



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** XII **Month of publication:** December 2025

DOI: <https://doi.org/10.22214/ijraset.2025.75934>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

A Review and Comparative Analysis of Foundational Shortest-Path Algorithms

Er. Harjasdeep Singh¹, Suneha Gumber², Rishu³

¹Assistant Professor, ^{2,3}Student, Dept. of CSE MIMIT, Malout

Abstract: *The shortest-path problem, a fundamental challenge in graph theory and computer science, seeks to find a path of minimum cumulative weight between vertices in a weighted graph. Its solution is critical to a vast array of applications, including network routing, logistics, robotics, and bioinformatics. This paper provides a comprehensive review and comparative analysis of four foundational algorithms that address this problem: Dijkstra's algorithm, the Bellman-Ford algorithm, the A* search algorithm, and the Floyd-Warshall algorithm. We delve into the historical context, theoretical underpinnings, and operational mechanics of each method. The analysis contrasts these algorithms across several key dimensions: algorithmic paradigm (greedy, dynamic programming, heuristic search), problem scope (single-source vs. all-pairs), handling of edge weights (including negative values), and computational complexity. By juxtaposing their respective strengths, weaknesses, and ideal use cases, this review illustrates that the selection of an optimal shortest-path algorithm is not a matter of absolute superiority but a nuanced decision contingent upon specific problem constraints such as graph structure, edge weight properties, and the required scope of the solution. The paper concludes by contextualizing these classic algorithms as essential building blocks for modern, more complex pathfinding solutions and highlights ongoing research that continues to refine our understanding of this classic computational problem.*

Index Terms: Graph Theory, Shortest Path, Dijkstra, Bellman-Ford, A* Search, Floyd-Warshall, Algorithm Analysis, Computational Complexity.

I. INTRODUCTION

The problem of finding the shortest path in a network is one of the most studied problems in computer science, operations research, and graph theory [1], [2]. In its most general form, the problem involves a weighted graph, $G=(V,E)$, where V is a set of vertices (or nodes) and E is a set of edges connecting pairs of vertices. Each edge $(u, v) \in E$ is assigned a numerical weight, $w(u,v)$, which can represent cost, distance, time, or any other quantifiable measure of traversing the edge [3], [4]. The objective is to find a path between two specified vertices—a source and a destination—such that the sum of the weights of the edges constituting the path is minimized.

The applications of shortest-path algorithms are pervasive and integral to modern technology and infrastructure. They form the backbone of digital mapping services like Google Maps for route planning, network routing protocols such as OSPF (Open Shortest Path First) for directing internet traffic, resource allocation in logistics and supply chain management, and pathfinding for autonomous agents in robotics and video games [4]–[6]. The versatility of the graph model allows vertices to represent not just physical locations but also abstract states, with edges representing transitions, making shortest-path algorithms a powerful tool for solving a wide range of optimization problems [3].

The literature presents a rich tapestry of algorithms designed to solve this problem, each tailored to different constraints and problem variations [1], [2]. These variations are broadly classified into two categories: the Single-Source Shortest Path (SSSP) problem, which aims to find the shortest paths from a single source vertex to all other vertices in the graph, and the All-Pairs Shortest Path (APSP) problem, which seeks to find the shortest path between every pair of vertices [1], [3].

This review focuses on four seminal algorithms that have become canonical in the study of this problem:

- 1) Dijkstra's Algorithm: A greedy algorithm that efficiently solves the SSSP problem for graphs with non-negative edge weights.
- 2) The Bellman-Ford Algorithm: A more robust SSSP algorithm based on dynamic programming, capable of handling graphs with negative edge weights and detecting negative-weight cycles.
- 3) The A* Search Algorithm: An informed or heuristic search algorithm that extends Dijkstra's approach to find the shortest path between a single pair of vertices more efficiently by using problem-specific knowledge to guide its search.
- 4) The Floyd-Warshall Algorithm: A dynamic programming algorithm that solves the APSP problem, elegantly handling both positive and negative edge weights (in the absence of negative-weight cycles).

The continued relevance of these distinct algorithms underscores a critical reality: there is no single "best" algorithm for all shortest-path scenarios. The optimal choice is contingent on the specific characteristics of the problem at hand, including the graph's size and density, the nature of its edge weights, and whether the solution requires paths from a single source or between all pairs of nodes. This paper aims to provide a rigorous comparative analysis of these four foundational algorithms. It will examine their historical origins, formalize their operational principles, analyze their computational complexity, and contrast their performance under various conditions. Through this detailed exploration, we seek to furnish a clear framework for understanding the trade-offs involved and for selecting the most appropriate algorithm for a given application.

