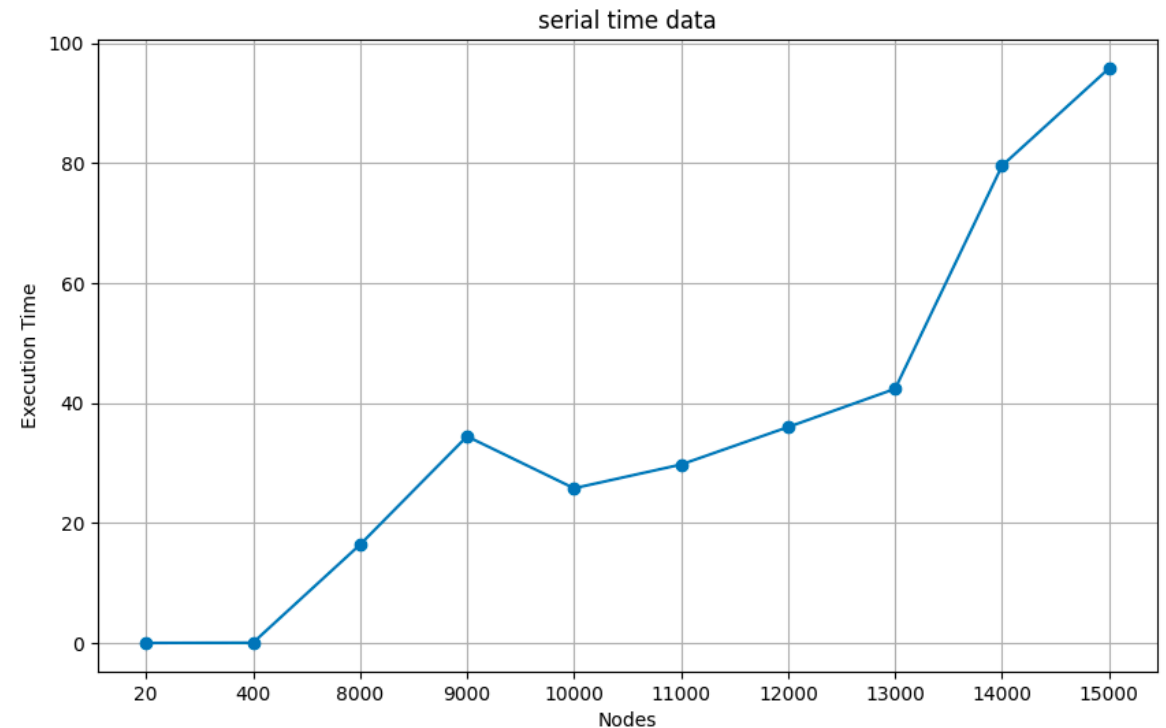
[https://en.wikipedia.org/wiki/Parallel\\_breadth-first\\_search](https://en.wikipedia.org/wiki/Parallel_breadth-first_search)

# Issues with serial implementation and Need for Parallelization

```

1  define serial_bfs(graph (V,E), source s):
2      for all v in V do
3          distance[v] = -1;
4      distance[s] = 0; level = 0; FS = {s}; NS = {}
5      while FS is not empty do
6          level = level + 1;
7          for u in FS do
8              for each neighbour v of u do
9                  if distance[v] = -1 then
10                     push(v, NS);
11                     distance[v] = level;
12             FS = NS, NS = {};
  
```

[wiki/Parallel\\_breadth-first\\_search](https://en.cppreference.com/w/cpp/algorithm/parallel_breadth-first-search)






