

Computer Science Extended Essay: Optimizing Shortest Path Algorithms in Weighted Graphs using Priority Queue Implementation

Research Question:

How Does the Use of a Priority Queue and the Implementation of Bellman-Ford and Dijkstra's Algorithm Affect the Time Complexity of Solving the Shortest Path Problem in Weighted Graphs?

Candidate Code: jpk691

Essay Word Count: 3999

Table of Contents

1. Introduction.....	1
2. Background Information.....	2
2.1 Weighted Graphs and The Shortest Path Problem.....	2
2.2 Shortest-Path Algorithms.....	3
2.2.1 Dijkstra's Algorithm.....	3
2.2.2 Bellman-Ford Algorithm.....	8
2.3 Priority Queues.....	13
3. Hypothesis and Applied Theory.....	15
4. Experimental Methodology.....	17
4.1 Weighted Graph Used.....	17
4.2 Independent Variables.....	20
4.3 Dependent Variables.....	21
4.4 Controlled Variables.....	21
4.5 Procedure.....	22
5. The Experimental Results.....	23
5.1 Tabular Data Presentation.....	23
5.2 Graphical Representation of Data.....	23
5.3 Data Analysis.....	26
5.3.1 Analyzing Dijkstra's Algorithm.....	26
5.3.2 Analyzing Bellman-Ford Algorithm.....	28
6. Limitations.....	30
7. Conclusion.....	32
8. Works Cited.....	34
9. Appendices.....	38
Appendix A: Python Code.....	38
Appendix B: Raw Data — Execution Times For Each Algorithm.....	45
Appendix C: Definitions and Key Terms.....	49

1. Introduction

Finding the most optimal route, whether it be networks, transportation or information flow, lies at the heart of efficient resource utilization and problem-solving in various domains. As the scale and complexity of these systems continue to expand, algorithms that determine the shortest path between points become indispensable. The Shortest Path Problem (SPP) is a well-known optimization problem in graph theory that deals with determining the minimum weight path between two vertices in a weighted graph. The efficiency of solving the SPP varies greatly based on the algorithms and data structures employed. The problem is considered computationally challenging due to its high complexity, and different algorithms have been proposed to solve it efficiently.

This essay will focus on investigating the time complexity (refer to **Appendix C** for definition) of pathfinding algorithms at solving the SPP on weighted graphs. This essay will specifically explore *Dijkstra's Algorithm* and the *Bellman–Ford Algorithm*, which are two prominent shortest-path algorithms. Both algorithms will also be compared with the implementation of a priority queue, which is an abstract data structure which, in the context of this essay, is used to keep track of the vertices to be explored. Which gives rise to the research question: “*How does the use of a priority queue and the implementation of Bellman-Ford and Dijkstra's algorithm affect the time complexity of solving the shortest path problem in weighted graphs?*”.

2. Background Information

2.1 Weighted Graphs and The Shortest Path Problem

A graph consists of points (vertices) connected by lines (edges), in weighted graphs, each edge has an associated weight. This weight can represent various things depending on the context, commonly distance or cost. There are two types of weighted graphs: directed (digraphs) and undirected graphs. Digraphs have edges that point from one vertex to another in a specific direction. Undirected weighted graphs have bidirectional edges that can be traversed in either direction between two vertices. Weighted digraphs are useful for modeling one-way relationships, such as transportation networks. Undirected weighted graphs are suited for modeling mutual relationships like social networks (example shown in **Figure 1**).




Figure 1: *Undirected Weighted Graph Consisting of 9 Vertices* (Virginia Tech: Department of Computer Science)

When the sum of the edge weights between two vertices is minimum (relative to all other possible routes), it is considered the shortest path. Given a weighted digraph $G = (V, E, w)$, to find the shortest path from a source vertex s to and an end vertex e , a path P from the start to end vertex needs to be found that minimizes the function:

$$w(P) := \sum_{u \rightarrow v \in P} w(u \rightarrow v)$$

(Erickson 273)

There are two main variations of the SPP: single-source (SSSP) and all-pairs (APSP), see **Appendix C** for definitions. This essay will be specifically focusing on the SSSP⁺ problem, where all weights are positive (i.e. $w : E \rightarrow \mathbb{R}^+$), as the APSP is comparatively more computationally taxing and time-consuming.

2.2 Shortest-Path Algorithms

2.2.1 Dijkstra's Algorithm

The most popular shortest-path algorithm widely used in applications such as Google Maps. It works by iteratively selecting the vertex with minimum weight, updating distances to its neighbors by considering the weights of the connecting edges, and repeating this process until all vertices have been visited.

Dijkstra's Algorithm is an example of a greedy algorithm (refer to **Appendix C** for definition). One stipulation to using the algorithm is that the graph needs to have a nonnegative weight on every edge (Abiy et al.).

In the context of this essay, a “naive” algorithm (refer to **Appendix C** for definition) is one with no optimizations (i.e. no priority queue).

```

function NaiveDijkstra(graph, source):
    dist = {} // dictionary to store the shortest distances
    visited = {} // dictionary to keep track of visited nodes

    for each vertex in graph:
        dist[vertex] = infinity // initialize all distances to infinity
        visited[vertex] = false // mark all nodes as not visited

    dist[source] = 0 // distance from source to itself is 0

    while there are unvisited nodes:
        current = node with the minimum distance from dist dictionary

        visited[current] = true // mark the current node as visited

        for each neighbor of current:
            if visited[neighbor] is false:
                // calculate the new distance from source to neighbor
                newDistance = dist[current] + edge_weight(current, neighbor)

                if newDistance < dist[neighbor]:
                    dist[neighbor] = newDistance // update the distance

    return dist

```

Figure 2: Pseudocode for Naive Dijkstra's Algorithm (ChatGPT, 2023)

The following is a breakdown of the pseudocode (**Figure 2** above):

1. Declaration: The 'dist' dictionary stores the shortest distances between source and every other vertex in the graph, while 'visited' keeps track of visited vertices.
2. Initialization: The distance from the source vertex to every other vertex is initially set as ∞ (unknown), except the source vertex, which is set to 0.
3. The algorithm will run while there are unvisited vertices.
 - a. Selecting Vertex With Minimum Distance: The algorithm selects the vertex 'current' with the minimum distance from the 'dist' dictionary among the unvisited vertices in each iteration, or $\text{current} = \arg \min(\text{dist}[\text{vertex}])$. The current vertex is marked as visited or *true*.
 - b. For each neighbor vertex of the current vertex:
 - i. Check if the neighbor vertex is unvisited: If `visited[neighbor]` is *false*, it means it has not been visited. 'newDistance' is a variable that will calculate and store the new distance from the source to the neighbor vertex.

- ii. The calculated distance is compared with the current distance in the 'dist[neighbor]' dictionary. If smaller, a shorter path has been found from the source vertex to the neighbor. 'dist[neighbor]' is updated with the new, smaller distance. Denoted mathematically as
- $$\text{dist}[\text{neighbor}] = \min(\text{dist}[\text{neighbor}], \text{newDistance}).$$
- 'min' is used here to choose the smaller of the two distances. By updating 'dist[neighbor]' with the new smaller distance, the edge between the current vertex and its neighbor vertex is "relaxed" (refer to **Appendix C** for definition). 'dist' is returned, which now contains all the shortest paths from the source vertex to every other vertex.

Consider the following graph:


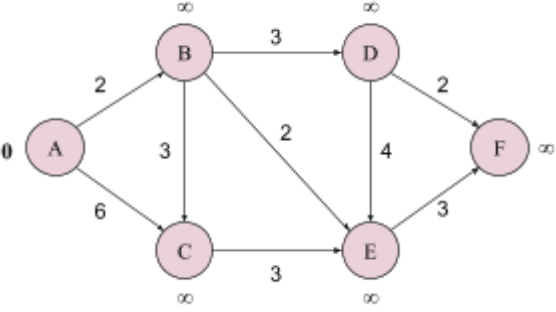
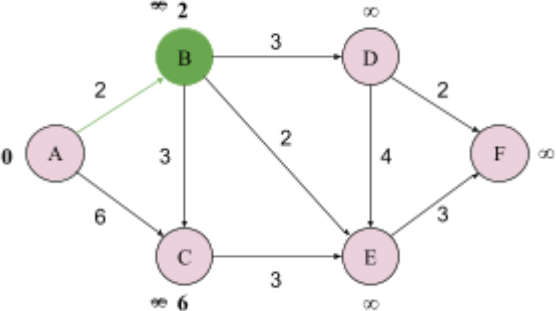
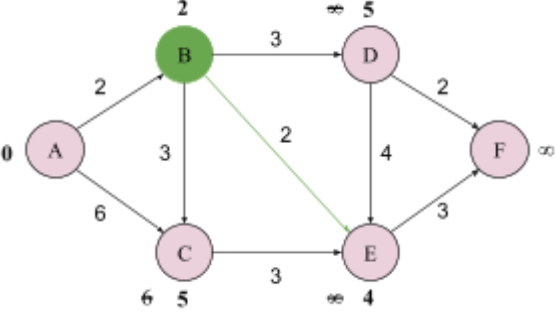
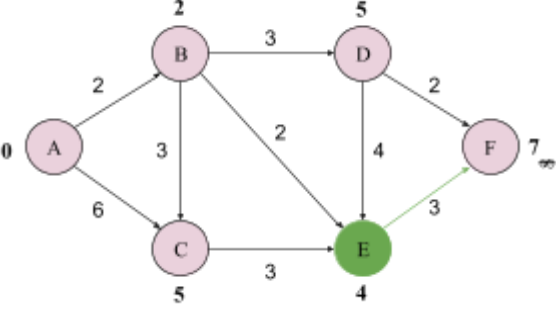
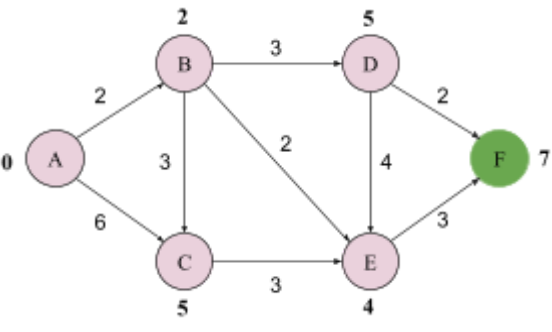


Figure 3: Example Weighted Digraph with 6 Vertices and 9 Edges

Below is an example use of the algorithm against **Figure 3**.