II. LITERATURE REVIEW AND HISTORICAL CONTEXT

The development of shortest-path algorithms is deeply intertwined with the history of computer science, reflecting a fascinating evolution of computational paradigms. The four algorithms under review did not emerge in a simple linear progression of improvement but rather represent distinct philosophical approaches to problem-solving: greedy optimization, dynamic programming, and heuristic search.

A. Dijkstra's Algorithm: The Greedy Approach

In 1959, Dutch computer scientist Edsger W. Dijkstra published "A Note on Two Problems in Connexion with Graphs," which introduced his now-famous algorithm for the SSSP problem [7]. The algorithm was conceived in just 20 minutes in 1956 as a demonstration of the capabilities of the ARMAC computer [1], [2]. Its design embodies the greedy paradigm: at each step, it makes the locally optimal choice by selecting the unvisited vertex closest to the source, with the assumption that this choice will lead to a globally optimal solution [3], [4]. This elegant and efficient approach established a benchmark for solving the SSSP problem on graphs with non-negative weights, and its core "relaxation" process remains a fundamental concept in the field [5], [8].

B. Bellman-Ford: The Power of Dynamic Programming

The ability to handle negative edge weights, a crucial requirement in fields like economic modeling and certain network protocols, was addressed through the principle of dynamic programming. The algorithm known today as Bellman-Ford has a dual origin. Richard Bellman, in his 1958 paper "On a Routing Problem," developed a functional equation based on dynamic programming to solve routing issues [9]. His method iteratively calculates the shortest path of at most k edges using the known shortest paths of at most $k-1$ edges [3], [4]. Independently, Lester R. Ford Jr., in his 1956 RAND Corporation report on network flow theory, developed a similar iterative method [10]. This methodical, bottom-up approach is more computationally intensive than Dijkstra's greedy strategy but offers greater versatility, including the critical ability to detect negative-weight cycles—a condition under which the shortest path is often undefined [5], [11].

C. A* Search: The Heuristic Infusion

While Dijkstra's and Bellman-Ford's algorithms explore paths radiating from the source, the A* search algorithm introduced a sense of direction. Developed by Peter E. Hart, Nils J. Nilsson, and Bertram Raphael at the Stanford Research Institute, it often explores a much smaller portion of the graph than its uninformed counterparts [5], [6]. This fusion of a formal graph search with domain-specific knowledge represented a significant step forward for goal-directed pathfinding.

D. Floyd-Warshall: The All-Pairs Perspective

The focus shifted from single source to all pairs shortest paths with the algorithm now widely attributed to Robert W. Floyd and Stephen Warshall. In 1962, Floyd published "Algorithm 97: Shortest Path," which presented a remarkably concise dynamic programming solution for the APSP problem [13]. Its recurrence relation is based on iterating through all possible intermediate vertices for each pair of start and end nodes [3], [4]. However, the algorithm's intellectual lineage is more complex. Bernard Roy had published a similar algorithm in 1959 in the French journal *Comptes Rendus de l'Académie des Sciences* [14]. Furthermore, Stephen Warshall, also in 1962, published an algorithm for finding the transitive closure of a graph ("A theorem on Boolean matrices"), which is structurally identical to Floyd's algorithm when applied to unweighted graphs [4], [15]. This history highlights a common theme in science: the simultaneous and independent discovery of fundamental ideas. The Floyd-Warshall algorithm, as it is now known, provides a powerful, if computationally demanding, method for comprehensive path analysis in dense graphs.

These historical threads reveal that the field did not simply iterate towards a single "best" algorithm. Instead, it branched out, exploring distinct computational philosophies to create a versatile toolkit, with each tool uniquely suited to a specific set of constraints and objectives.

III. METHODOLOGICAL FRAMEWORK FOR ANALYSIS

To conduct a rigorous and systematic comparison of the four selected algorithms, this paper establishes a formal framework grounded in graph theory and computational complexity analysis. This framework defines the problem domain and outlines the criteria against which each algorithm will be evaluated.

A. Graph Theory Preliminaries

Let $G=(V,E)$ be a directed graph, where V is a finite set of vertices (nodes), $|V|=n$, and $E \subseteq V \times V$ is a set of edges, $|E|=m$. Each edge $(u,v) \in E$ is associated with a real-valued weight $w(u,v)$. A path p from a vertex v_0 to a vertex v_k is a sequence of vertices $\langle v_0, v_1, \dots, v_k \rangle$ such that $(v_{i-1}, v_i) \in E$ for all $1 \leq i \leq k$. The weight of a path is the sum of the weights of its constituent edges: $w(p) = \sum_{i=1}^k w(v_{i-1}, v_i)$.