# Parallel BFS implementation

```
1 define 1_D_distributed_BFS(graph(V,E), source s, rank):
2   for all v in V do
3     distance[v] = -1;
4   level = 0; FS = {}; NS = {};
5   if find_owner(s) = rank then
6     FS = {s}; distance[s] = 0;
7   global_FS_is_not_empty = true
8   while global_FS_is_not_empty do
9     level = level + 1;
10    FS = {set of local vertices}
11    for each u in FS do
12      for each neighbor v of u do
13        j = find_owner(v)
14        push(v, send_buffer[j])
15    // all-to-all communication
16    for 0 <= j < p do
17      if j != rank then
18        send send_buffer[j] to j
19        recv recv_buffer[j] from j
20    NS = {neighbors of vertices in FS including non local}
21    for each u in NS and distance[u] == -1 do
22      distance[u] = level
23      push(u, FS)
24    NS = {};
25    global_FS_is_not_empty = AllReduce(FS.size(), SUM) == 0
```

## Modifications:

- Similar to serial BFS implementation, but instead of checking the queue of vertices sequentially, we implement this in parallel across all the vertices at the same level.
- A neighbor vertex from one processor may belong to other processor. Hence each processor needs to communicate with all others.
- The algorithm ends when global size of frontier across all processors is zero.

# Dry Run



8 Vertices divided between 4 processors

- Processor 0: {0, 1}
- Processor 1: {2, 3}
- Processor 2: {4, 5}
- Processor 3: {6, 7}

Processor 0:

Iteration 1:

FS = {0}

NS = {1}

All visited = {1,2,3}

Iteration 2:

FS = {1}

NS = {}

All visited={4}

Processor 1:

Iteration 1:

FS = {}

NS = {}

Iteration 2:

FS = {2,3}

NS = {}

All visited={5,6,7}

Processor 2:

Iteration 1:

FS = {}

NS = {}

Iteration 2:

FS = {}

NS = {}

Iteration 3:

FS = {4,5}

NS = {}

Processor 3:

Iteration 1:

FS = {}

NS = {}

Iteration 2:

FS = {}

NS = {}

Iteration 3:

FS = {6,7}

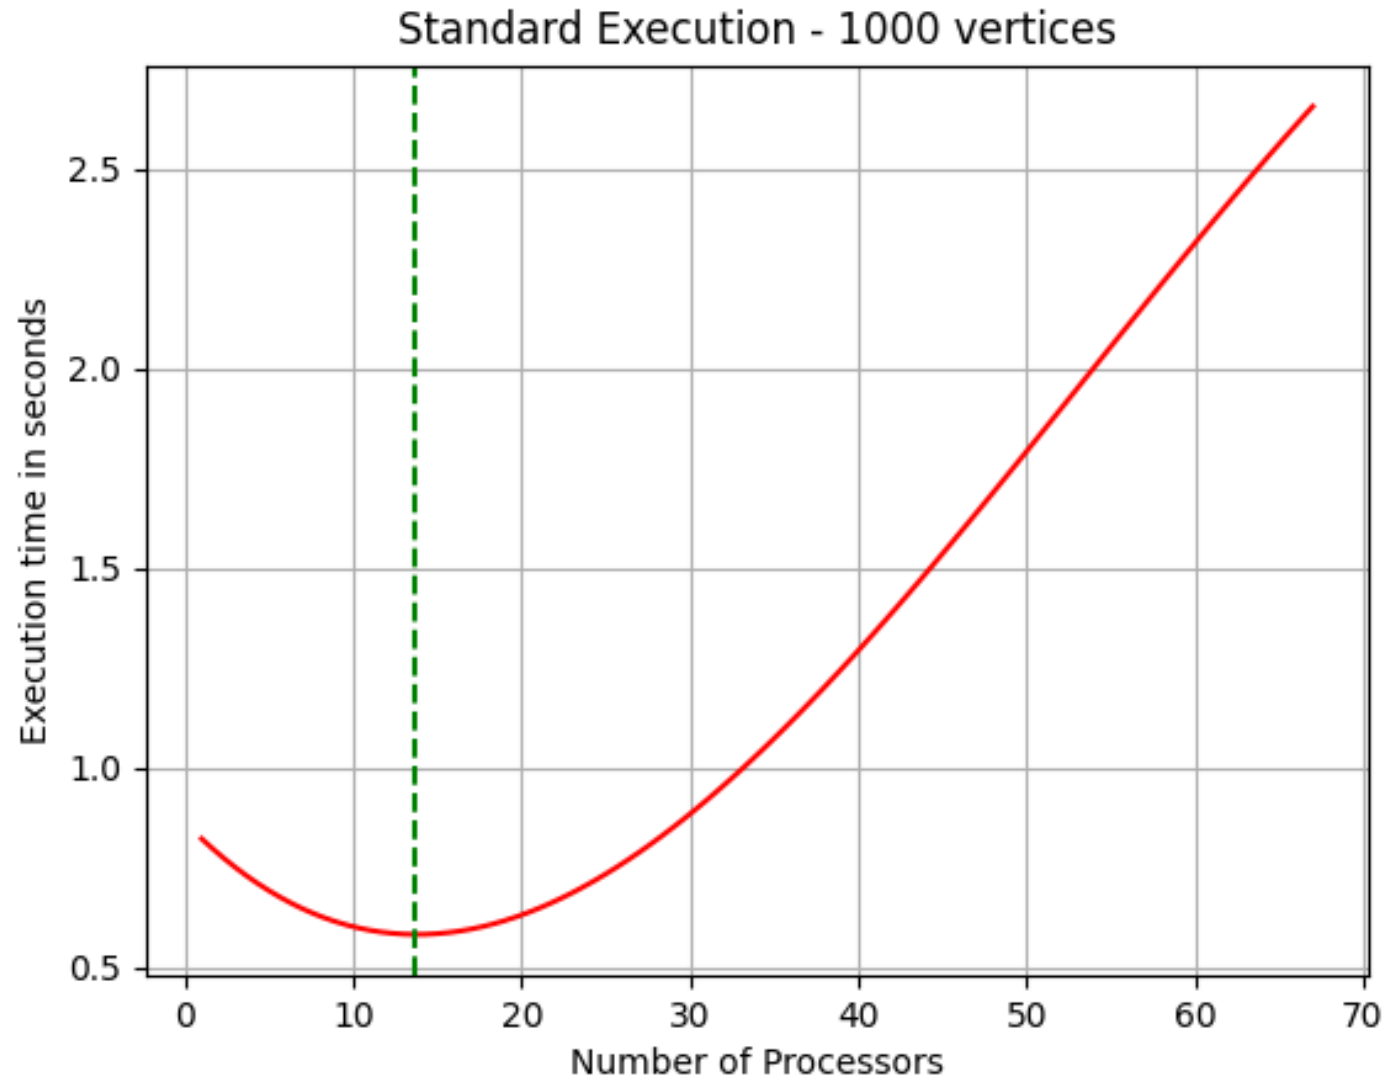
NS = {}




# Advantages of Parallel over Serial Implementation

- **Efficiency:** Parallel BFS improves performance by processing multiple vertices in parallel, significantly enhancing overall efficiency.
- **Scalability:** It is highly scalable and can handle large scale graphs.
- **Concurrency:** Parallel BFS allows for concurrent exploration, minimizing idle time and maximizing resource utilization.
- **Load balancing:** This ensures efficient utilization of computational resources.


```
$ slurm.sh
1  #!/bin/bash
2
3  #SBATCH --nodes=40
4  #SBATCH --ntasks-per-node=1
5  #SBATCH --constraint=IB|OPA
6  #SBATCH --time=00:10:00
7  #SBATCH --partition=general-compute
8  #SBATCH --qos=general-compute
9  #SBATCH --job-name="bfs-4000-vertices-40-nodecore"
10 #SBATCH --output=output-4000-vertices-40-nodecore.txt
11 #SBATCH --error=output-4000-vertices-40-error.txt
12 #SBATCH --exclusive
13
14 module load ccrsoft/2023.01
15 module load gcccore/11.2.0
16 module load intel
17 module load python/3.9.6
18
19 export I_MPI_PMI_LIBRARY=/opt/software/slurm/lib64/libpmi.so
20 srun pip install mpi4py numpy > /dev/null 2>&1
21
22 srun -n 40 python parallel-bfs.py 4000 60
23
```