Graph	Explanation
 <p>Figure 4: Distances Initially Unknown</p>	<p>The distance from every vertex to the start vertex is initialized to be ∞ except the starting vertex itself, which would be 0.</p> <p>The only vertex visited so far is the starting vertex itself, A.</p> <p>$\text{dist} = \{A: 0, B: \infty, C: \infty, D: \infty, E: \infty, F: \infty\}$ $\text{visited} = \{A: \text{True}, B: \text{False}, C: \text{False}, D: \text{False}, E: \text{False}, F: \text{False}\}$</p>
 <p>Figure 5: Exploring Neighbors of Vertex A</p>	<p>The neighbors of the current vertex (A) are explored, which are B and C. Now, the edge from A to B is relaxed. Currently, the distance known from A to B is ∞, and the new potential distance is $0 + 2 = 2$. Since $2 < \infty$, the shortest distance from A to B is updated as 2. This can also be represented mathematically as:</p> $\begin{aligned} \text{dist}[B] &= \min(\text{dist}[B], \text{dist}[A] + \text{edge_weight}(A, B)) \\ &= \min(\infty, 0 + 2) \\ &= 2 \end{aligned}$ <p>The edge from A to C is to be relaxed, the known distance from A to C is ∞, the new potential distance is $0 + 6 = 6$. Since $6 < \infty$, the shortest distance from A to C is updated as 6. This can be represented mathematically in a similar way as the above.</p> <p>Now, the current vertex is set to the neighbor vertex with the smallest distance, which is B (since $2 < 6$).</p> <p>$\text{current} = B$ $\text{visited} = \{A: \text{True}, B: \text{True}, C: \text{False}, D: \text{False}, E: \text{False}, F: \text{False}\}$ $\text{dist} = \{A: 0, B: 2, C: 6, D: \infty, E: \infty, F: \infty\}$</p>

Graph	Explanation
 <p>Figure 6: Exploring Neighbors of Vertex B</p>	<p>B is marked as visited in the ‘visited’ dictionary by setting it to <i>true</i>. The neighbors of B are to be explored, which are C, D and E. The process of edge relaxation is repeated. Now, the edge from B to C is relaxed. The current known shortest distance from A to C is 6, but the new potential distance through B is $2 + 3 = 5$. Since $5 < 6$, the shortest distance to C can be updated to be 5. Denoted mathematically as:</p> $\text{dist}[C] = \min(\text{dist}[C], \text{dist}[B] + \text{edge_weight}(B, C))$ $= \min(6, 2 + 3)$ $= 5$ <p>The edge from B to E is to be relaxed, the current distance is ∞ from A to E, the new potential distance through B is $2 + 2 = 4$. Since $4 < \infty$, the shortest distance from A to E is 4.</p> <p>Finally, the edge from B to D is relaxed; the current distance is ∞, and the new potential distance through B is $2 + 3 = 5$. Since $5 < \infty$, the shortest distance from A to D is 5. Now, the current vertex is set to the neighbor vertex with the smallest distance, which is E (since $2 < 3$).</p> <p>$\text{dist} = \{A: 0, B: 2, C: 5, D: 5, E: 4, F: \infty\}$ $\text{visited} = \{A: \text{True}, B: \text{True}, C: \text{True}, D: \text{True}, E: \text{True}, F: \text{False}\}$</p>
 <p>Figure 7: Exploring Neighbors of Vertex E</p>	<p>$\text{current} = E$</p> <p>Now explore the neighbors of E, which would be only F. The edge from E to F is to be relaxed. The current known distance from A to F is ∞ and the potential distance will be $2 + 2 + 3 = 7$. Since $7 < \infty$, the shortest distance from A to F is updated to be 7.</p> <p>The current vertex is set to the neighbor vertex with the smallest distance, which is simply F.</p> <p>$\text{dist} = \{A: 0, B: 2, C: 5, D: 5, E: 4, F: \infty\}$</p>

Graph	Explanation
 <p data-bbox="245 629 743 696">Figure 8: End of Loop, Shortest Paths Found</p>	<p data-bbox="810 280 1342 387">current = F visited = {A: True, B: True, C: True, D: True, E: True, F: True}</p> <p data-bbox="810 436 1374 611">Now that all the vertices have been visited, the loop terminates and the 'dist' dictionary containing the shortest paths from A to every other vertex is returned. As shown below:</p> <p data-bbox="810 651 1374 689">dist = {A: 0, B: 2, C: 5, D: 5, E: 4, F: 7}</p>

This essay will use Big-O notation (refer to **Appendix C** for definition) for measuring time complexity (as opposed to Big-Ω or Big-Θ), as it aims to evaluate the performances of both algorithms under maximum stress (worst-case).

The worst-case time complexity of Dijkstra's Algorithm without the use of a priority queue is $O(V^2)$, where V represents the number of vertices in the graph. In each iteration, the algorithm needs to find the vertex with minimum distance among the unvisited vertices. In a naive implementation, this requires iterating over all vertices, resulting in a time complexity of $O(V)$ for this operation. As this process is repeated for each vertex ($O(V)$), the overall time complexity becomes $O(V^2)$ (since $O(V) \times O(V)$).

2.2.2 Bellman-Ford Algorithm

Bellman-Ford's algorithm finds applications in networking, particularly in a distance-vector routing protocol. This protocol decides how to route packets of data on a network (Chumbley et al.).

Similar to Dijkstra's Algorithm, it uses the idea of relaxation but doesn't use [it] with [a] greedy technique (Morampudi). It works by keeping track of weight distance from the origin

and previous node in the shortest path, looping over the edges/connections for n times (n being the number of nodes/vertices), and updating the fastest route to the destination (stevenard). In comparison to Dijkstra's algorithm, the Bellman-Ford algorithm admits or acknowledges the edges with negative weights (Magzhan and Mat).

```
function bellmanFord(graph, source):
    distance = {} // dictionary to store the shortest distance from the source

    // Step 1: Initialization
    for each vertex v in graph:
        distance[v] = infinity
    distance[source] = 0

    // Step 2: Relax edges repeatedly
    for i from 1 to |V|-1: // |V| is the number of vertices in the graph
        for each edge (u, v, w) in graph:
            if distance[u] + w < distance[v]:
                distance[v] = distance[u] + w

    // Step 3: Check for negative weight cycles
    for each edge (u, v, w) in graph:
        if distance[u] + w < distance[v]:
            return "Graph contains a negative weight cycle"

    return distance
```

Figure 9: Pseudocode for Naive Bellman-Ford Algorithm (ChatGPT, 2023)

The following is a breakdown of the pseudocode (**Figure 9** above):

1. Declaration: A dictionary called 'distance' is created to store the shortest distance from the source vertex to every other vertex.
2. Initialization: Set the distance from the start vertex to every other vertex to be ∞ , except the start vertex, which is set to 0.
3. Continuously Relax Edges: Iteratively update the distances until the optimal solution is found. Let $|V|$ be the total number of vertices in the graph. This process is repeated $|V| - 1$ times.
 - a. For every edge (u, v, w) in the graph, where u and v are vertices and w is the weight of the edge, the distance to v is checked to see if it can be minimized by going through u .

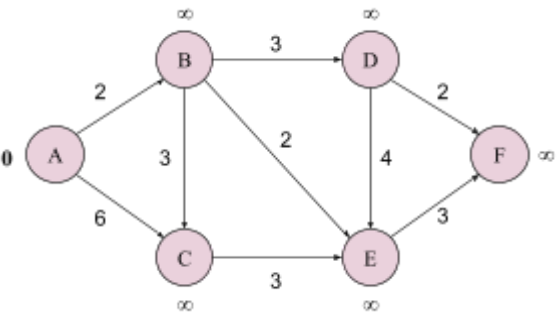
- i. If this condition is satisfied, the distance of vertex v is updated to be the sum of the distance to vertex u and the weight of the edge (u, v) .

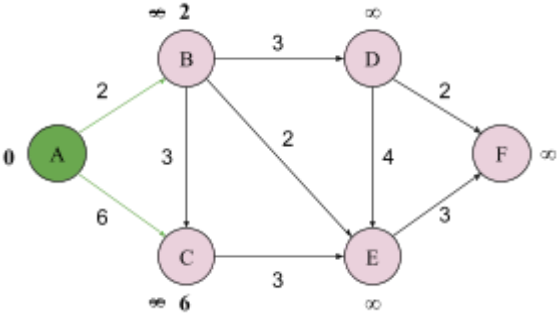
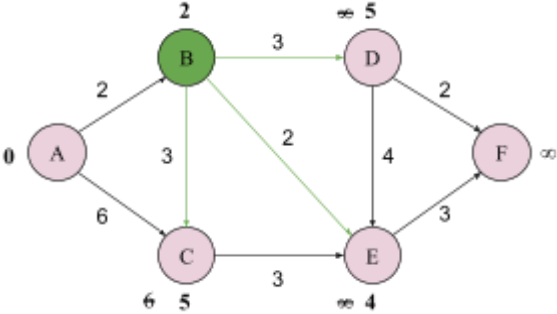
The relaxation can be denoted mathematically as:

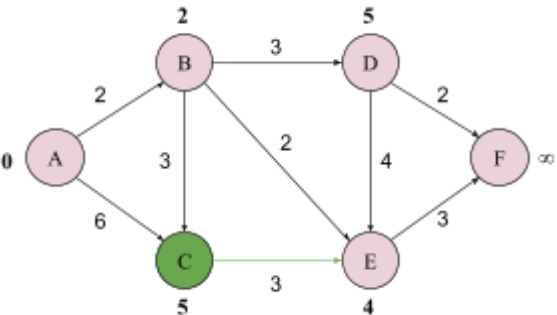
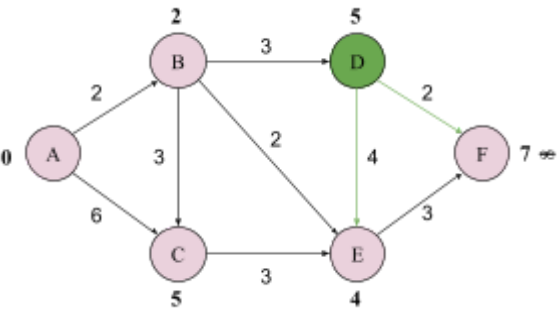
$$\text{distance}[v] = \min(\text{distance}[v], \text{distance}[u] + w)$$