Institute and published in their 1968 paper, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," A* enhanced Dijkstra's greedy framework by incorporating a heuristic function. The shortest-path weight $\delta(u,v)$ from vertex u to vertex v is defined as:

$\delta(u,v) = \min \{ w(p) \mid \text{if a path exists from } u \text{ to } v \}$, which estimates the cost from a given vertex n to the destination [4], [6]. By prioritizing vertices that appear to be on the most promising path to the goal (i.e., those with the lowest $f(n) = g(n) + h(n)$),

A* intelligently guides its search, $\delta(u,v) = \infty$ otherwise. A shortest path from u to v is any path p with weight $w(p) = \delta(u,v)$.

B. Problem Variants

This review addresses two primary variants of the shortest-path problem [1], [3]:

- 1) Single-Source Shortest Path (SSSP): Given a graph G and a source vertex $s \in V$, find the shortest-path weight $\delta(s, v)$ for all $v \in V$. Dijkstra's, Bellman-Ford, and A* (for a single destination) address this problem.
- 2) All-Pairs Shortest Path (APSP): Find the shortest-path weight $\delta(u,v)$ for every pair of vertices $(u,v) \in V \times V$. The Floyd-Warshall algorithm is designed for this problem.

C. Criteria for Comparative Analysis

Each algorithm will be evaluated based on the following criteria:

- 1) Algorithmic Paradigm: The underlying design strategy (e.g., greedy, dynamic programming, heuristic search).
- 2) Time Complexity: The asymptotic upper bound on the algorithm's running time, expressed in Big-O notation as a function of n and m .
- 3) Space Complexity: The asymptotic upper bound on the memory required by the algorithm.
- 4) Edge Weight Constraints: The algorithm's ability to function correctly with non-negative, negative, or zero-weight edges.
- 5) Negative Cycle Detection: The capability of the algorithm to detect the presence of a negative-weight cycle, a cycle whose edges sum to a negative value. The shortest path is undefined in graphs containing a negative cycle reachable from the source, as traversing the cycle indefinitely would lead to an infinitely small path weight [3], [11].

D. The Role of Data Structures and Graph Density

The theoretical complexity of an algorithm is an abstraction that can be significantly influenced by implementation choices, particularly the data structure used to represent the graph and the graph's density. A graph is considered **dense** if m is close to its maximum possible value, $O(n^2)$, and **sparse** if m is much smaller, often close to $O(n)$ [4].

- 1) Adjacency Matrix: An $n \times n$ matrix where the entry at (i,j) stores the weight of edge (i,j) . It requires $O(n^2)$ space and is efficient for dense graphs, as checking for an edge is an $O(1)$ operation.
- 2) Adjacency List: An array of lists, where each index i corresponds to a vertex and stores a list of its neighbors. It requires $O(n+m)$ space and is more efficient for sparse graphs.

The practical performance of algorithms like Dijkstra's, whose complexity is often given as $O(m \log n)$, is highly dependent on this choice. For a dense graph, this becomes $O(n^2 \log n)$, which may be less efficient than a simpler $O(n^2)$ implementation.

Therefore, our analysis will consider performance on both sparse and dense graphs, acknowledging the critical interplay between algorithm, data structure, and graph topology.

IV. IN-DEPTH ANALYSIS OF SHORTEST-PATH ALGORITHMS

This section provides a detailed technical examination of each of the four algorithms, covering their operational principles, formal pseudocode, and a rigorous analysis of their complexity and limitations.

A. Dijkstra's Algorithm

1) *Principle of Operation:* Dijkstra's algorithm solves the SSSP problem for a weighted directed graph with non-negative edge weights [3], [4]. It operates on a greedy principle, iteratively building a set of vertices, S , for which the shortest path from the source s is known. Initially, S is empty, and the distance to all vertices is set to infinity, except for the source, whose distance is 0 [5].

