Improving Performance of BFS Algorithm with GPU Computing using CUDA

Freya Amaranta Sutedja
Jakarta Intercultural School
Jakarta, Indonesia
freyamaranta@gmail.com

Olivia Tansi
Jakarta Nanyang School
Kabupaten Tangerang, Indonesia
oliviatansi79@gmail.com

Claudia Angelica
St. Joseph's Institution International
Singapore, Singapore
angelsmail157@gmail.com

Michael Nathan Pamurahardjo
SpringField School
Jakarta, Indonesia
michaelnathanpam@gmail.com

***Abstract*—This paper aimed to implement both sequential and parallel implementations using CUDA on Breadth First Search (BFS) to see the differences and effects of it, followed by an analysis of the result. We used the algorithm as that will be elaborated more on the paper, here we would like to generally compare its memory consumption and run time. It is found out that parallel implementation runs faster on an average of 31.23 times compared to sequential implementation running the same task. This is possible due to complete mapping on each column operations into thread memory on GPU, as it gives constant time complexity to each queue operation on its Traversal. Optimum threads to be used in parallel programming is also needed to be found, here we try to find it with trial-and-error testing. Further research would involve more complex parallel programming implementation and a more controlled testing environment.**

Keywords—Breadth First Search, Sequential, Parallel Computing, CUDA, NVIDIA

I. INTRODUCTION

Parallel processing is the division of process into different parts, which performed concurrently by different processors in a computer [1]. Parallel programming is a concept of using two or more processors to complete a task. Parallel programming comes with various benefits from solving larger problems, doing things faster in a more reasonable time, and more cases can be finished [2]. Set side by side with CPU, GPU is taking into consideration to work with parallel computations as it suits more. GPU nowadays have considerably developed from recent years not just for gaming as its early purpose, with the extend of General-purpose computations on GPUs (GPGPU), as one of them is NVIDIA CUDA [3].

In memory hierarchy, on parallel program execution the threads will access data from some memory spaces. Each thread will have both access to local and global memory. This will make memory consumption on parallel computing consumes more rather than serial, because there are a lot of bandwidth usage that happens upon memory transfers [4].

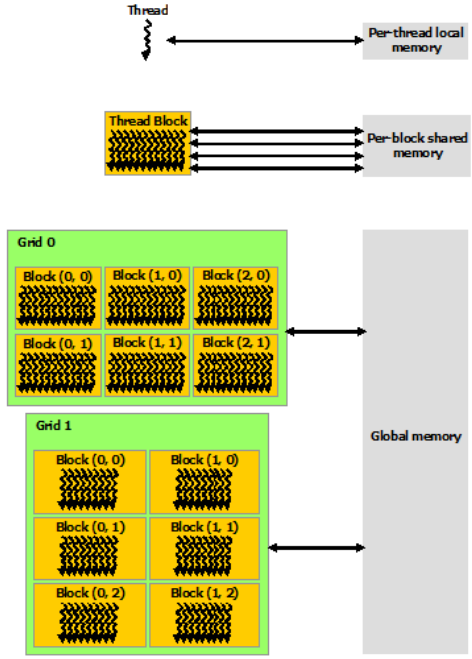

Fig 1.1. Memory Hierarchy [4]

Compute Unified Device Architecture or better known as CUDA is a platform for developer to perform parallel computing developed by NVIDIA on graphical processing units (GPUs). CUDA enables computing to be speed up in a dramatic manner leveraging the power of GPUs a machine has. Developing in CUDA can be done in popular languages such as C, C++, Python, etc [5]. Analyzing tool is also provided using the CUDA Profiling Tools to investigate the details of our program [6].

Here what we are going to analyze is the usage of parallel programming using CUDA compared to the sequential programming in performing Breadth-First Traversal. The BFS algorithm will be done onn adjacency square matrix (N×N matrices), which then the run time and memory consumption will be recorded and compared. Both sequential and parallel solutions will be developed in C programming language. The machine used for running all the solutions is Lenovo Ideapad Gaming Laptop 3, with the details as such: AMD Ryzen 5 4600H Processor @ 3.0GHz (12CPUs), 16384MB RAM, and NVIDIA® GeForce® GTX 1650Ti Graphics (4 GB GDDR5 dedicated).