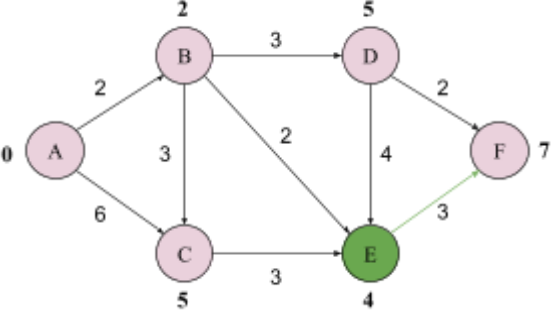
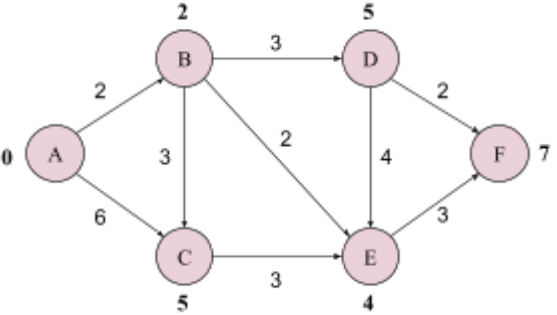
4. Check for negative weight cycles: If the sum of the weights of the edges along the cycle is negative, it is defined as a negative weight cycle. For every edge (u, v, w) , the distance from the start vertex to vertex v ($\text{distance}[u] + w$) is checked to see if it is smaller than the current shortest distance to v ($\text{distance}[v]$). If a negative weight cycle is present, the algorithm will fail to provide a solution.
5. If the condition above is not satisfied, the 'distance' dictionary is returned which now contains all the shortest paths.

The algorithm will not be demonstrated using a graph with negative weights as it is not relevant to this experiment. The following example demonstrates the algorithm's use on the same graph (**Figure 3**) as discussed in the previous section.

Graph	Explanation
 <p>Figure 10: Distances Initially Unknown</p>	<p>The distance from every vertex to the start vertex is initialized to be ∞ except the starting vertex itself, which would be 0.</p> <p>The only vertex visited so far is the starting vertex itself, A.</p> <p>$\text{distance} = \{A: 0, B: \infty, C: \infty, D: \infty, E: \infty, F: \infty\}$</p>

Graph	Explanation
 <p data-bbox="341 965 675 999">Figure 11: First Iteration</p>	<p data-bbox="836 297 1369 405">The main loop is entered, where edges in the graph are continuously relaxed for a total of 5 iterations in this case.</p> <p data-bbox="836 443 1326 477">For every edge (u, v, w) in the graph:</p> <p data-bbox="836 512 1198 629"> $(A, B, 2): \text{distance}[A] + 2$ $= 0 + 2$ $= 2$ </p> <p data-bbox="836 636 1369 779">Now if this distance is smaller than the current $\text{distance}[B]$, which is infinity, the distance from A to B is updated as 2 ($\text{distance}[B] = 2$).</p> <p data-bbox="836 815 1198 931"> $(A, C, 6): \text{distance}[A] + 6$ $= 0 + 6$ $= 6$ </p> <p data-bbox="836 938 1369 1055">Similar to the above, $\text{distance}[C]$ is currently infinity, which is greater than 6. Thus $\text{distance}[C] = 6$.</p> <p data-bbox="836 1090 1369 1198">This is the first iteration of Step 3 (refer to pseudocode). The known distances are as shown below.</p> <p data-bbox="836 1234 1326 1312"> $\text{distance} = \{A: 0, B: 2, C: 6, D: \infty, E: \infty, F: \infty\}$ </p>
 <p data-bbox="325 1839 691 1872">Figure 12: Second Iteration</p>	<p data-bbox="836 1350 1198 1467"> $(B, C, 3): \text{distance}[B] + 3$ $= 2 + 3$ $= 5$ </p> <p data-bbox="836 1473 1369 1617">Now, compare this with the current $\text{distance}[C]$, which is 6. Since $5 < 6$, $\text{distance}[C]$ is updated to be 5, which is the shorter path.</p> <p data-bbox="836 1659 1198 1776"> $(B, E, 2): \text{distance}[B] + 2$ $= 2 + 2$ $= 4$ </p> <p data-bbox="836 1783 1369 1861">Since $4 < \infty$, $\text{distance}[E]$ is updated to be 4.</p> <p data-bbox="836 1897 1198 2013"> $(B, D, 3): \text{distance}[B] + 3$ $= 2 + 3$ $= 5$ </p>

Graph	Explanation
	<p>Since $5 < \infty$, $\text{distance}[D]$ is updated to be 5.</p> <p>This is the second iteration of Step 3. The new distances are shown as below:</p> <p>$\text{distance} = \{A: 0, B: 2, C: 5, D: 5, E: 4, F: \infty\}$</p>
 <p>Figure 13: Third Iteration</p>	<p>$(C, E, 3): \text{distance}[C] + 3$ $= 5 + 3$ $= 8$</p> <p>Since 8 is greater than the current $\text{distance}[E]$, which is 4. The distance dictionary is not updated.</p> <p>This is the third iteration of Step 3. The new distances are shown as below:</p> <p>$\text{distance} = \{A: 0, B: 2, C: 5, D: 5, E: 4, F: \infty\}$</p>
 <p>Figure 14: Fourth Iteration</p>	<p>$(D, E, 4): \text{distance}[D] + 4$ $= 5 + 4$ $= 9$</p> <p>Since 9 is greater than the current $\text{distance}[E]$, which is 4. The distance dictionary is not updated.</p> <p>$(D, F, 2): \text{distance}[D] + 2$ $= 5 + 2$ $= 7$</p> <p>Since 7 is less than the current value of $\text{distance}[F]$, which is ∞. $\text{distance}[F]$ is updated to be 7.</p> <p>This is the fourth iteration of Step 3. The new distances are shown as below:</p> <p>$\text{distance} = \{A: 0, B: 2, C: 5, D: 5, E: 4, F: 7\}$</p>

Graph	Explanation
 <p data-bbox="343 705 678 739">Figure 15: Fifth Iteration</p>	<p data-bbox="837 324 1197 448"> $(E, F, 3): \text{distance}[E] + 3$ $= 4 + 3$ $= 7$ </p> <p data-bbox="837 481 1380 672"> No updates occur in the fifth and final iteration, all edges have been successfully relaxed. This is the fifth iteration of Step 3. The final distances are shown as below: </p> <p data-bbox="837 705 1388 784"> $\text{distance} = \{A: 0, B: 2, C: 5, D: 5, E: 4, F: 7\}$ </p>
 <p data-bbox="207 1176 813 1209">Figure 16: End of Loop, Shortest Paths Found</p>	<p data-bbox="837 828 1308 907"> Now, check for any negative weight cycles, for which there are none </p> <p data-bbox="837 940 1372 1019"> Finally, the ‘distance’ dictionary is returned containing all the shortest paths: </p> <p data-bbox="837 1052 1388 1131"> $\text{distance} = \{A: 0, B: 2, C: 5, D: 5, E: 4, F: 7\}$ </p>

The time complexity of the naive Bellman-Ford Algorithm is $O(VE)$. Where V is the number of vertices and E is the number of edges. The algorithm repeatedly relaxes each edge for $V - 1$ iterations, in each iteration, it checks all edges E within the graph and updates the distance to each vertex in case a shorter path is found. Resulting in $(V - 1) \times E$ relaxation operations, the constant -1 can be ignored. Resulting in a total time complexity of $O(VE)$.

2.3 Priority Queues

A priority queue is a variation of the traditional queue data structure where each element in the queue is associated with a priority value, elements are attended to in the queue based on this priority value. Conventionally, the element at the front of the queue has the highest priority value.

Priority queues are usually implemented using heaps (refer to **Appendix C** for definition).

There are two different types of heaps, a max heap and min heap. In a min heap, the value of the parent node is less than or equal to the value of the child node, for all nodes, this property is known as the heap invariant. The max heap is simply the opposite. The value of the nodes in the context of this experiment being distances. **Figure 16** shows a min heap of height 3 satisfying the heap invariant, while **Figure 17** shows a violated heap invariant where the child of parent node 3 is less than its parent.




Figure 17: *Valid Min Heap*




Figure 18: *Invalid Min Heap*

The two main operations of the priority queue that concern this experiment are insertion and extraction (removal). The worst-case time complexity of both operations is $O(\log_2 N)$ (Garg), where N is the number of elements in the heap. Since the min heap can be visualized as a complete binary tree, the tree's height can be expressed as a logarithm of N (e.g. **Figure 18** above has 7 nodes, therefore $\lceil \log_2 7 \rceil = 3$). Leading to a logarithmic time complexity for both insertion and extraction.

In this experiment, the priority queue being implemented into the algorithms will be using a min heap, as Python has a built-in data structure which is essential to prioritize the smallest distance at each localized step of the algorithm, in order to find an optimal path.

3. Hypothesis and Applied Theory

The experiment will find the relationship between execution time in nanoseconds (y) and the number of vertices in the digraph (x). By increasing the number of vertices, a clear correlation can be drawn between both variables and how the relationship differs when a priority queue is implemented within both algorithms.

The graph used in this experiment is such that $|E| > |V|$. Thus, recalling the time complexities for both algorithms, it is hypothesized that for naive implementations, Dijkstra's Algorithm will run with a lower execution time. The number of edges in the graph are such that $|V| \times 1.5 \approx |E|$, so the time complexity of the Bellman-Ford can be approximated to be a quadratic as $O(V \times V \times 1.5) \approx O(V^2)$. Therefore, for both naive algorithms, there is predicted to be a quadratic relationship between x and y .