The algorithm maintains a priority queue, Q , of vertices that have been discovered but are not yet in S . The priority of each vertex u in Q is its current known shortest distance from the source, $d[u]$. In each iteration, the algorithm extracts the vertex u with the minimum distance from Q , adds it to S , and performs a "relaxation" step for each of its neighbors, v [3], [4]. Relaxation of an edge (u, v) consists of checking whether the path to v through u is shorter than the currently known path to v . If $d[u] + w(u, v) < d[v]$, the distance $d[v]$ is updated, and the predecessor of v is set to u [5]. This process continues until the priority queue is empty, at which point the shortest paths to all reachable vertices have been determined. The greedy choice is justified because, with non-negative weights, the first time a path to a vertex is finalized (by moving it from Q to S), it is guaranteed to be the shortest one [3].

2) *Pseudocode:* Algorithm 1 presents the formal pseudocode for Dijkstra's algorithm using a min-priority queue.

Algorithm 1 Dijkstra's Algorithm Dijkstra G, w, s Graph $G = (V, E)$, weight function w , source vertex s
 Distances $d[v]$ and predecessors $\pi[v]$ for all $v \in V$ Initialize-Single-Source(G, s) $S \leftarrow \emptyset$ $Q \leftarrow V$ Min-priority queue of vertices
 $Q \neq \emptyset \leftarrow \text{Extract-Min}(Q)$ $S \leftarrow S \cup \{u\}$ each vertex $v \in \text{Adj}[u]$ Relax(u, v, w) Initialize-Single-Source G, s each vertex $v \in V$ $d[v] \leftarrow \infty$
 $\pi[v] \leftarrow \text{NIL}$ $d[s] \leftarrow 0$ Relax u, v, w $d[v] > d[u] + w(u, v)$
 $d[v] \leftarrow d[u] + w(u, v)$ $\pi[v] \leftarrow u$ Decrease-Key($Q, v, d[v]$)

3) *Complexity Analysis:* The time complexity of Dijkstra's algorithm is critically dependent on the implementation of the min-priority queue Q [4].

- **Simple Array:** If Q is an unsorted array, 'Extract-Min' takes $O(n)$ time and 'Decrease-Key' takes $O(1)$ time. Since 'Extract-Min' is called n times and 'Relax' (which calls 'Decrease-Key') is called at most m times, the total complexity is $O(n^2 + m) = O(n^2)$. This is efficient for dense graphs where $m \approx n^2$.
- **Binary Heap:** With a binary heap, both 'Extract-Min' and 'Decrease-Key' take $O(\log n)$ time. The total complexity becomes $O(n \log n + m \log n) = O((n + m) \log n)$. For connected graphs ($m \geq n - 1$), this is commonly simplified to $O(m \log n)$ [3], [4]. This is the standard implementation and is highly efficient for sparse graphs.
- **Fibonacci Heap:** A more advanced data structure, the Fibonacci heap, provides an amortized time complexity of $O(\log n)$ for 'Extract-Min' and $O(1)$ for 'Decrease-Key'. This yields the best-known theoretical worst-case complexity for Dijkstra's algorithm: $O(m + n \log n)$ [3], [4].

The space complexity is $O(n + m)$ to store the graph (using an adjacency list) and the distance/predecessor arrays.

4) *Limitations:* The primary limitation of Dijkstra's algorithm is its inability to handle negative edge weights. The greedy strategy relies on the assumption that once a vertex u is extracted from the priority queue, its shortest path $d[u]$ is finalized. A negative edge could violate this: a path through a later-explored vertex could lead back to u or one of its neighbors via a negative edge, creating a shorter path than the one already found. This invalidates the core assumption of the algorithm [3], [11].

B. The Bellman-Ford Algorithm

1) *Principle of Operation:* The Bellman-Ford algorithm solves the SSSP problem in graphs that may contain negative-weight edges [4], [5]. It employs a dynamic programming approach based on the principle of relaxation. The algorithm iteratively relaxes all edges in the graph. After the first pass of relaxing all m edges, it guarantees to find all shortest paths of length 1 (i.e., one edge). After the second pass, it finds all shortest paths of length at most 2, and so on [3].

Since any simple shortest path can have at most $n-1$ edges, the algorithm repeats this relaxation process $n-1$ times. After $n-1$ iterations, it is guaranteed to have found the shortest path for all reachable vertices, provided there are no negative-weight cycles [3], [4]. A key feature of Bellman-Ford is its ability to detect a loop that runs $n-1$ times and an inner loop that iterates through all edges. This results in a time complexity of $O((n-1) \cdot m) = O(nm)$. The second part, for negative cycle detection, iterates through all edges once, taking $O(m)$ time. Thus, the total time complexity is $O(nm)$ [3], [4]. This holds for both sparse and dense graphs, becoming $O(n^3)$ in the dense case.