II. BREADTH-FIRST TRAVERSAL

We have done several methods on calculating and solving the Breadth First Traversal to obtain algorithm with think suits the best and fastest running time compared to the other. The goal is to compute the Breadth-First Traversal, with N×N adjacency matrix, and N is scaled as big as possible.

A. Sequential Algorithm

Breadth first search (BFS) is a general algorithm that is often used to search for the shortest path in an unweighted graph using a vertex-based technique. It is optimal for finding the shortest path. This algorithm is present in real life settings such as transportation, communication, information, and social networks.

Given an undirected or directed bipartite graph G = (V,E), where V stands for vertex or node and E stands for edge, and a distinguished starting node s, BFS finds the smallest number of edges to every node reachable from s. It visits every node and edge of G to determine whether the graph is connected. Nodes on the same level (siblings) are visited first before moving onto the next (children) and nodes that have been visited more than once will be deleted from the queue.

BFS begins at s and inspects every adjacent node the starting node s in a left to right or top to bottom order. It will do the same for every unvisited adjacent node and repeat the same action until all the nodes are visited. Its time complexity is O(V+E) when using an adjacency list and P(V2) when using an adjacency matrix. The number of columns and rows in the matrix depends on the number of nodes. As the size of the matrices increase. The number of nodes and edges will increase, and so does the time it takes for the BFS to traverse through all the nodes following the time complexity formula O(V+E).

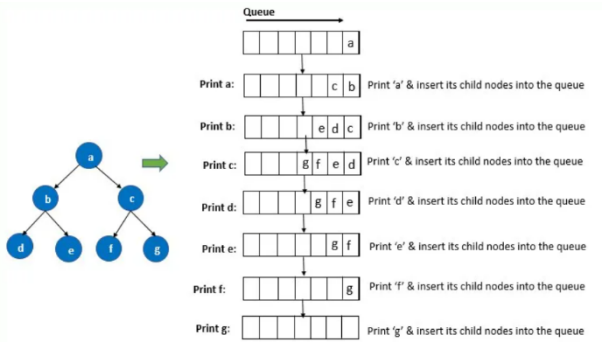

Fig 2.1. Algorithm Illustration [7]

As we want to have the matrices as large as possible, it is a problem when storing the numbers in an array. As an array have certain limitation, it is not possible to use fixed multi-dimensional array. So, in order to tackle this problem, we proposed to use a dynamic multi-dimensional array as shown below. The following function have been tested and able to store up to N=8192.

```
int** create_matrix(int rows, int cols) {
     int** mat = new int* [rows];
     for (int i = 0; i < rows; ++i) {
          mat[i] = new int[cols]();
     }

     return mat;
}
```

After matrix is created, it will then be filled up on another function with symmetrical binary values to represent the vertices and edges. The following code shows code that uses iterative method by implementing data structure queue with supportive function enqueue and dequeue, we are holding the value of current vertices stored in the queue array before getting evaluated. We can see the illustration on Figure 2.1. Result verification for large input data is done by tracing up the vertices and edges printed until n = 64 for the number of vertices.

```
        printf("%d ", i);
        visited[i] = 1;
        enqueue(&q, i); // Enqueue i for exploration
        while (!isEmpty(&q))
        {
            int node = dequeue(&q);
            for (int j = 0; j < 7; j++)
            {
                if(a[node][j] ==1 && visited[j] == 0){
                    printf("%d ", j);
                    visited[j] = 1;
                    enqueue(&q, j);
                }
            }
        }
```