For Dijkstra's Algorithm, the priority queue is queried to extract the vertex with the smallest current distance, taking $O(\log_2 V)$ time. Then, the algorithm performs edge relaxation on the neighboring vertices of the extracted vertex, potentially updating their current distances. This is repeated until all vertices have been processed or the priority queue is empty, improving the time complexity to $O((V + E) \log_2 V)$, as the extraction of the vertex with the smallest distance becomes optimized through the logarithmic time complexity of the min heap extraction.

As mentioned in **Section 2.2.2**, the number of relaxation operations in the Bellman-Ford Algorithm is $O(VE)$. With a priority queue, each relaxation operation involves inserting or extracting an element in the priority queue, which is time complexity $O(\log_2 V)$, leading to the relaxation operations having a time complexity of $O((VE) \log_2 V)$. The priority queue

benefits the selection of the vertex with the smallest distance in each iteration which can be done in $O(\log_2 V)$ time. Combining this time complexity with the relaxation operations results in $O((VE) \log_2 V \times \log_2 V) = O((V+E) \log_2 V)$.

It appears that both min heap implemented algorithms have the same time complexity, however, the graph that will be used in the experiment is such that there are no negative weights, and Dijkstra's Algorithm tends to perform better in this environment. Furthermore, Dijkstra's Algorithm does not iterate through all the edges, while the Bellman-Ford Algorithm does multiple times.

Thus, it is hypothesized that with the priority queue implementations, both algorithms will reduce in execution time, specifically, Dijkstra's Algorithm will run with a lower execution time compared to the Bellman-Ford Algorithm. For both priority queue implemented algorithms, there is predicted to be a logarithmic relationship.

4. Experimental Methodology

4.1 Weighted Graph Used

As a resident of Singapore, the author of this paper frequently relies on the public bus network for their transportation needs. Therefore, for this experiment, they decided to represent the Singapore Bus Network as a weighted digraph, which can be used to evaluate the execution times of the algorithms. The vertices being bus stops and the weights being the travel distances between the stops in kilometers.

The dataset used in the experiment's code originates from the Land Transport Authority; this data was mostly scraped from the website and compiled into a data repository on Github (Aun). Two files were saved locally on the author's system, stops.json and services.json.

```
8 # Constants for Haversine formula
9 AVG_EARTH_RADIUS = 6371 # in kilometers
10
11 # Step 1: Parse services.json
12 with open('services.json') as services_file:
13     services_data = json.load(services_file)
14
15 # Step 2: Parse stops.json
16 with open('stops.json') as stops_file:
17     stops_data = json.load(stops_file)
18
19 # Step 3: Create a dictionary of stop IDs to coordinates
20 stop_coordinates = {}
21 for stop_id, stop_info in stops_data.items():
22     longitude, latitude, _, _ = stop_info
23     stop_coordinates[stop_id] = (float(latitude), float(longitude))
```

Figure 19: Code Snippet Showing The Parsing of JSON Data and Dictionary Population

Both files services.json and stops.json are read and deserialized into Python objects.

The services.json file contains data about bus services in Singapore, each identified by a distinct service ID, such as 3 or 4, and including information about its name and routes. The stops.json file contains information regarding the bus stops themselves, such as bus stop ID,

latitude and longitude. A dictionary called `stop_coordinates` is created and is populated with bus stop IDs as the keys and their corresponding latitude and longitude coordinates as values.

```

25 # Function to calculate Haversine distance
26 def haversine_distance(coord1, coord2):
27     lat1, lon1 = coord1
28     lat2, lon2 = coord2
29
30     dlat = math.radians(lat2 - lat1)
31     dlon = math.radians(lon2 - lon1)
32
33     a = math.sin(dlat / 2) * math.sin(dlat / 2) + math.cos(
34         math.radians(lat1)) * math.cos(math.radians(lat2)) * math.sin(
35         dlon / 2) * math.sin(dlon / 2)
36     c = 2 * math.atan2(math.sqrt(a), math.sqrt(1 - a))
37
38     distance = AVG_EARTH_RADIUS * c
39     return distance

```

Figure 20: Code Snippet Showing The Haversine Distance Formula

The weights are calculated using the Haversine distance formula (shown below), which accurately calculates the distances between two points on the surface of a sphere (approximating the shape of Earth as a sphere) given the latitude and longitude of the two points.

$$d = 2r \arcsin \left(\sqrt{\sin^2 \left(\frac{\Phi_2 - \Phi_1}{2} \right) + \cos(\Phi_1) \cos(\Phi_2) \sin^2 \left(\frac{\lambda_2 - \lambda_1}{2} \right)} \right)$$

where d = distance between bus stops, r = radius of earth, Φ_1 = latitude of first point, Φ_2 = latitude of second point, λ_1 = longitude of first point, λ_2 = longitude of second point

(Sydorenko)

The weights do not reflect the real-life distances between the bus stops but rather provide an approximation.

```

41 # Step 4: Create an empty directed graph
42 graph = nx.DiGraph()
43
44 # Step 5: Add edges to the graph with weights
45 stop_ids = list(stop_coordinates.keys())[:len(stop_coordinates)//1]
46
47 for service_id, service_info in services_data.items():
48     routes = service_info['routes']
49     for route in routes:
50         for i in range(len(route) - 1):
51             start_stop_id = route[i]
52             end_stop_id = route[i + 1]
53
54             if start_stop_id in stop_ids and end_stop_id in stop_ids:
55                 start_coordinates = stop_coordinates[start_stop_id]
56                 end_coordinates = stop_coordinates[end_stop_id]
57                 distance = haversine_distance(start_coordinates, end_coordinates)
58                 graph.add_edge(start_stop_id, end_stop_id, weight=distance)

```

Figure 21: Code Snippet Showing The Creation of The Digraph

A digraph is created using the NetworkX library. The `services_data` dictionary is iterated over, and for each service, the routes are examined. Within each route, consecutive pairs of stops are considered. If both the start and end stops are present in the `stop_ids` list, their corresponding coordinates are retrieved from the `stop_coordinates` dictionary. The distance between the start and end coordinates are calculated with the aforementioned haversine distance function. Finally, an edge is added to the graph, connecting the start and end stop IDs, with the calculated distance as the weight. Resulting in **Figure 22** below:




Figure 22: Visualization of All The Bus Stops in Singapore Using the Kamada-Kawai Layout

4.2 Independent Variables

The size of the weighted digraph will be changed, specifically, the number of vertices and edges in the graph. The idea is to iterate through increasing fractions of the stops.json file so that the number of vertices and edges in the graph can be increased, which is a convenient and logical way of changing the independent variables in this experiment. This will be done by changing this variable:

```
45. stop_ids = list(stop_coordinates.keys())[:len(stop_coordinates)//x]
```

By changing the x after the floor division operator, the fraction of the stops included in the graph can be controlled. For example, when $x = 10$, only one-tenth of the stops will be considered, resulting in a smaller graph size with a reduced number of vertices and edges.

The value of x in this experiment will be varied from 10 to 1, with 1 being inclusive of all 5083 bus stops. The purpose is to showcase algorithm performance under escalating computational stress, revealing the correlation between execution time and input size. **Figure 23** shows the various number of vertices and edges that will be used to test the algorithms:

Vertices	Edges
508	734
564	816
633	911
726	1012
847	1186
1010	1363
1269	1681
1694	2241
2540	3454
5083	7420

Figure 23: *Range of Values for Independent Variable*

4.3 Dependent Variables

The only dependent variable is the execution time for the given algorithm to find the shortest path from the source vertex to all the other vertices (SSSP) in nanoseconds. This will be measured using the `time.perf_counter_ns()` function from the `time` module in Python, which provides a high-resolution timer for accurate timing measurements in nanoseconds.

4.4 Controlled Variables

Variable	Description	Specification
IDE Used	The program will be run on the same IDE. To minimize variations in code execution and optimization, which can prevent systematic errors in execution times between the algorithms.	IDE: Visual Studio Code 1.79.0 (user setup) Python 3.11.4 <ul style="list-style-type: none">- NetworkX 3.1- json 2.0.9
Computer and OS	The program will be run on a Razer Book 13. This will minimize systematic and random error as there are no variations in hardware and OS.	OS: Microsoft Windows 11 Home Version 10.0.22621 Build 22621 Processor: Processor 11th Gen Intel(R) Core(TM) i7-1165G7 @ 2.80GHz, 2803 Mhz, 4 Core(s), 8 Logical Processor(s) Memory: 16GB DDR4 SDRAM 4267MHz
Start Vertex	The first bus stop id found in the <code>stops.json</code> file will be used, so that the algorithms are subjected to the same initial conditions, reducing random error in results due to different starting points.	id: 10009
Input Dataset	The same data will be used to construct the weighted graph, reducing systematic error and ensuring a consistent basis for algorithmic comparison.	<code>stops.json</code> and <code>services.json</code>

Variable	Description	Specification
Graph Representation	The graph will be represented using the NetworkX library, the same graph structure and connectivity are used consistently throughout this experiment.	nx.DiGraph()
Graph	The same graph will be used to benchmark the algorithms. The distances between each vertex will be kept constant, reducing systematic error.	Weighted Digraph consisting of 5083 vertices and 7420 edges in total

4.5 Procedure

1. Install the following library: NetworkX
2. Open the file ee.py (see **Appendix A**) in Visual Studio Code version 1.79.0. Place stops.json and services.json in the same folder as the program.
3. Run the program by clicking the arrow button or pressing F5.
4. Perform 10 trials for each value of the independent variable (re-run the program). For each trial, record the execution times of the Dijkstra's algorithm (naive), Bellman-Ford algorithm (naive), Dijkstra's algorithm with a priority queue, and Bellman-Ford algorithm with a priority queue. Note down the results in a spreadsheet software.

```

PS D:\> & C:/Users/ /AppData/Local/Microsoft/WindowsApps/python3.11.exe d:/holder.py
Number of vertices: 5083
Number of edges: 7420
Dijkstra's algorithm (without priority queue):
Execution time: 1574833200 nanoseconds
Bellman-Ford algorithm (without priority queue):
Execution time: 13775991300 nanoseconds
Dijkstra's algorithm with a priority queue:
Execution time: 20048000 nanoseconds
Bellman-Ford algorithm with a priority queue:
Execution time: 10272326900 nanoseconds

```

Figure 24: Results Outputted in the Terminal

5. Calculate the average execution times for each algorithm (naive and priority queue) based on the recorded data.

5. The Experimental Results

5.1 Tabular Data Presentation

Figure 25 below shows the average execution time in nanoseconds for each algorithm to find the SSSP in the weighted digraph with increasing vertices and edges. See **Appendix B** for the raw data.