Space Complexity: The algorithm requires storage for the distance and predecessor arrays, leading to a space complexity of $O(n)$ [4]. If an adjacency list is used for the graph, the total space is $O(n+m)$.

C. The A* Search Algorithm

1) **Principle of Operation:** A* is an informed search algorithm that aims to find the shortest path from a single source to a single destination [6], [12]. It can be seen as an extension of Dijkstra's algorithm. Like Dijkstra's, it uses a priority queue (often called the "open set") to explore the graph. However, the priority of a vertex n is not just its distance from the source, $g(n)$, but an evaluation function $f(n) = g(n) + h(n)$ [4], [6].

Here, $g(n)$ is the cost of the currently known shortest path from the source to n , and $h(n)$ is a heuristic function that estimates the cost of the shortest path from n to the destination. The heuristic guides the search towards the destination, allowing A* to avoid exploring paths that are moving away from the goal [6], [12]. The algorithm also maintains a "closed set" of vertices that have already been fully explored.

For A* to be optimal (i.e., guaranteed to find the shortest path), the heuristic function $h(n)$ must be admissible, meaning it never overestimates the true cost to reach the goal. That is, $h(n) \leq \delta(n, \text{goal})$ for all vertices n [6], [12]. A common example of an admissible heuristic is the straight-line Euclidean distance in a geometric pathfinding problem.

2) **Pseudocode:** Algorithm 3 details the A* search procedure.

negative-weight cycles. If, after $n-1$ iterations, a further (i.e., n -th) iteration still results in a relaxation (a shortening of a path), it implies that a negative-weight cycle exists in the graph and is reachable from the source [3], [5].

3) **Pseudocode:** Algorithm 2 describes the Bellman-Ford procedure, including the step for negative cycle detection.

Algorithm 2 The Bellman-Ford Algorithm

Bellman-Ford G, w, s Graph $G = (V, E)$, weight function w , source vertex s 'true' if no negative cycle, 'false' otherwise. Distances $d[v]$ and predecessors $\pi[v]$. Initialize- Single-Source(G, s)
 $i \leftarrow 1$ to $|V| - 1$ each edge $(u, v) \in E$ Relax(u, v, w) Standard relaxation each edge $(u, v) \in E$ Check for negative cycles $d[v] > d[u] + w(u, v)$ 'false' Negative cycle detected 'true'

4) **Complexity Analysis:** The structure of the Bellman-Ford algorithm is straightforward, leading to a simple complexity analysis.

Time Complexity: The algorithm consists of two main parts. The first part has a nested loop structure: an outer

Algorithm 3 The A* Search Algorithm

A-Star start, goal, h Start vertex start, goal vertex goal, heuristic h The shortest path from start to goal openSet $\leftarrow \{start\}$ Priority queue with f -scores
cameFrom \leftarrow an empty map $gScore[v] \leftarrow \infty$ for all $v \in V$ $gScore[start] \leftarrow 0$ $fScore[v] \leftarrow \infty$ for empty $current \leftarrow$ vertex in openSet with the lowest
 $fScore_{current} = goal$ reconstruct path (cameFrom, current) Remove $current$ from openSet each neighbor
of $current$ tentative $gScore \leftarrow gScore[current] + w(current, neighbor)$ tentative $gScore <$
 $gScore[neighbor]$ cameFrom[neighbor] $\leftarrow current$ $gScore[neighbor] \leftarrow tentative gScore$
 $fScore[neighbor] \leftarrow gScore[neighbor] + h(neighbor)$ neighbor not in openSet Add neighbor to openSet failure No path found

5) **Complexity Analysis:** The performance of A* is highly sensitive to the quality of its heuristic.

- **Time Complexity:** In the worst case, with a non-informative heuristic (e.g., $h(n) = 0$ for all n), A* becomes equivalent to Dijkstra's algorithm, and its time complexity is $O(m \log n)$ with a binary heap [6]. As the heuristic becomes more accurate, the number of explored nodes decreases, and the performance improves significantly. With a perfect heuristic ($h(n) = \delta(n, \text{goal})$), A* only explores nodes on the optimal path, leading to a complexity close to $O(n+m)$ in practice.

- Space Complexity: A^* must store all generated nodes in the open and closed sets. In the worst case, this can be the entire graph, leading to a space complexity of $O(n+m)$. For problems with large state spaces, such as in game AI or robotics, this exponential space requirement can be a significant drawback [5], [6].