B. Parallel Algorithm

Identical to the Sequential Algorithm purpose and goal, which is to perform Breadth First Traversal with matrix size as large as possible. The code is modified from Harish & Narayan [8]. The dataset for parallel algorithm is the same unweighted and undirected graph generated for sequential implementation. However, the data is converted into a compact adjacency list to match the architecture of GPU Thread Memory as proposed by Harish & Narayan. The following code shows the implemented kernel:

```
__global__ void CUDA_BFS_KERNEL(Node *Va, int *Ea, bool *Fa, bool *Xa, int
*Ca,bool *done) // memory races on both Xa and Ca

{

    int id = threadIdx.x + blockIdx.x * blockDim.x;

    if (id > NUM_NODES) *done = false;

    if (Fa[id] == true && Xa[id] == false) {

    Fa[id] = false; Xa[id] = true;

      __syncthreads();

      printf("%d ", id); //Change CUDA settings for printf support

      int k = 0; int i;

      int start = Va[id].start;

      int end = start + Va[id].length;

      for (int i = start; i < end; i++)

      {

        int nid = Ea[i];

        if (Xa[nid] == false)

        {

          Ca[nid] = Ca[id] + 1;

          Fa[nid] = true;

          *done = false;

        }

      }

    }

}
```

The printf in the above code is needed to print the vertices that appear in the order of BFS. `__syncthreads()` call is needed as well, since once an active node has been processed, the cost associated with its neighbours need to be processed only after the visited of all threads is updated accordingly. Hence `__synchthreads` prevents racing between `Xa` and `Ca`.

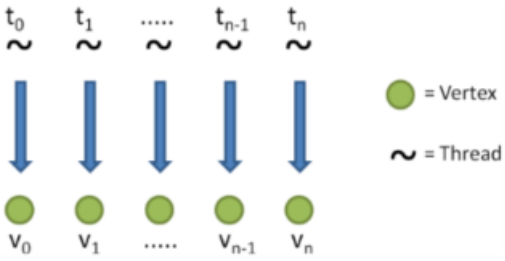

Fig 2.2. Example Mapping Threads into Vertices [3]

In this version of implementation, we mapped 1 thread per vertex, demonstrated in figure above, by using various block size to test the capability of GPU to process each vertex. Thus, this shows that there is potential to have 511 redundant threads that are idle during kernel execution.

III. HYPOTHESIS

A. *Grid and Block*

Definitions of grid and block are interconnected one another. Block is a programming abstraction that represents a group of thread, which can be executed either in a serial or parallel way. Grid have the same concept, but it is formed by group of blocks which can be executed in just one kernel [4].

Besides, the following shows the results for the machine *deviceQuery* which explain lot more information about our GPU and its CUDA processing capability. At first, the authors thought that the maximum capable thread to be used in the parallel algorithm will be equal to 1024 according to *Maximum number of threads per block*. Eventually, on reality the capable thread being used is as on stated in the *Warp Size* as the authors tried to put in the value accordingly to the given code.

```
Device 0: "NVIDIA GeForce GTX 1650 Ti"
  CUDA Driver Version / Runtime Version          11.4 / 11.1
  CUDA Capability Major/Minor version number:    7.5
  Total amount of global memory:                 4096 MBytes (4294967296 bytes)
  (16) Multiprocessors, ( 64) CUDA Cores/MP:     1024 CUDA Cores
  GPU Max Clock rate:                            1485 MHz (1.49 GHz)
  Memory Clock rate:                             6001 Mhz
  Memory Bus Width:                              128-bit
  L2 Cache Size:                                 1048576 bytes
  Maximum Texture Dimension Size (x,y,z)         1D=(131072), 2D=(131072, 65536), 3D=(16384, 16384, 16384)
  Maximum Layered 1D Texture Size, (num) layers  1D=(32768), 2048 layers
  Maximum Layered 2D Texture Size, (num) layers  2D=(32768, 32768), 2048 layers
  Total amount of constant memory:               65536 bytes
  Total amount of shared memory per block:       49152 bytes
  Total shared memory per multiprocessor:        65536 bytes
  Total number of registers available per block: 65536
  Warp size:                                     32
  Maximum number of threads per multiprocessor:  1024
  Maximum number of threads per block:           1024
  Max dimension size of a thread block (x,y,z): (1024, 1024, 64)
  Max dimension size of a grid size    (x,y,z): (2147483647, 65535, 65535)
  Maximum memory pitch:                          2147483647 bytes
  Texture alignment:                             512 bytes
  Concurrent copy and kernel execution:          Yes with 2 copy engine(s)
  Run time limit on kernels:                     Yes
  Integrated GPU sharing Host Memory:            No
  Support host page-locked memory mapping:       Yes
  Alignment requirement for Surfaces:            Yes
  Device has ECC support:                        Disabled
  CUDA Device Driver Mode (TCC or WDDM):         WDDM (Windows Display Driver Model)
  Device supports Unified Addressing (UVA):      Yes
  Device supports Managed Memory:                Yes
  Device supports Compute Preemption:            Yes
  Supports Cooperative Kernel Launch:            Yes
  Supports MultiDevice Co-op Kernel Launch:      No
  Device PCI Domain ID / Bus ID / location ID:   0 / 1 / 0
  Compute Mode:
     < Default (multiple host threads can use ::cudaSetDevice() with device simultaneously) >

deviceQuery, CUDA Driver = CUDART, CUDA Driver Version = 11.4, CUDA Runtime Version = 11.1, NumDevs = 1
Result = PASS
```