Standard Execution - 2000 vertices



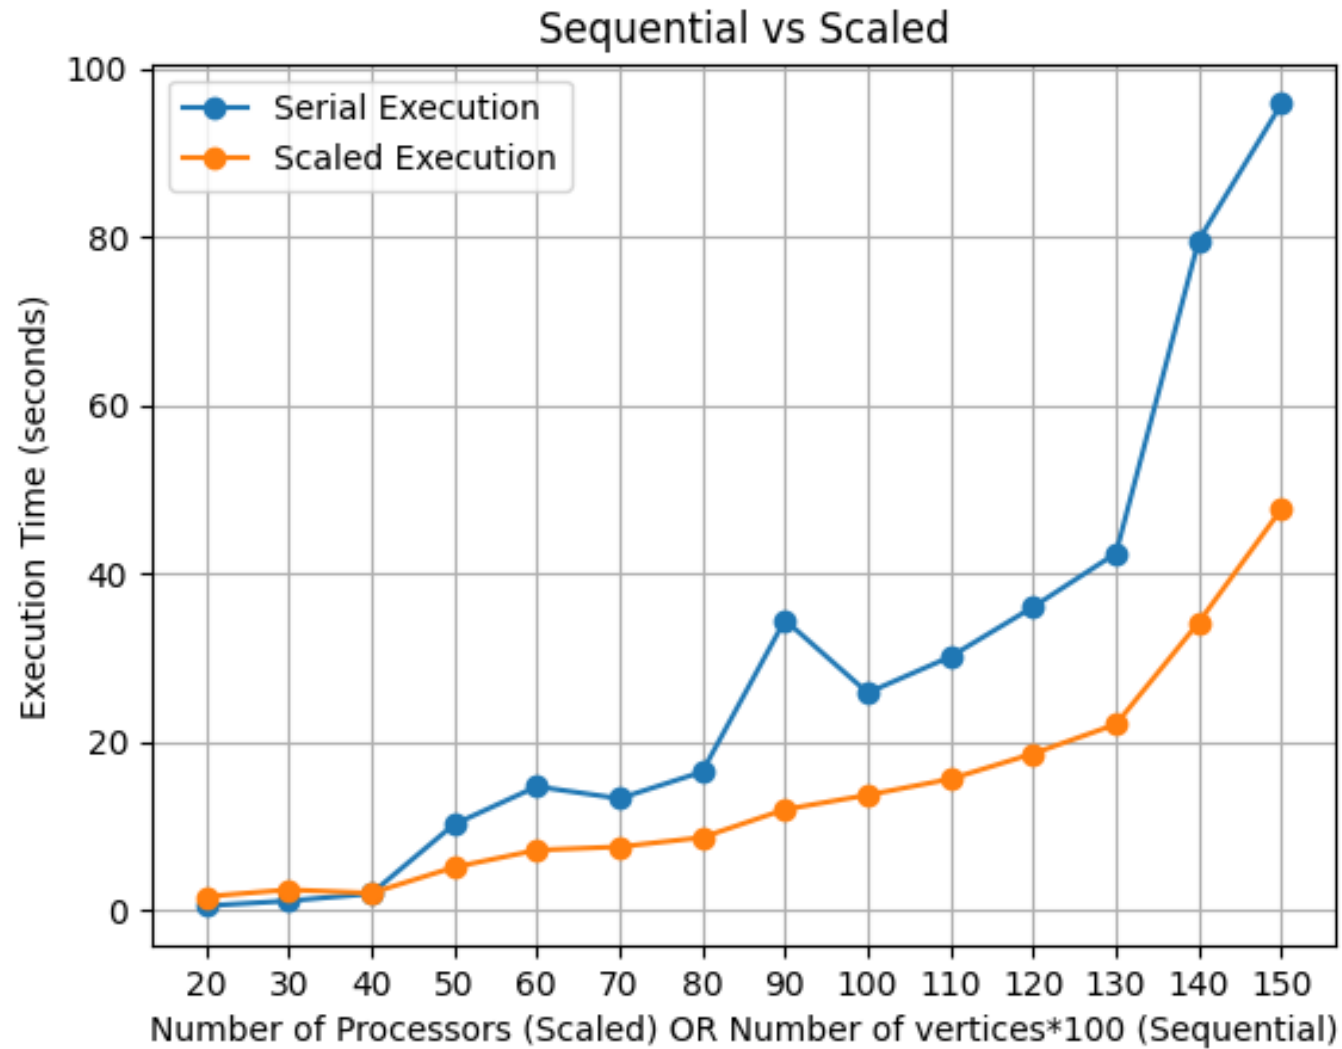
Standard Execution - 4000 vertices



Standard Execution - 8000 vertices



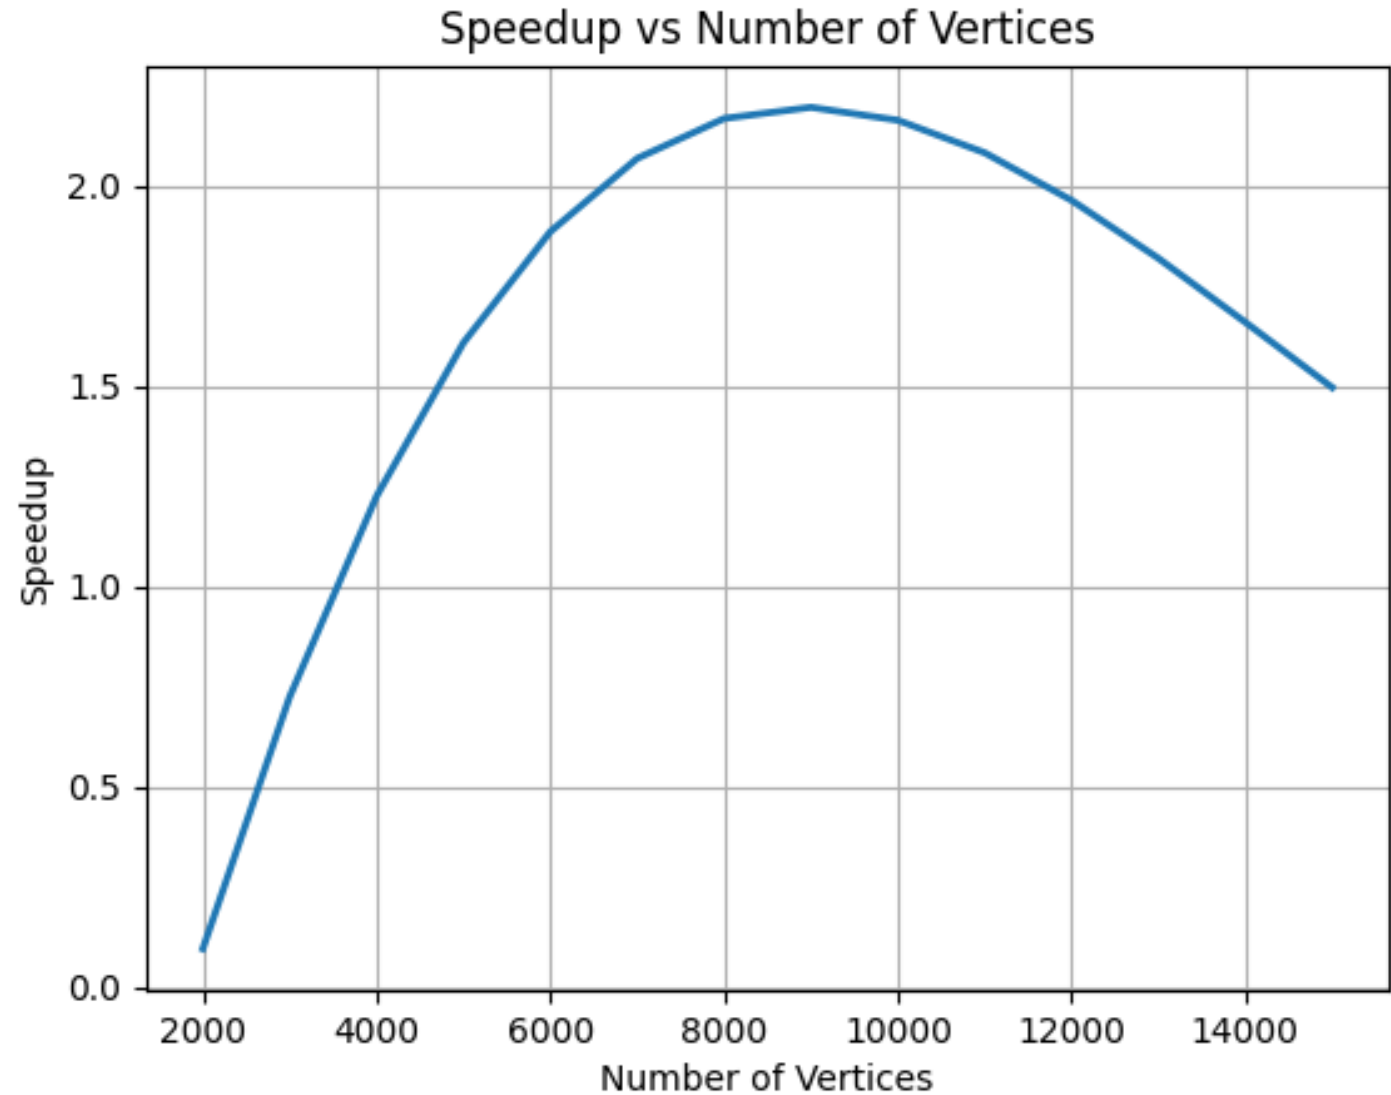




$$\text{Speed up} = T_{\text{seq}} / T_p$$

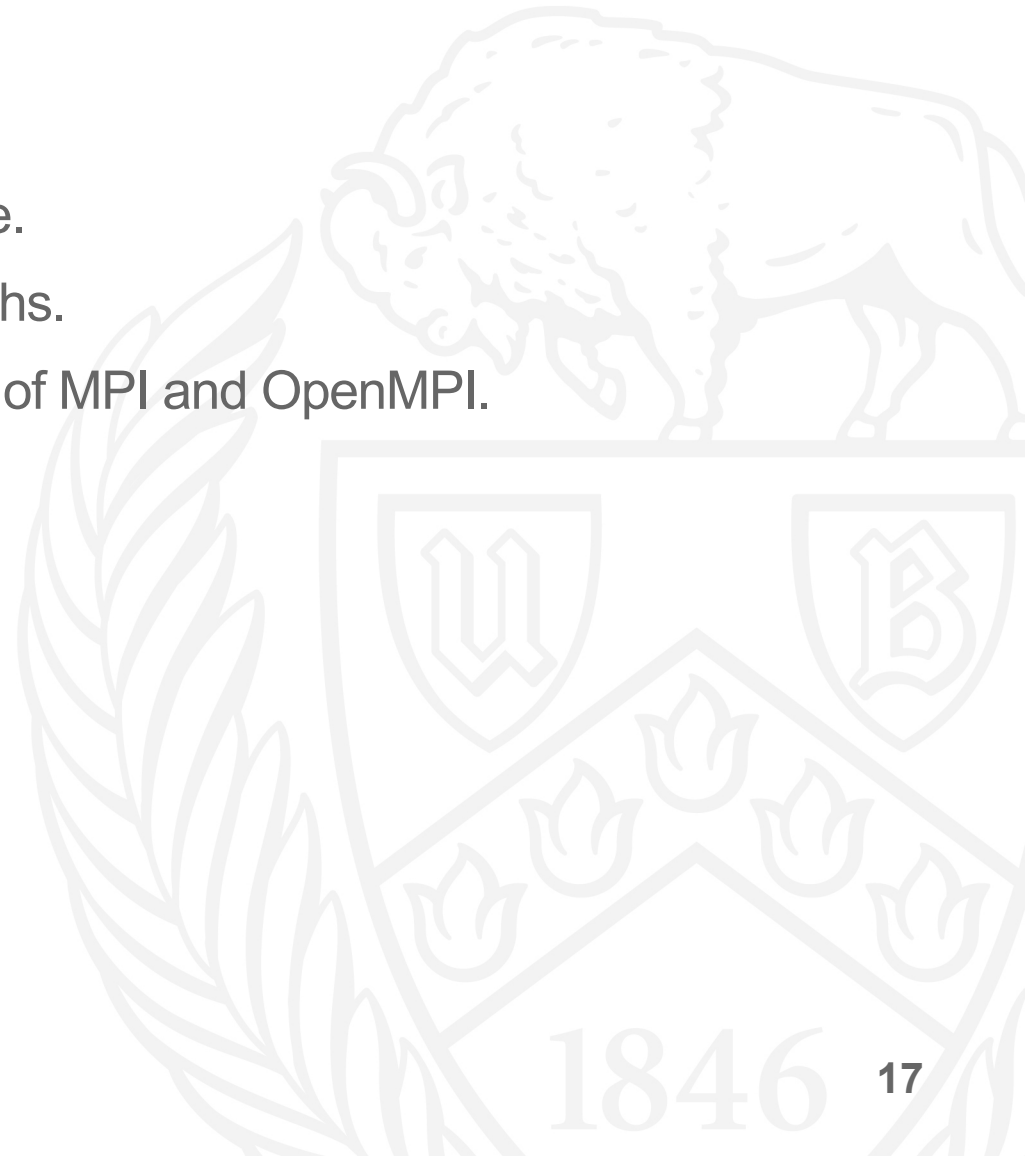
$T_{\text{seq}}$  is the execution time of sequential algorithm.