Vertices	Edges	Average Execution Time (nanoseconds)			
		Dijkstra's Algorithm (Naive)	Bellman-Ford Algorithm (Naive)	Dijkstra's Algorithm (Priority Queue)	Bellman-Ford Algorithm (Priority Queue)
508	734	18277760	145268750	1823730	92678050
564	816	20924270	167572360	1902800	102746450
633	911	26766110	226592440	2109500	131542560
726	1012	33335310	270494600	2194940	159646660
847	1186	48589910	365931930	2879720	247789270
1010	1363	66669930	507861230	3299370	344395450
1269	1681	105192020	734278890	3651060	466298250
1694	2241	174840830	1373807150	6027200	1015853050
2540	3454	388870850	3093997800	9051150	2287816100
5083	7420	1598144270	13528484770	17772930	9412515240

Figure 25: *Average Execution Time for Each Algorithm*

5.2 Graphical Representation of Data

Execution time was chosen to be graphed against the number of vertices instead of edges as it better reflects the complexity of the graph and enables analysis of algorithm efficiency and scalability with increasing graph size.

A log scale for the x axis was due to the wide and unevenly distributed range of values. The reason for this uneven distribution was justified in **Section 4.2**. Using a standard scale would result in a compressed representation, restricting effective observation of trends and patterns.




Figure 26: Execution Time of Dijkstra's Algorithm (Naive) Against the Number of Vertices




Figure 27: Execution Time of Dijkstra's Algorithm (Priority Queue) Against the Number of Vertices




Figure 28: Execution Time of the Bellman-Ford Algorithm (Naive) Against the Number of Vertices

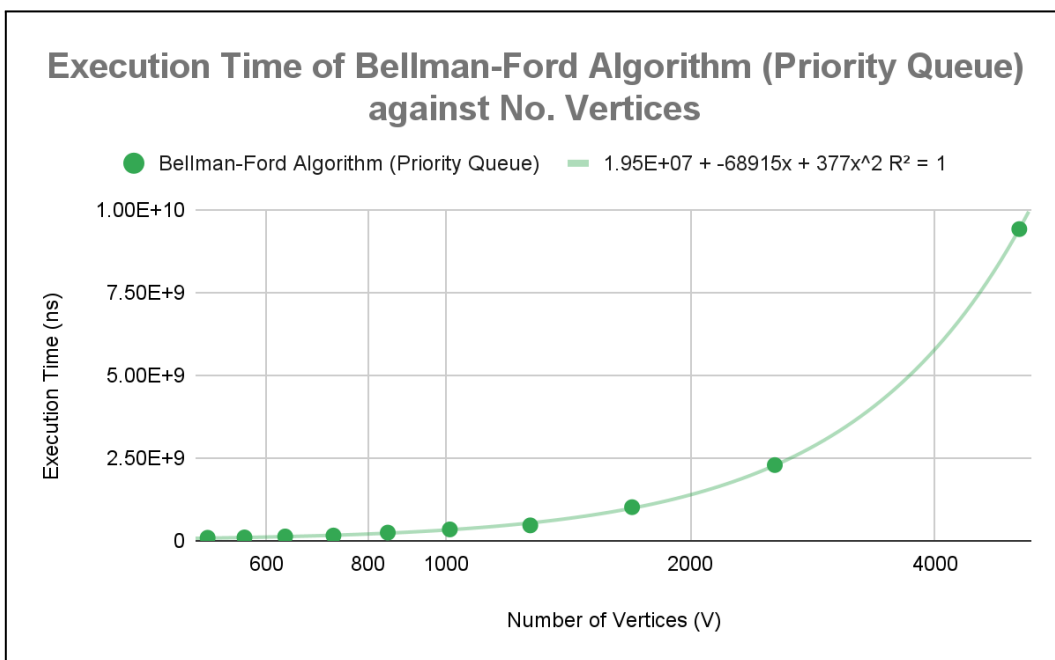


Figure 29: Execution Time of the Bellman-Ford Algorithm (Priority Queue) Against the Number of Vertices

5.3 Data Analysis

5.3.1 Analyzing Dijkstra's Algorithm

The hypothesis regarding the quadratic relationship between x and y has proven to be correct for the naive implementation. The R^2 value in **Figure 26** of exactly 1 indicates that the data points are a perfect fit to the quadratic curve. The equation in **Figure 26** can be re-written in the general form $ax^2 + bx + c$ as $63.3x^2 - 9042x + 8.43 \times 10^6$. The positive a coefficient suggests that as the number of vertices in the graph increases, the execution time of the algorithm will increase quadratically. The negative b value suggests that as the number of vertices increases, the execution time of the algorithm will decrease linearly (by a relatively small amount).

The hypothesis regarding the logarithmic relationship for the priority queue implemented algorithm is partially correct. **Figure 27** shows the trend line to have equation $(-5.23 \times 10^{-3})x^2 + 3576x - 226186$. The negligible a value for this equation suggests that the relationship between x and y is more linear. As mentioned previously, the time complexity for the priority queue implemented Dijkstra's Algorithm is $O((V + E) \log_2 V)$, using the previous approximation $|V| \times 1.5 \approx |E|$, the time complexity can be written as $O(2.5V \log_2 V)$. Graphing this out against $y = x$, it is observed that the time complexity is superlinear (**Figure 30** below). A superlinear relationship is a non-linear function that appears to grow linearly.




Figure 30: Graph of $O(2.5V \log_2 V)$ in Red and $y = x$ in Blue (Desmos)

The relationship between x and y is indeed logarithmic, but appears to be linear. This can be due to the relatively small input size for the algorithm; the graph used to benchmark the algorithm contained not as many vertices and edges needed to observe a clear logarithmic relationship. Furthermore, the graph used was sparse, which was in favor of the algorithm since there were fewer edges to explore. However, for the size of the dataset used, the relationship can be better modeled using a linear function.

Looking back at **Figure 25**, taking the average execution times for both algorithms for the entire graph (5083 vertices) and applying the percentage decrease formula:

$$\frac{\text{initial} - \text{final}}{\text{initial}} \times 100$$

$$\frac{1598144270 - 17772930}{1598144270} \times 100$$

Which is approximately **98.9%**. It is evident that the min heap priority queue implementation has greatly optimized Dijkstra's Algorithm, resulting in significantly improved execution times.

5.3.2 Analyzing Bellman-Ford Algorithm

The hypothesis regarding the quadratic relationship has proven to be correct for the naive implementation. The R^2 value in **Figure 28** indicates that the data points are a perfect fit to the quadratic curve. It behaves similarly to Dijkstra's Algorithm (naive), however, with a longer execution time overall.

The hypothesis regarding the logarithmic relationship for the priority queue implementation was incorrect, as the $R^2 = 1$ value shown in **Figure 29** suggests that the data points perfectly fit the shape of a quadratic curve. As a consequence of the sparse graph, the number of priority queue operations is performed a relatively smaller number of times compared to relaxation operations and iterations. As fewer edges need to be considered during each iteration, the $O(\log_2 V)$ component of the time complexity becomes negligible, which results in a more quadratic relationship.

The a value of a quadratic represents how wide/narrow the parabola is, it can be interpreted as the rate at which the execution time increases, a higher a value exemplifying a steeper increase in execution time against the number of vertices. Using the percentage decrease formula with the a values as the input, it is observed that the Bellman-Ford Algorithm decreased in execution time by about **34.2%** after the priority queue was implemented. This is due to several reasons. With the implementation of a priority queue, the algorithm can terminate early once all shortest paths have been found, it guarantees that once a vertex's shortest distance is finalized, it will not be updated further. Furthermore, relaxation operations are optimized as the vertex with the smallest distance is always selected first, which reduces the number of comparisons required, thus improving execution time.

Applying the second derivative test on the trendlines in **Figures 28** and **29**, the rate of execution time change with input size increase is determined, aiding in identifying the algorithm with lower execution times as input size expands.

Trendline in **Figure 28**:

$$y = 573x^2 - 285238x + 1.74 \times 10^8$$

$$\frac{dy}{dx} = 1146x - 285238$$

$$\frac{d^2y}{dx^2} = 1146$$

Trendline in **Figure 29**:

$$y = 377x^2 - 68915x + 1.95 \times 10^7$$

$$\frac{dy}{dx} = 754x - 68915$$

$$\frac{d^2y}{dx^2} = 754$$

As $1146 > 754$, it is concluded that the Bellman-Ford algorithm with a priority queue is more efficient and scalable as input size increases.

6. Limitations

There were various experiment-related limitations that could potentially have impeded the achievement of better results.

The graph used in this experiment is Singapore's Bus Network. In **Figure 22**, it can be seen that each vertex only points towards one other vertex, indicating the graph was sparse.

Empirically, most real-world graphs are sparse by nature. The number of edges is within a constant multiple of the number of vertices (Cook). The main issue is the algorithms may not encounter enough complexity to demonstrate the improvement of priority queue implementation. There were less relaxation and priority queue operations overall, which can be why all the algorithms mostly exhibited quadratic time complexity behaviors.

As Singapore is a small country, the graph contained only 5083 vertices and 7420 edges. This became an issue when collecting results for Dijkstra's Algorithm with a priority queue, as the relationship was observed to be initially linear when graphed without using a log scale (**Figure 31**). The use of a graph with more vertices (and preferably more edges) is speculated to eliminate this issue.