Fig 3.1. Device Query

B. *Sample Program and Memory Allocation on GPU*

The sample program we tried on is addition on grid and block [9], the first thing that we create is a function that has this ability.

```
__global__ void arradd(int* md, int* nd, int* pd) {
        int myid = blockIdx.x * blockDim.x + threadIdx.x;
        pd[myid] = md[myid] + nd[myid];
        printf("Block Number: %d Thread number : %d.\n",
blockIdx.x, threadIdx.x);
}
```

Afterwards, we can start to declare variables that we needed to work with in the main function. Continue the process with *cudaMalloc* that functions to allocate memory to GPU. Together with using *cudaMemcpy*, we could copy the data in array from CPU to the memory in GPU.

```
int size = MAXN * sizeof(int);
int m[MAXN], n[MAXN], p[MAXN], * md, * nd, * pd;
int i = 0;

for (i = 0; i < MAXN; i++) {
        m[i] = ;
}

cudaMalloc(&md, size);
cudaMemcpy(md, m, size, cudaMemcpyHostToDevice);
```

Next thing to work with is declaring grid and block dimension, whereas in the code uses *dim3* which is an integer vector and uses *Block* and *Thread*. After creating grid and block dimension we use *arradd* function, a kernel that has we created before in CUDA. This function will add *m* and *n* then put the results in array *pd* according to the size of grid and block dimension declared. After finishing the process, we can free the memory as we have done in parallel algorithm.

```
dim3 Block(4);
dim3 Thread(5);
arradd << < Block, Thread >> > (md, nd, pd);
cudaMemcpy(p, pd, size, cudaMemcpyDeviceToHost);

cudaFree(md);
cudaFree(nd);
cudaFree(pd);

for (int i = 0; i < 20; i++) {
        printf("\t%d", p[i]);
}
```

Here on the sample program, we learned that *Block* and *Thread* affect the data being performed which is only equals to *Block x Thread*, we think that if the resource is not allocated then the process cannot be done. Besides, the running *Block* is not may always be in order. For example, *Block 3* can run first than *Block 0*. While for the BFS algorithm, the following code is used:

C. *Execution Time and Memory Consumption Hypothesis*

Several other case that we have tried to implement sequential Breadth First Traversal on C, we think that it might take a long time to execute even on small matrices. We predict that the memory consumed is more or less the product of N×N which one of the number will represent 4 bytes [10], with the total memory consumption is N × N × 4 bytes. We see that from the sequential algorithm the largest cost is on

the computation to fill up the resulting matrix, which is a quadratic time of $O(n^2)$.

On parallel implementation, we think that it will always be faster than the sequential algorithm as it leverages GPU not just the CPU. Thread will also come to play a role in the execution time of parallel programming, as every single thread is mapped into a vertex and the sequential best case running time of adjacency list implementation is O(V+E) then the larger the thread, as long as it can map every single vertex on the graph, it is expected to achieve constant running time - O(1).Other than that, for parallel implementation memory, we guess that it will just be more or less than two times the memory needed for sequential one. The guess came from that one is used in CPU while the other is consumed on GPU.