$T_p$  is the execution time of the parallel algorithm with  
 $p$  Processors



# Future Work

- Access nodes greater than 143 nodes with 1 core per node.
- Test performance by changing density of edges in the graphs.
- Implement my parallel approach using OpenMPI or Hybrid of MPI and OpenMPI.



# References

- Wikipedia [https://en.wikipedia.org/wiki/Parallel\\_breadth-first\\_search](https://en.wikipedia.org/wiki/Parallel_breadth-first_search)
- BFS <https://www.hackerearth.com/practice/algorithms/graphs/breadth-first-search/tutorial/>
- Parallel BFS on Distributed Memory Systems [https://people.eecs.berkeley.edu/~aydin/sc11\\_bfs.pdf](https://people.eecs.berkeley.edu/~aydin/sc11_bfs.pdf)
- Distributed BFS Algorithm, IIT Delhi <https://www.youtube.com/watch?v=wpWvCabHqQU>
- Applications <https://www.ijcsma.com/articles/graph-traversals-and-its-applications-in-graph-theory.pdf>
- CCR Docs <https://docs.ccr.buffalo.edu/en/latest/>
- MPI for Python <https://mpi4py.readthedocs.io/en/stable/>
- MPI python [https://www.youtube.com/watch?v=36nCgG40DJo&ab\\_channel=SharcnetHPC](https://www.youtube.com/watch?v=36nCgG40DJo&ab_channel=SharcnetHPC)