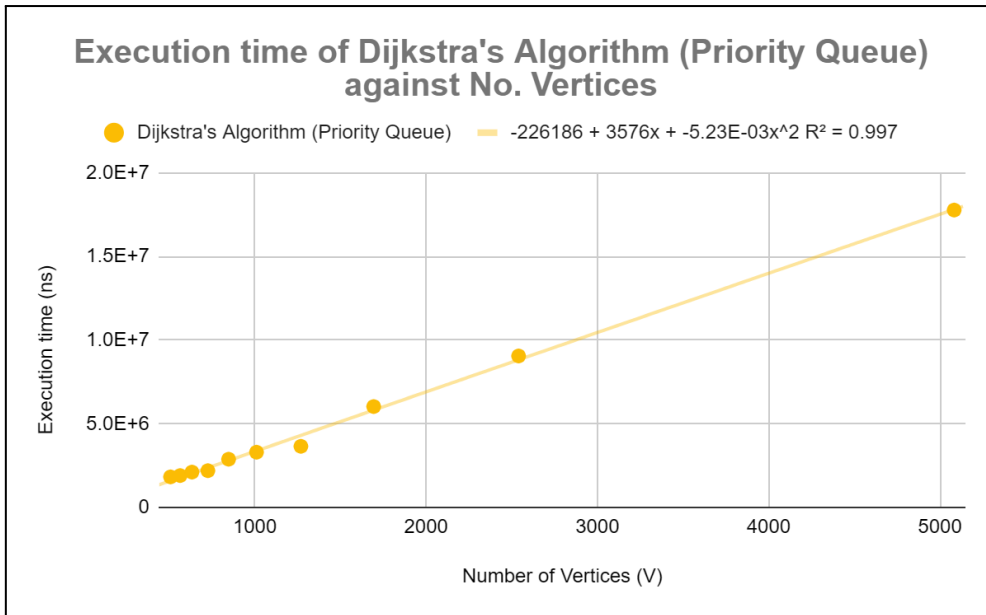


Figure 31: *Figure 27 without a Log Scale*

Hardware limitations can influence execution time, potentially causing fluctuations in results. The CPU of the laptop used was also being utilized by other processes such as Windows service host processes, which adds random error to the algorithm execution time. This is an inherent limitation to the device that was used in this experiment, as it is not meant to be a dedicated computing system for intricate algorithmic calculations.

7. Conclusion

The results clearly show that with the implementation of a min heap priority queue, both the Bellman-Ford and Dijkstra's Algorithm reduce in time complexity and execution time.

By implementing a min heap priority queue into Dijkstra's Algorithm, the selection and extraction of the minimum distance vertex during each iteration improves, reducing the number of relaxations and comparisons. This makes it more scalable as the number of vertices increases in a graph.

The Bellman-Ford Algorithm improves for the same reasons, however, it is not as pronounced due to the nature of the algorithm being dependent on the number of iterations and relaxation operations which overshadows the time complexity in sparse graph environments.

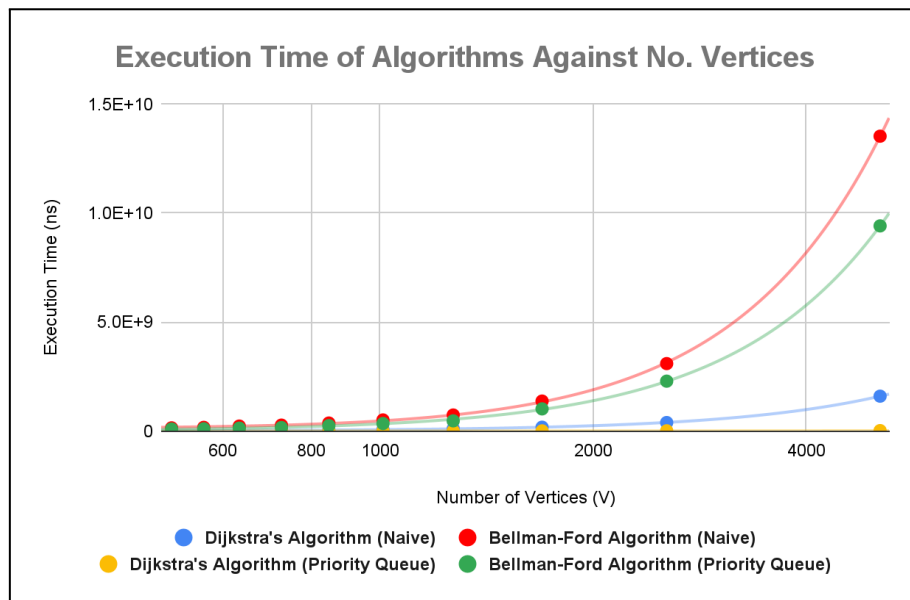


Figure 31: Combined Graph Showing the Execution Times of All Algorithms Against the Number of Vertices

Looking at **Figure 31** above, it can be concluded that Dijkstra's Algorithm is better at solving the SSSP problem in weighted graphs compared to the Bellman-Ford Algorithm. This is supported by the steepness of both the red and green lines (Bellman-Ford) as compared to the blue and yellow lines (Dijkstra), the blue and yellow lines are significantly more shallow, suggesting that Dijkstra's Algorithm is more efficient in terms of execution time as the number of vertices increases.

This paper hopes to prove useful to computer scientists collaborating with transportation planners and urban developers, providing them with a deeper understanding of how algorithmic optimizations can lead to more efficient transportation networks in urban environments. This research may serve as a foundational resource for future studies aiming to improve various graph-based algorithms in diverse applications, from network routing to logistics and beyond.

8. Works Cited

Abiy, Thaddeus, et al. "Dijkstra's Shortest Path Algorithm | Brilliant Math & Science Wiki."

Brilliant.org, 2016, brilliant.org/wiki/dijkstras-short-path-finder/. Accessed 25 July 2023.

Ashwani K. "Complete Tutorial on Big O (Big Oh) Notation - DevOpsSchool.com."

DevOpsSchool.com, 15 June 2021, www.devopsschool.com/blog/complete-tutorial-on-big-o-big-oh-notation/. Accessed 25 June 2023.

Aun, Chee. "SG Bus Data." *GitHub*, 26 May 2023, github.com/cheeaun/sdbusdata. Accessed

29 June 2023.

ChatGPT. "Prompt: Pseudocode Naive Bellman Ford using 'distance' dictionary for SSSP.

Using graph and source as the parameters. Please do not put any comments." *openai.com*,

2023. ChatGPT May. 24 Version.

<https://chat.openai.com/share/7a25254e-ed33-4cdd-bc6d-211c218619cc>. Accessed 22 Jun.

2023.

ChatGPT. "Prompt: Pseudocode Naive Dijkstra using dictionary for SSSP with no comments using a "visited" variable" *openai.com*, 2023. ChatGPT May. 24 Version.

<https://chat.openai.com/share/6cf54f75-c948-402e-aca9-2413aec013c2>. Accessed 20 Jun.

2023.

ChatGPT. "Prompt: The research question for my IB Computer Science Extended Essay is "How does the use of a priority queue and the implementation of Bellman-Ford and Dijkstra's algorithm affect the time complexity of solving the shortest path problem in weighted graphs?". For the weighted directed graph, I intend to use Singapore's Bus Network."

openai.com, 2023. ChatGPT Aug. 3 Version.

<https://chat.openai.com/share/f17c220d-4473-42bb-ae0c-eb426002e334>. Accessed 15 Aug.

2023.

Chumbley, Alex, et al. "Bellman-Ford Algorithm | Brilliant Math & Science Wiki."

Brilliant.org,

brilliant.org/wiki/bellman-ford-algorithm/#:~:text=%5Chspace%7B12mm%7D-

Accessed 2 July 2023.

Cook, John. "Real World Graphs." *Www.johndcook.com*, 21 July 2013,

www.johndcook.com/blog/2013/07/21/three-properties-of-real-world-graphs/.

Accessed 20 July 2023.

Desmos. "Desmos | Graphing Calculator." *Desmos*, www.desmos.com/calculator/azlovhgobb.

Accessed 8 Aug. 2023.

Erickson, Jeff. *Algorithms*. Erscheinungsort Nicht Ermittelbar] Jeff Erickson June 13, 2019,

p. 273, jeffe.cs.illinois.edu/teaching/algorithms/book/Algorithms-JeffE.pdf. Accessed

16 Sept. 2023.

Garg, Prateek. "Heaps and Priority Queues - Prateek Garg." *HackerEarth*, 2019,

www.hackerearth.com/practice/notes/heaps-and-priority-queues/. Accessed 13 July

2023.

Jaludi, Mariam. "Data Structures: Heaps." *The Startup*, 6 Oct. 2019,

medium.com/@mariam.jaludi/data-structures-heaps-b039868a521b. Accessed 5 July

2023.

Magzhan, Kairanbay, and Hajar Mat. "A Review and Evaluations of Shortest Path

Algorithms." *International Journal of Scientific & Technology Research*, vol. 2, no. 6,

Jan. 2013,

www.ijstr.org/final-print/june2013/A-Review-And-Evaluations-Of-Shortest-Path-Algorithms.pdf.

Accessed 11 July 2023.

Morampudi, Shanmukh Chowdary. "Bellman-Ford Algorithm." *Medium*, 25 Mar. 2021, shanmukhchowdary147.medium.com/bellman-ford-algorithm-21ac42a8cf3d. Accessed 28 June 2023.

Nvidia. "Shortest Path Problem." *NVIDIA Developer*, 29 July 2016, developer.nvidia.com/discover/shortest-path-problem. Accessed 20 June 2023.

Perplexity AI. "Prompt: What is a naive implementation of an algorithm?" Perplexity AI, n.d., <https://www.perplexity.ai/search/e5d545d8-3b3e-4dba-82fd-f576752f4643?s=c>. Accessed 19 Jun. 2023

SaturnCloud. "Relaxation of an Edge in Dijkstra's Algorithm: A Guide for Data Scientists | Saturn Cloud Blog." *Saturncloud.io*, 18 July 2023, saturncloud.io/blog/relaxation-of-an-edge-in-dijkstras-algorithm-a-guide-for-data-scientists/#:~:text=We%27ve%20learned%20that%20relaxation. Accessed 15 Aug. 2023.

stevenard, jb. "Bellman-Ford." *Medium*, 30 Sept. 2022, levelup.gitconnected.com/bellman-ford-6bd907c6c4c0. Accessed 28 June 2023.