IV. Result and Analysis

After the project was done completely, finally we can evaluate the results of each algorithm on its run time and memory consumption. The testing method is described in each sub chapter correspondingly, as it uses different methodology to run on sequential and parallel implementation.

Please be informed that all the calculations and results provided may not be fully accurate as there may be technical errors and many other things to consider going on in the machine (e.g. running other programs together, plugged in to electricity, etc.)

A. Sequential Implementation

To obtain the running time of sequential implementation and to get the best and most accurate result, for each N×N matrices it is being executed for 30 times. To do it automatically, we put all the algorithm of Breadth-First Traversal described before on chapter two to a function that accepts two parameters, which is its row and column. The function also contains the code below to record time.

Doing the execution on main and other function does not give a significant run time difference on the Traversal. Thus, after running the function 30 times using a for loop. The result of average run time is shown below in a table. N.B. n/a means that it is not available as the answer cannot be provided (equal to 0).

TABLE I.
Sequential Run Time

N	***Time in second***
2, 4, 8, 16	n/a
32	1×10^{-4}
64	9×10^{-4}
128	6.999×10^{-3}
256	5.723×10^{-2}
512	5.400×10^{-1}
1024	5.814×10^{0}
2048	7.108×10^{1}
4096	6.27×10^{2}
8192	4.520×10^{3}

On the other hand, the recording of memory consumption used the default profiler given on Visual Studio 2019. The number displayed on the table is the peak of the memory consumption. While on the process it is have several steps especially when filling up matrix a, b, and c which have significant increases on memory consumption. The result is recorded as below.

TABLE II.
Sequential Memory Consumption

N	***Memory in MB***
2, 4, 8, 16, 32, 64	n/a
128	1.4×10^{0}
256	2.2×10^{0}
512	4.8×10^{0}
1024	1.49×10^{1}
2048	5.59×10^{1}
4096	1.94×10^{2}
8192	7.725×10^{2}

On the table it is shown that from N=2 until N=64, the memory being consumed is less than 1MB. We conclude that way because it is that even when a 2x2 matrices, it only needs roughly 16bytes (assume each int needs 4bytes). The result cannot be provided as it requires a program to run for one second or more to see diagnostic details on Visual Studio 2019.

B. Parallel Implementation

Corresponding to the sequential implementation, here we will also discuss the result of run time of parallel algorithm implementation on Breadth First Traversal. Unlike sequential programming, we now need to determine the number of threads to use for parallel programming. The following are the tables of result of traversing N×N matrices using different number of threads.

TABLE III.
Threads Effect on Run Time (in second)

Thread	***1024***	***2048***	***4096***	***8192***
2	2.48×10^{0}	9.59×10^{0}	6.37×10^{1}	4.92×10^{2}
4	1.78×10^{0}	4.02×10^{0}	1.93×10^{1}	1.34×10^{2}
8	1.64×10^{0}	2.64×10^{0}	7.99×10^{0}	4.30×10^{1}
16	1.55×10^{0}	2.45×10^{0}	7.95×10^{0}	4.25×10^{1}
32	1.49×10^{0}	2.56×10^{0}	7.83×10^{0}	4.25×10^{1}

Based on the table, the upfront prediction on hypothesis holds as threads are affecting run time of a parallel implementation. It is important to use optimal threads as it is significant and can reduce time up to an average of 73.36% (comparing the slowest and the fastest). We also found out that the optimal thread on the machine and algorithm is 16 threads on some cases. There is also a stagnant level around threads 8, 16, and 32 compared to using 2 or 4 threads with significant differences. There are two columns to describe the run time in parallel implementation. The “Time in second” column is the product from Visual Studio 2019 profiler. While the “NVProf” column comes from using CUDA Event API. Where on the first few line we will create initializer for the counter and start to record the time as well, as the code below [11].

```
cudaEvent_t start, stop;
float milliseconds=0;
cudaEventCreate(&start);
cudaEventCreate(&stop);
cudaEventRecord(start);
// Parallel Algorithm Here
cudaEventRecord(stop);
cudaEventSynchronize(stop);
cudaEventElapsedTime(&milliseconds, start, stop);
```

The results will be printed on the variable milliseconds, as it is also recorded on milliseconds instead of second. The results will be shown on the following table, where the milliseconds have been converted to seconds. Where all the time taken is using 32 threads on the execution.