D. The Floyd-Warshall Algorithm

- 1) *Principle of Operation:* The Floyd-Warshall algorithm solves the APSP problem using a dynamic programming approach [4], [13]. It works by iteratively considering each vertex in the graph as a potential intermediate vertex in the shortest paths between all other pairs of vertices.

Let $d^{(k)}$ be the weight of the shortest path from vertex i to vertex j using only intermediate vertices from the set $\{1, 2, \dots, k\}$. The algorithm computes a sequence of $n \times n$ matrices, D^0, D^1, \dots, D^n , where $D^k[i, j] = d^{(k)}$. The

- 2) *Negative Cycle Detection:* The Floyd-Warshall algorithm can be used to detect negative-weight cycles. After the algorithm completes, if any diagonal element d_{ii} is negative, it indicates that vertex i is part of or can reach a negative-weight cycle [4], [11]. This is because the path from a vertex to itself should be 0, and a negative value implies that the path can be made shorter by traversing a cycle.

V. COMPARATIVE ANALYSIS AND APPLICATIONS

The choice of a shortest-path algorithm is a critical design decision that hinges on the specific constraints of the problem. This section provides a direct comparison of the four algorithms based on the framework established earlier and discusses their principal real-world applications. A summary of this comparison is presented in Table I.

A. SSSP vs. APSP

The most fundamental distinction is between algorithms that solve the SSSP problem and those that solve the APSP problem. For APSP, one could simply run an SSSP algorithm, like Dijkstra's, once for each of the n vertices.

- 1) For sparse graphs ($m \approx n$): Running Dijkstra's n times yields a total time complexity of $n \times O(m \log n) = O(nm \log n)$. This is generally more efficient than Floyd-Warshall's $O(n^3)$ complexity. If negative weights are present, Johnson's algorithm, which uses Bellman-Ford

base case, D^0 , is the initial adjacency matrix of the graph.

Warshall's $O(n^3)$ complexity. If negative weights are

The core of the algorithm is the following recurrence relation

present, Johnson's algorithm, which uses Bellman-Ford

[4]:

(k)

(k-1)

(k-1)

(k-1)

once and Dijkstra's n times, achieves $O(nm + n^2 \log n)$

$d_{ij} := \min(d_{ij}^{(k-1)}, d_{ik}^{(k-1)} + d_{kj}^{(k-1)})$

+ $d_{kj}^{(k-1)}$)

$d_{ij}^{(k)} = \min(d_{ij}^{(k-1)}, d_{ik}^{(k-1)} + d_{kj}^{(k-1)})$

This formula states that the shortest path from i to j using intermediate vertices up to k is either the shortest path using only vertices up to $k-1$, or it is the path that goes from i to k and then from k to j (both using intermediate vertices up to $k-1$). By iterating k from 1 to n , the final matrix $D^{(n)}$ contains the shortest-path weights between all pairs of vertices [13].

- 2) *Pseudocode:* The implementation of the Floyd-Warshall algorithm is notably compact, as shown in Algorithm 4, and is also superior to Floyd-Warshall on sparse graphs [3].

- For dense graphs ($m \approx n^2$): Running Dijkstra's n times results in $O(n \cdot n^2) = O(n^3)$ with a simple array priority queue, or $O(n \cdot n^2 \log n)$ with a binary heap. In this scenario, Floyd-Warshall's $O(n^3)$ complexity is competitive and often preferred due to its simpler implementation and better memory locality [4], [13].

B. Handling of Negative Weights

Algorithm 4 The Floyd-Warshall Algorithm

Floyd-Warshall W is an $n \times n$ weight matrix of a graph with n

vertices $n \times n$ matrix of shortest-path weights $D^{(0)}$ $W_k \leftarrow 1$

for $i \leftarrow 1$ to n for $j \leftarrow 1$ to n $d^{(k)} \leftarrow \min(d^{(k-1)}, d^{(k-1)} +$

The presence of negative edge weights is a major differentiating factor.

• Dijkstra's algorithm fails because its greedy assumption

$d^{(k-1)}(n)$ ij ij ik

is violated. It is fundamentally unsuited for this problem

kj

3) Complexity Analysis:

- Time Complexity: The algorithm's structure consists of three nested loops, each iterating from 1 to n . The inner operation is constant time. Therefore, the time complexity is consistently $\Theta(n^3)$ regardless of the graph's structure or density [4], [13].
- Space Complexity: The algorithm requires an $n \times n$ matrix to store the distances. With a space-optimized implementation that updates the matrix in place, the space complexity is $\Theta(n^2)$ [4]. An additional $\Theta(n^2)$ matrix can be used to reconstruct the paths themselves. domain.
- Bellman-Ford is the classic SSSP solution for graphs with negative weights. Its iterative nature ensures that it correctly propagates the effects of negative edges throughout the graph [4], [5]. Its ability to detect negative cycles is crucial in applications where such cycles render the problem ill-defined, such as in economic arbitrage analysis [3].
- Floyd-Warshall also correctly handles negative weights for the APSP problem and provides a simple mechanism for detecting the presence of any negative cycle in the graph by inspecting the diagonal of the final distance matrix [4].