Sydorenko, Daniil. "GitHub - DaniilSydorenko/Haversine-Geolocation: Get Distances between Two Points or Get Closest Position to Current Point. Based on the Haversine Formula." *GitHub*, 2020, github.com/DaniilSydorenko/haversine-geolocation. Accessed 16 Sept. 2023.

Team, Great Learning. "What Is Time Complexity and Why Is It Essential?" *GreatLearning Blog: Free Resources What Matters to Shape Your Career!*, 14 July 2022, www.mygreatlearning.com/blog/why-is-time-complexity-essential/#:~:text=Time%20complexity%20is%20defined%20as. Accessed 15 Aug. 2023.

Virginia Tech: Department of Computer Science. "Undirected Weighted Graph Containing 9 Vertices." *Virginia Tech*, 3 Dec. 2009,

courses.cs.vt.edu/~cs3114/Fall10/Notes/T22.WeightedGraphs.pdf. Accessed 3 June 2023.

YOON MI KIM. “Dijkstra Algorithm: Key to Finding the Shortest Path, Google Map to Waze.” *Medium*, Medium, 13 June 2019, medium.com/@yk392/dijkstra-algorithm-key-to-finding-the-shortest-path-google-map-to-waze-56ff3d9f92f0. Accessed 24 June 2023.

9. Appendices

Appendix A: Python Code

The code was initially generated using the Large Language Model (LLM) “ChatGPT”, however, underwent further modifications by the author (ChatGPT, 2023). Below is a table containing the significant prompts and their respective explanation as to why they were entered.

Prompt	Explanation
<p>The research question for my IB Computer Science Extended (sic) Essay is "How does the use of a priority queue and the implementation of Bellman-Ford and Dijkstra's algorithm affect the time complexity of solving the shortest path problem in weighted graphs?". For the weighted directed graph, I intend to use Singapore's Bus Network.</p>	<p>This was the first prompt fed into the LLM.</p> <p>The purpose of this prompt was to introduce ChatGPT, the core research question of the essay as well as provide context for the rest of the conversation.</p>
<p>There are two files you will need, "services.json" and "stops.json". Below are sample data from each file.</p> <pre> services.json { "2": { "name": "Changi Village Ter ⇌ Kampong Bahru Ter", "routes": [[...]] } } stops.json { "10009": [103.81722, 1.2821, "Bt Merah Int", "Bt Merah Ctrl"], ... } </pre>	<p>The purpose of this prompt was to introduce the data files “services.json” and “stops.json”, which was the information used to construct the weighted digraph.</p> <p>Sample data was provided from both files which enabled the LLM to understand how to deserialize the files into Python objects.</p>

Prompt	Explanation
<p>I want you to do this in python, and use the networkx library to create the digraph, and calculate the weights using the haversine distance</p>	<p>The purpose of this prompt was to give the specification that the program should be created in Python.</p> <p>It also specifies that the digraph is to be created with NetworkX, as opposed to igraph or PyGraphviz. It also specifies for the weights of the digraph to be created by applying the haversine formula on the latitude and longitude values.</p>
<p>Let me explain what I need fully:</p> <ol style="list-style-type: none"> 1. Parsetops.json and services.json in python objects using the json library 2. Create a dictionary of stop IDs to coordinates 3. create a function to calculate haversine distance, with the paramters being coord1 and coord2 4. create an empty digraph using networkx 5. add edges to the graph with weights 6. Create the following algorithms: Dijkstra's Algorithm Naive, Dijkstra's Algorithm with a min heap priority queue, Bellman-Ford Algorithm Naive, Bellman-Ford Algorithm with a min heap priority queue. All of these algorithms will solve the SSSP problem. Implement the min heap priority queue using "from queue import PriorityQueue" 7. Measure the execution time of each algorithm in nanoseconds using time.perf_counter_ns() from the time library 	<p>This lengthy prompt helped generate most of the code, I gave instructions in the order of how I wanted them to be implemented within the program.</p> <p>In steps 1-5, I explain to the LLM how the digraph should be constructed using the stops.json and services.json file. These steps involve parsing the data, mapping stop IDs to coordinates, calculating distances, creating the digraph, and adding weighted edges to represent the transportation network.</p> <p>In steps 6 and 7, I ask the LLM to write the implementations for the naive and priority queue-implemented algorithms, as well as the specification to measure the execution time using the time library in Python.</p>
<pre># Measure the execution time of each algorithm in nanoseconds source_stop = "10009" # Example source stop ID target_stop = "10041" # Example target stop ID</pre>	<p>The purpose of this was to correct the LLM, as it initially thought that the problem to be solved was the Single-Pair Shortest Path (SPSP), which finds the shortest path only between a single pair of vertices.</p>

Prompt	Explanation
since i am intending the alogirhtms (sic) to solve the SSSP problem, a target stop wont be needed.	However, the problem being focused on in this essay was the SSSP, therefore, ChatGPT was corrected accordingly.

Below is the Python code used in the experiment to benchmark the algorithms and collect the raw data for the execution times in nanoseconds.

```

1. # filename: ee.py
2. import json
3. import networkx as nx
4. import math
5. from queue import PriorityQueue
6. import time
7.
8. # Constants for Haversine formula
9. AVG_EARTH_RADIUS = 6371 # in kilometers
10.
11. # Step 1: Parse services.json
12. with open('services.json') as services_file:
13.     services_data = json.load(services_file)
14.
15. # Step 2: Parse stops.json
16. with open('stops.json') as stops_file:
17.     stops_data = json.load(stops_file)
18.
19. # Step 3: Create a dictionary of stop IDs to coordinates
20. stop_coordinates = {}
21. for stop_id, stop_info in stops_data.items():
22.     longitude, latitude, _, _ = stop_info
23.     stop_coordinates[stop_id] = (float(latitude), float(longitude))
24.
25. # Function to calculate Haversine distance
26. def haversine_distance(coord1, coord2):
27.     lat1, lon1 = coord1
28.     lat2, lon2 = coord2
29.
30.     dlat = math.radians(lat2 - lat1)

```

```

31.     dlon = math.radians(lon2 - lon1)
32.
33.     a = math.sin(dlat / 2) * math.sin(dlat / 2) + math.cos(
34.         math.radians(lat1)) * math.cos(math.radians(lat2)) * math.sin(
35.         dlon / 2) * math.sin(dlon / 2)
36.     c = 2 * math.atan2(math.sqrt(a), math.sqrt(1 - a))
37.
38.     distance = AVG_EARTH_RADIUS * c
39.     return distance
40.
41. # Step 4: Create an empty directed graph
42. graph = nx.DiGraph()
43.
44. # Step 5: Add edges to the graph with weights
45. stop_ids = list(stop_coordinates.keys())[:len(stop_coordinates)//1]
46.
47. for service_id, service_info in services_data.items():
48.     routes = service_info['routes']
49.     for route in routes:
50.         for i in range(len(route) - 1):
51.             start_stop_id = route[i]
52.             end_stop_id = route[i + 1]
53.
54.             if start_stop_id in stop_ids and end_stop_id in stop_ids:
55.                 start_coordinates = stop_coordinates[start_stop_id]
56.                 end_coordinates = stop_coordinates[end_stop_id]
57.                 distance = haversine_distance(start_coordinates,
58. end_coordinates)
59.                 graph.add_edge(start_stop_id, end_stop_id,
weight=distance)
60.
61. # Dijkstra's algorithm (without priority queue)
62. def dijkstra(graph, start_node):
63.     distances = {node: float('inf') for node in graph.nodes}
64.     distances[start_node] = 0
65.
66.     visited = set()
67.
68.     while len(visited) < len(graph.nodes):
69.         current_node = min((node for node in graph.nodes if node not in
70. visited), key=distances.get)

```

```

71.     visited.add(current_node)
72.
73.     for neighbor, edge_data in graph[current_node].items():
74.         weight = edge_data['weight']
75.         distance = distances[current_node] + weight
76.
77.         if distance < distances[neighbor]:
78.             distances[neighbor] = distance
79.
80.     return distances
81.
82. # Bellman-Ford algorithm (without priority queue)
83. def bellman_ford(graph, start_node):
84.     distances = {node: float('inf') for node in graph.nodes}
85.     distances[start_node] = 0
86.
87.     for _ in range(len(graph.nodes) - 1):
88.         for u, v, edge_data in graph.edges(data=True):
89.             weight = edge_data['weight']
90.             if distances[u] + weight < distances[v]:
91.                 distances[v] = distances[u] + weight
92.
93.     return distances
94.
95. # Dijkstra's algorithm with a priority queue
96. def dijkstra_priority_queue(graph, start_node):
97.     distances = {node: float('inf') for node in graph.nodes}
98.     distances[start_node] = 0
99.
100.    pq = PriorityQueue()
101.    pq.put((0, start_node))
102.
103.    while not pq.empty():
104.        current_distance, current_node = pq.get()
105.
106.        if current_distance > distances[current_node]:
107.            continue
108.
109.        for neighbor, edge_data in graph[current_node].items():
110.            weight = edge_data['weight']
111.            distance = current_distance + weight

```

```

112.
113.         if distance < distances[neighbor]:
114.             distances[neighbor] = distance
115.             pq.put((distance, neighbor))
116.
117.     return distances
118.
119. # Bellman-Ford algorithm with a priority queue
120. def bellman_ford_priority_queue(graph, start_node):
121.     distances = {node: float('inf') for node in graph.nodes}
122.     distances[start_node] = 0
123.
124.     pq = PriorityQueue()
125.     pq.put((0, start_node))
126.
127.     while not pq.empty():
128.         current_distance, current_node = pq.get()
129.
130.         if current_distance > distances[current_node]:
131.             continue
132.
133.         for u, v, edge_data in graph.edges(data=True):
134.             if u != current_node:
135.                 continue
136.
137.             weight = edge_data['weight']
138.             distance = current_distance + weight
139.
140.             if distance < distances[v]:
141.                 distances[v] = distance
142.                 pq.put((distance, v))
143.
144.     return distances
145.
146. print("Number of vertices:", graph.number_of_nodes())
147. print("Number of edges:", graph.number_of_edges())
148.
149. # Measure execution time of Dijkstra's algorithm (without priority queue)
150. start_time = time.perf_counter_ns()
151. shortest_paths_dijkstra = dijkstra(graph, '10009')
152. end_time = time.perf_counter_ns()