TABLE IV.
PARALLEL RUN TIME

N	*Time in second*	*NVProf in second*
2, 4, 8, 16, 32, 64	n/a	n/a
128	1.185×10^{0}	5×10^{-4}
256	1.423×10^{0}	1.6×10^{-3}
512	1.425×10^{0}	1.1×10^{-2}
1024	1.487×10^{0}	8.25×10^{-2}
2048	2.56×10^{0}	5.506×10^{-1}
4096	7.831×10^{0}	3.950×10^{0}
8192	4.248×10^{1}	3.134×10^{1}
16384	2.943×10^{2}	2.537×10^{2}

It is obvious that the parallel algorithm runs faster for Breadth First Traversal rather than the sequential one. The run time can be easily noticeable especially on bigger N. When N≤256, the CUDA role is not so significant compared to the total time. On larger N, it is shown that CUDA plays a big role as on N=16384 it consumes 85,99% of the total time. Rows with smaller number of vertices are not compared due to the assumptions that the computation on GPU will produce a similar performance with CPU. This is because the data movement from board memory to GPU memory most likely will give overhead problem. On the other hand, larger N needs more computational power, which parallel programming provide through the usage of GPU. But there are some trade-offs from the fast run time of parallel algorithm that is shown below on the memory consumption.

TABLE V.
PARALLEL MEMORY CONSUMPTION

N	*Memory in MB*
2, 4, 8, 16, 32, 64	n/a
128	7.338×10^{1}
256	2.776×10^{1}
512	3.902×10^{1}
1024	9.133×10^{1}
2048	2.095×10^{2}
4096	4.975×10^{2}
8192	1.6×10^{3}
16384	6.1×10^{3}

On the table it is shown that from N=2 until N=64, the memory being consumed is not available using the default profiler on Visual Studio 2019. As can be seen on N=128, it consumes memory more than N=256 and N=512, of course it is a weird phenomenon. There has been some spike data on N=128 until N=512, which later on will be more stable on N=1024 and so. The memory needed for such errors sometimes give less than 10MB on an execution and sometimes even more than 50MB on the other tries. It is remained unknown what cause this to happen, we speculate that it might be that there may be errors when placing or filling the data.

C. Comparison and Analysis

This sub chapter will compare the results of both sequential and parallel algorithm implementation by putting the results side by side. First, it is obvious that the parallel algorithm cost less time in the execution compared to the sequential algorithm. We will perform calculation of the total speedup using the following equation.

$$total\ speedup = \frac{T_s}{T_p} \quad (1)$$

The equation above is Amdahl's Law [12]. Let T_s be the computation time needed in sequential (without parallel computation) and T_p be the computation time needed in parallel computation. The comparison and speedup calculation are given in the table below.

TABLE VI.
RUN TIME COMPARISON AND SPEED UP

N	*Sequential Run in second*	*Parallel Run in second*	*Total Speedup in multiple*
2, 4, 8, 16	n/a	n/a	n/a
32	1×10^{-4}	n/a	n/a
64	9×10^{-4}	n/a	n/a
128	6.999×10^{-3}	1.185×10^{0}	0.006
256	5.723×10^{-2}	1.423×10^{0}	0.040
512	5.400×10^{-1}	1.425×10^{0}	0.379
1024	5.814×10^{0}	1.487×10^{0}	3.911
2048	7.108×10^{1}	2.56×10^{0}	27.764
4096	6.27×10^{2}	7.831×10^{0}	80.123
8192	4.520×10^{3}	4.248×10^{1}	106.403