TABLE I

COMPARATIVE SUMMARY OF FOUNDATIONAL SHORTEST-PATH ALGORITHMS

Criterion	Dijkstra's Algorithm	Bellman-Ford Algorithm	A* Search Algorithm	Floyd-Warshall Algorithm
Problem Type	SSSP	SSSP	Single-Pair (typically)	APSP
Algorithmic Paradigm	Greedy	Dynamic Programming	Heuristic Search (Informed)	Dynamic Programming
Edge Weights	Non-negative only	Positive or negative	Non-negative (typically)	Positive or negative
Negative Cycle	Cannot handle	Detects reachable cycles	Not designed for this	Detects all cycles
Time Complexity	$O(m \log n)$ (Binary Heap)	$O(nm)$	Varies (heuristic-dependent)	$\Theta(n^3)$
Space Complexity	$O(n+m)$	$O(n+m)$	$O(n+m)$ (can be exponential)	$\Theta(n^2)$
Primary Application	Network routing (OSPF), GPS	Distance-vector protocols (RIP)	Game AI, Robotics, Pathfinding	All-pairs analysis, Transitive closure

C. Heuristic vs. Uninformed Search

A* stands apart from Dijkstra's by incorporating domain knowledge. While Dijkstra's algorithm explores outwards uniformly from the source, A* focuses its search towards the goal [6], [12]. This makes it exceptionally powerful for single-pair pathfinding in large state spaces, such as maps or game levels. The trade-off is the cost of designing and computing an effective heuristic. A poor heuristic can make A* perform no better, or even worse, than Dijkstra's, while a highly accurate heuristic can lead to near-optimal search performance.

D. Real-World Applications

The theoretical differences between the algorithms directly map to their suitability for different real-world applications.

- 1) Dijkstra's Algorithm: Its efficiency and simplicity make it the standard for SSSP problems with non-negative costs. It is famously used in the OSPF (Open Shortest Path First) internet routing protocol, where link costs are positive metrics [3]. It is also fundamental to GPS navigation and mapping services for calculating the fastest or shortest route between two points [4], [5].

- 2) Bellman-Ford Algorithm: Its ability to handle negative weights makes it suitable for distance-vector routing protocols like the Routing Information Protocol (RIP). While RIP itself does not use negative weights, the Bellman-Ford structure is robust to the dynamic updates and potential routing loops that can occur in such networks [3], [4]. It is also used in analyzing networks for arbitrage opportunities where negative cycles represent profitable loops.
- 3) A* Search Algorithm: A* is the dominant algorithm in any domain requiring goal-directed pathfinding. This includes video game development for NPC (non-player character) pathfinding, robotics for navigation and motion planning, and logistics applications for finding optimal routes for autonomous vehicles like drones or warehouse robots [4], [6].
- 4) Floyd-Warshall Algorithm: Its $O(n^3)$ complexity limits its use to smaller or denser graphs. It is valuable for computing the **transitive closure** of a graph (determining reachability between all pairs of nodes) [11]. Other applications include computing all-pairs distances in social networks to find the "degrees of separation" between all members, or in bioinformatics for sequence alignment [3], [4].

VI. CONCLUSION AND FUTURE DIRECTIONS

This review has systematically analyzed four of the most foundational algorithms for solving the shortest-path problem: Dijkstra's, Bellman-Ford, A*, and Floyd-Warshall. The analysis demonstrates that no single algorithm is universally superior. Instead, each represents a unique solution tailored to a specific set of problem constraints, embodying a distinct trade-off between speed, generality, and problem scope.

Dijkstra's algorithm offers unparalleled efficiency for the SSSP problem on graphs with non-negative weights, making it the default choice for a wide range of applications. The Bellman-Ford algorithm, while slower, provides the necessary robustness to handle negative edge weights and the critical capability to detect negative-weight cycles. The A* search algorithm showcases the power of heuristic guidance, dramatically accelerating single-pair pathfinding in applications where domain knowledge can be leveraged to direct the search. Finally, the Floyd-Warshall algorithm provides a comprehensive, albeit computationally intensive, solution for the APSP problem, proving invaluable for dense graphs and problems requiring a complete distance matrix.