```

```

153. execution_time_dijkstra = end_time - start_time # Time in nanoseconds
154.
155. # Measure execution time of Bellman-Ford algorithm (without priority
queue)
156. start_time = time.perf_counter_ns()
157. shortest_paths_bellman_ford = bellman_ford(graph, '10009')
158. end_time = time.perf_counter_ns()
159. execution_time_bellman_ford = end_time - start_time
160.
161. # Measure execution time of Dijkstra's algorithm with priority queue
162. start_time = time.perf_counter_ns()
163. shortest_paths_dijkstra_pq = dijkstra_priority_queue(graph, '10009')
164. end_time = time.perf_counter_ns()
165. execution_time_dijkstra_pq = end_time - start_time
166.
167. # Measure execution time of Bellman-Ford algorithm with priority queue
168. start_time = time.perf_counter_ns()
169. shortest_paths_bellman_ford_pq = bellman_ford_priority_queue(graph,
'10009')
170. end_time = time.perf_counter_ns()
171. execution_time_bellman_ford_pq = end_time - start_time
172.
173. # Print the results
174. print("Dijkstra's algorithm (without priority queue):")
175. print("Execution time:", execution_time_dijkstra, "nanoseconds")
176.
177. print("Bellman-Ford algorithm (without priority queue):")
178. print("Execution time:", execution_time_bellman_ford, "nanoseconds")
179.
180. print("Dijkstra's algorithm with a priority queue:")
181. print("Execution time:", execution_time_dijkstra_pq, "nanoseconds")
182.
183. print("Bellman-Ford algorithm with a priority queue:")
184. print("Execution time:", execution_time_bellman_ford_pq, "nanoseconds")

```

Appendix B: Raw Data — Execution Times For Each Algorithm

Below are four tables containing the raw data for the execution times for all both algorithms naively and priority queue implemented.

Dijkstra's Algorithm (Naive)												
Vertices	Edges	Execution Time (nanoseconds)										Average
		Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	
508	734	11981200	21934400	18719100	18874700	20814800	21334600	19235900	18074100	14620800	17188000	18277760
564	816	22998800	23953800	16544800	17178300	17298700	20640300	22623900	16552200	25723100	25728800	20924270
633	911	29409900	21111900	28999300	21792300	30336900	29728900	29359800	25534200	29384200	22003700	26766110
726	1012	35152900	39225900	37455000	39047500	28401400	32461400	28173600	33840000	30231900	29363500	33335310
847	1186	52392900	46105900	53701300	34253700	45261300	48200400	52807600	50324000	52461800	50390200	48589910
1010	1363	66005400	66154700	64456600	74033600	64706200	72119000	74553200	53278300	63849000	67543300	66669930
1269	1681	95114300	103253100	83637100	114374500	110008000	120537000	109009200	105947500	108108000	101931500	105192020
1694	2241	175072500	181875000	170464400	206013000	138782000	192145500	188640400	173174800	139807800	182432900	174840830
2540	3454	377145300	365253100	483338100	418696100	383313600	383556500	348133600	319151700	434086500	376034000	388870850
5083	7420	1884426500	1658591600	1321206300	1388226400	1547314400	1538802100	1656388100	1574833200	1712586800	1699067300	1598144270

Bellman-Ford Algorithm (Naive)												
Vertices	Edges	Execution Time (nanoseconds)										Average
		Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	
508	734	138645100	144163600	159186900	140928100	136946800	158509300	134072000	144617900	145289700	150328100	145268750

Bellman-Ford Algorithm (Naive)												
Vertices	Edges	Execution Time (nanoseconds)										
		Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	Average
564	816	204098200	203175500	167630900	123132900	147859500	200917900	171318100	116900800	179134100	161555700	167572360
633	911	258275300	217986500	235156500	239027400	235583800	215396400	230625000	211644700	191907700	230321100	226592440
726	1012	318274600	289944200	295647900	314248900	227696100	255524200	175536900	301477900	295709100	230886200	270494600
847	1186	396026400	387980000	306934900	377343200	389627900	357523100	318928100	361835200	356576900	406543600	365931930
1010	1363	508380200	549141700	471250500	553523900	496799700	520952800	537441800	512110500	569653100	359358100	507861230
1269	1681	811700400	694252700	640477700	745964900	802252600	806862000	623195600	652351200	714230200	851501600	734278890
1694	2241	1628369200	1337397300	1370179800	1173142100	1260270300	1441802600	1483032600	1351741800	1309590500	1382545300	1373807150
2540	3454	3374066700	3368888000	3098249400	2799032800	2512895700	3315074500	3142171700	3233649900	3320771300	2775178000	3093997800
5083	7420	14316089000	13657371800	14277746400	14967629500	12846789100	11772173400	13503970400	13775991300	13590726300	12576360500	13528484770

Dijkstra's Algorithm with Priority Queue												
Vertices	Edges	Execution Time (nanoseconds)										
		Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	Average
508	734	1616000	1850700	1514500	1852300	2562500	1629000	1911800	1728600	1849000	1722900	1823730
564	816	2181800	2178300	1649100	1929400	2415600	2120200	1734700	1507700	1665800	1645400	1902800
633	911	2485600	2140200	2187800	1988500	2269300	1820100	2218700	1664800	2225200	2094800	2109500
726	1012	2210600	2258500	2333200	2435600	1865200	2706800	1697000	2435800	2342000	1664700	2194940

Dijkstra's Algorithm with Priority Queue												
Vertices	Edges	Execution Time (nanoseconds)										Average
		Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	
847	1186	283500 0	298220 0	241620 0	361400 0	301410 0	299460 0	286570 0	214600 0	344920 0	248020 0	2879720
1010	1363	343860 0	343120 0	253720 0	364750 0	410240 0	267020 0	366310 0	335420 0	339080 0	275850 0	3299370
1269	1681	321950 0	330530 0	338590 0	435790 0	431480 0	357350 0	337350 0	340310 0	324800 0	432910 0	3651060
1694	2241	679170 0	605580 0	604460 0	590120 0	593800 0	579850 0	571280 0	579570 0	584160 0	639210 0	6027200
2540	3454	730390 0	106532 00	917690 0	707160 0	901610 0	111858 00	729970 0	109760 00	890260 0	892570 0	9051150
5083	7420	205864 00	202696 00	180116 00	147505 00	199088 00	148806 00	145941 00	200480 00	194513 00	152284 00	17772930

Bellman-Ford Algorithm with Priority Queue												
Vertices	Edges	Execution Time (nanoseconds)										Average
		Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	
508	734	892207 00	993532 00	655643 00	867326 00	109478 200	980603 00	980867 00	912629 00	960335 00	929881 00	92678050
564	816	121085 600	128722 100	115316 400	851093 00	133008 300	950124 00	782480 00	830666 00	105363 400	825324 00	102746450
633	911	150521 800	148400 500	152023 300	103652 400	154740 200	123229 300	156813 600	110479 600	109957 600	105607 300	131542560
726	1012	130372 600	178411 600	183438 200	179431 300	139177 600	147948 500	162878 800	180355 000	189677 800	104775 200	159646660
847	1186	274182 800	270720 700	262788 300	260363 300	257145 000	214754 900	268824 600	164592 900	211039 500	293480 700	247789270
1010	1363	377515 000	378410 000	384107 200	328014 100	384777 200	307763 200	229923 900	355940 000	333227 000	364276 900	344395450
1269	1681	408273 800	413298 300	519683 400	564734 100	553576 900	484097 200	374041 300	572099 200	382550 700	390627 600	466298250
1694	2241	113576	113024	103404	930110	102354	105627	961803	105721	902341	927181	1015853050

Bellman-Ford Algorithm with Priority Queue												
Vertices	Edges	Execution Time (nanoseconds)										
		Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	Average
		2200	1700	7300	200	9500	8400	400	4900	400	500	
2540	3454	223602	236622	190106	189384	224222	198364	267176	207517	322457	228363	2287816100
		7300	0200	0900	0400	3900	4200	4600	4100	3400	2000	
5083	7420	104838	917024	858195	931121	992269	891821	875521	102723	103902	831910	9412515240
		94700	8900	1000	6500	7100	8400	9000	26900	75600	4300	

Appendix C: Definitions and Key Terms

Below are some technical terms used in this essay that can be used to support the reader's comprehension:

Time Complexity: Time complexity is defined as the amount of time taken by an algorithm to run, as a function of the length of the input (Team).

Single-Source Shortest Path (SSSP): [The] Single-source shortest path (or SSSP) problem requires finding the shortest path from a source vertex to all other vertices in a weighted graph (Nvidia).

All-Pairs Shortest Path (APSP): [The] All-pairs shortest path (or APSP) problem requires finding the shortest path between all pairs of vertices in a graph (Nvidia).

Greedy Algorithm: A greedy algorithm is a simple, intuitive algorithm that is used in optimization problems. The algorithm makes the optimal choice at each step as it attempts to find the overall optimal way to solve the entire problem (YOON MI KIM).

Naive Algorithm: A naive implementation [of an algorithm] is a programming technique that prioritizes imperfect shortcuts for the sake of speed, simplicity, or lack of knowledge (Perplexity AI. "Prompt: What is a naive implementation of an algorithm?").

Relax/Relaxing/Relaxation: Relaxation is the process of updating the distance of a [vertex] if a shorter path is found (SaturnCloud).

Big-O Notation: Big O notation is a mathematical notation describing a function's limiting behavior when the argument goes towards a certain value or infinity (Ashwani K).

Heap: A Heap is a complete binary tree-based data structure (Jaludi).