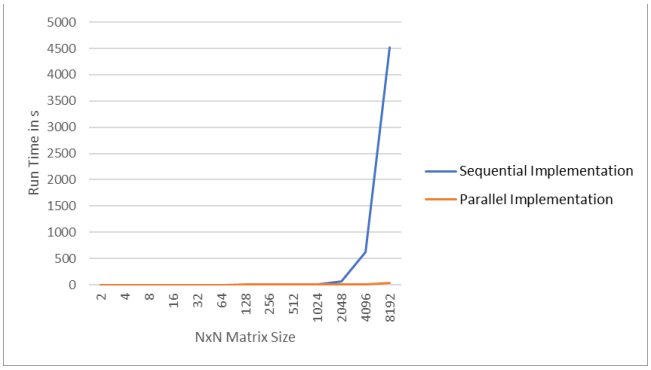

Figure 4.1. Run Time Comparison

From the table and figure presented above, we can see that by making a program parallel is very significant to its run time. Even though, on N≤512 it takes longer time to finish

computation with parallel solution as the $total\ speedup < 1$. We indicate that it may be because of overhead in parallelism in the task (thread or process) start up and termination cost and maybe combined with other problems [13][14]. So, our hypothesis on the previous chapter is incorrect as we predict that parallel will functions well and faster no matter the matrices size is. But in average, the speedup done by making a program parallel is 31.232 times faster. Thus, using parallel implementation in solving problems is still recommended.

Besides comparing the run time, the memory consumption is also clearly stated head-to-head in Table VII. In addition to the memory consumption needed, a new column called "Ratio" was added to show how many times more do the parallel memory consumption needs compared to sequential memory consumption.

$$Ratio = \frac{MC_p}{MC_s} \quad (2)$$

The equation above is made to easily identify the ratio of memory consumed on each process. Where MC_p denotes the memory consumption on a parallel computation. On the other hand, MC_s denotes the memory consumption for a sequential computation.

TABLE VII.
MEMORY CONSUMPTION AND MULTIPLE

N	***Sequential Run in MB***	***Parallel Run in MB***	***Ratio***
2, 4, 8, 16, 32, 64	n/a	n/a	n/a
128	1.4×10^0	7.338×10^1	52.411
256	2.2×10^0	2.776×10^1	12.618
512	4.8×10^0	3.902×10^1	8.129
1024	1.49×10^1	9.133×10^1	6.130
2048	5.59×10^1	2.095×10^2	3.478
4096	1.94×10^2	4.975×10^2	2.562
8192	7.725×10^2	1.6×10^3	2.071

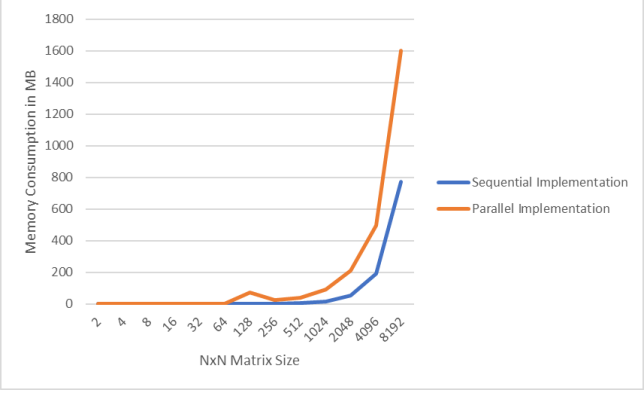

Fig 4.2. Memory Consumption Comparison

On the table and figure above, it is assumed for the memory consumed on N≤64 is 0 as it is not available to do the calculation. It is also shown on both on the table and figure that the difference is quite much whether it is on the run time or on the memory consumption. But vice versa to the run time, memory consumption for parallel computation cost more than sequential. The amount of memory required for parallel computation can be greater than serial, due to the need of replicating data and overheads associated with libraries used and subsystems [14].