The true legacy of these algorithms extends beyond their direct applications. They serve as fundamental building blocks and conceptual baselines for the ongoing development of more advanced pathfinding techniques. This is evident in several key areas of modern research:

- **Hybrid Algorithms:** More sophisticated algorithms often combine the strengths of these foundational methods. A prime example is Johnson's algorithm, which cleverly uses Bellman-Ford as a preprocessing step to re-weight a graph, eliminating negative edges. This allows the more efficient Dijkstra's algorithm to be run subsequently for each vertex, making it one of the fastest solutions for APSP on sparse graphs with negative weights [3].
- **Parallelization:** As datasets grow, parallel computation becomes essential. The inherent structure of these algorithms dictates their suitability for parallelization. The nested loops of Floyd-Warshall, for instance, are highly amenable to parallel execution through techniques like 2D block mapping [1], [2]. In contrast, the inherently sequential nature of Dijkstra's algorithm—where each step depends on the minimum-priority vertex from the previous step—presents a greater challenge for parallel implementation [2], [5].
- **New Theoretical Frontiers:** Even after decades of study, the theoretical limits of shortest-path computation are still being explored. Recent work by Duan et al. has introduced a novel algorithm that breaks the long-standing "sorting barrier" associated with Dijkstra's algorithm, achieving a slightly faster time complexity for the SSSP problem on directed graphs [5], [8]. This breakthrough, which itself combines ideas from both Dijkstra's and Bellman-Ford's algorithms, demonstrates that this field remains a vibrant area of active research.

In conclusion, the algorithms of Dijkstra, Bellman, Ford, Hart, Nilsson, Raphael, Floyd, and Warshall are not merely historical artifacts. They are enduring pillars of computer science that continue to solve countless real-world problems today. More importantly, they provide the core principles—greedy selection, iterative relaxation, dynamic programming, and heuristic guidance—that inspire and enable the next generation of algorithms designed to navigate the ever-more-complex networks of our digital world.

REFERENCES

- [1] A. Madkour, W. G. Aref, F. U. Rehman, M. A. Rahman, and S. Basalamah, "A survey of shortest-path algorithms," arXiv preprint arXiv:1705.02044, 2017.
- [2] L. dos Santos and L. M. de Assis, "Parallelization of shortest path algorithms: a comparative analysis," Pesquisa Operacional, vol. 43, 2023.
- [3] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, Introduction to Algorithms, 3rd ed. MIT Press, 2009.



- [4] R. Sedgewick and K. Wayne, Algorithms, 4th ed. Addison-Wesley Professional, 2011.
- [5] B. V. Cherkassky, A. V. Goldberg, and T. Radzik, "Shortest path algorithms: Theory and experimental evaluation," Mathematical programming, vol. 73, no. 2, pp. 129–174, 1996.
- [6] N. J. Nilsson, Principles of artificial intelligence. Tioga Publishing Company, 1980.
- [7] E.W.Dijkstra, "A note on two problems in connexion with graphs," Numerische Mathematik, vol. 1, no. 1, pp. 269–271, 1959.
- [8] R. Duan, G. He, M. Lyu, Y. Wu, and R. Xie, "Breaking the sorting barrier for shortest paths on directed graphs," in Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). SIAM, 2023, pp. 136–152.
- [9] R. Bellman, "On a routing problem," Quarterly of Applied Mathematics, vol. 16, no. 1, pp. 87–90, 1958.
- [10] L. R. Ford, Jr., "Network flow theory," RAND Corporation, Tech. Rep. P-923, 1956.
- [11] S. Hougardy, "The floyd-warshall algorithm on graphs with negative cycles," Information Processing Letters, vol. 110, no. 8-9, pp. 279–281, 2010.
- [12] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," IEEE Transactions on Systems Science and Cybernetics, vol. 4, no. 2, pp. 100–107, 1968.
- [13] R. W. Floyd, "Algorithm 97: Shortest path," Communications of the ACM, vol. 5, no. 6, p. 345, 1962.
- [14] B. Roy, "Transitivite et connexite," Comptes Rendus de l'Academie des Sciences, vol. 249, pp. 216–218, 1959.
- [15] S. Warshall, "A theorem on boolean matrices," Journal of the ACM, vol. 9, no. 1, pp. 11–12, 1962.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)