V. CONCLUSION

We present Breadth First Traversal on both sequential and parallel model to see the differences. To conclude, it is recommended to use parallel algorithm with optimal thread numbers on bigger cases (breaking point should be found) as it solves things faster with speedup on average of 31.23, instead of the sequential. Some of the important notes to be taken for parallel usage is that it needs more advanced memory handling as the dataset is initially taken from a file, created by CPU. Although running time in parallel computation time would not be a problem, but memory of the machine (CPU) would be the more complex part of solving cases. On further research it should also be explored more about giving constant value on *cudaMemcpy* and parallel implementations for machine learning or deep learning as well as it consumes lots of time and resources which parallel programming can deal with [15].

There are also some errors that remained unsolved on the project as the spike data and the data provided may not be 100% accurate. Some of the recommendations that we might give is to use profiler other than default from Visual Studio 2019 as it may not be that accurate. The machine quality also needs to be maintained (e.g. no running software when executing code, plugged in electricity, heat sink quality, etc.). Implementing Breadth First Traversal either on sequential and parallel may use a better algorithm rather than the one that provided.

REFERENCES

[1] "parallel-processing noun - Definition, pictures, pronunciation and usage notes | Oxford Advanced Learner's Dictionary at OxfordLearnersDictionaries.com." [Online]. Available: https://www.oxfordlearnersdictionaries.com/definition/english/parallel-processing. [Accessed: 05-Sep-2022].

[2] L. Woodard, "Introduction to Parallel Programming What is Parallel Programming?," 2013.

[3] Chetan Pise, Shailendra W. Shende, "NVIDIA Graphics Card for Parallelization of BFS Graphs Algorithm using CUDA" *Elsevier Proceedings of 3rd International Conference on Recent Trends in Engineering & Technology (ICRTET'2014)*, 2014.

[4] "Programming Guide :: CUDA Toolkit Documentation." [Online]. Available: https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html#scalable-programming-model. [Accessed: 13-Sep-2022].

[5] "CUDA Zone | NVIDIA Developer." [Online]. Available: https://developer.nvidia.com/cuda-zone. [Accessed: 05-Sep-2022].

[6] "Profiler :: CUDA Toolkit Documentation." [Online]. Available: https://docs.nvidia.com/cuda/profiler-users-guide/index.html. [Accessed: 15-Sep-2022].

[7] Zulaikha Lateef, "All You Need To Know About The Breadth First Search Algorithm" Available: https://www.edureka.co/blog/breadth-first-search-algorithm/

[8] "Implementing Breadth First Search in CUDA" [Online]. Available: https://siddharths2710.wordpress.com/2017/05/16/implementing-breadth-first-search-in-cuda/ [Accessed: 08-Okt-2022].

[9] Aulia Hening Darmasti, "Cuda Programming Tutorial (Bahasa Indonesia) - YouTube," 21-Mar-2022. [Online]. Available: https://www.youtube.com/watch?v=JjQwNlz_NTo. [Accessed: 13-Sep-2022].

[10] "C++ Data Types - GeeksforGeeks." [Online]. Available: https://www.geeksforgeeks.org/c-data-types/. [Accessed: 13-Sep-2022].

[11] "How to Implement Performance Metrics in CUDA C/C++ | NVIDIA Developer Blog." [Online]. Available: https://devblogs.nvidia.com/how-implement-performance-metrics-cuda-cc/. [Accessed: 11-Okt-2022].

[12] "Parallel Speedup — Parallel Computing Concepts." [Online].

Available: http://selkie.macalester.edu/csinparallel/modules/IntermediateIntroduction/build/html/ParallelSpeedup/ParallelSpeedup.html. [Accessed: 09-Sep-2022].
[13] “Overhead of Parallelism - Chanaka Balasooriya - Medium.” [Online]. Available: https://medium.com/@chanakadkb/overhead-of-parallelism-d1d3c43abadd. [Accessed: 16-Sep-2022].
[14] “Introduction to Parallel Computing.” [Online]. Available: https://computing.llnl.gov/tutorials/parallel_comp/. [Accessed: 16-Sep-2022].
[15] V. Hegde and S. Usmani, "Parallel and Distributed Deep Learning." [Online]. Available: https://web.stanford.edu/~rezab/classes/cme323/S16/projects_reports/hedge_usmani.pdf. [Accessed: 01-Sep-2022]