


Shortest Paths



shortest path from Princeton CS department to Einstein's house

Robert Sedgewick and Kevin Wayne · Copyright © 2005 · <http://www.Princeton.EDU/~cos226>

Brief History

Shimbel (1955). Information networks.

Ford (1956). RAND, economics of transportation.

Leyzorek, Gray, Johnson, Ladew, Meaker, Petry, Seitz (1957).
Combat Development Dept. of the Army Electronic Proving Ground.

Dantzig (1958). Simplex method for linear programming.

Bellman (1958). Dynamic programming.


Moore (1959). Routing long-distance telephone calls for Bell Labs.

Dijkstra (1959). Simpler and faster version of Ford's algorithm.

Shortest Path Problem

Shortest path problem. Given a weighted digraph, find the shortest directed path from s to t .

cost of path = sum of edge costs in path



Path: $s \rightarrow 6 \rightarrow 3 \rightarrow 5 \rightarrow t$
Cost: $14 + 18 + 2 + 16 = 50$

Versions.

- Point-to-point, single source, all pairs.
- Nonnegative edge weights, arbitrary weights, Euclidean weights.

Applications

More applications.

- Robot navigation.
- Texture mapping.
- Typesetting in TeX.
- Urban traffic planning.
- Optimal pipelining of VLSI chip.
- Telemarketer operator scheduling.
- Subroutine in higher level algorithms.
- Routing of telecommunications messages.
- Approximating piecewise linear functions.
- Network routing protocols (OSPF, BGP, RIP).
- Exploiting **arbitrage** opportunities in currency exchange.
- Optimal truck routing through given traffic congestion pattern.

Reference: *Network Flows: Theory, Algorithms, and Applications*, R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, Prentice Hall, 1993.


Dijkstra's Algorithm

Single Source Shortest Path

Assumptions.

- Digraph G .
- Single source s .
- Edge weights $c(v, w)$ are nonnegative.

Goal. Find shortest path from s to every other vertex.



Edge Relaxation


Valid weights. For all vertices v , $\pi(v)$ is length of some path from s to v .

Edge relaxation.

- Consider edge $e = v \rightarrow w$.
- If current path from s to v plus edge $v \rightarrow w$ is shorter than current path to w , then update current path to w .

```

if ( $\pi(w) > \pi(v) + c(v, w)$ ) {
     $\pi(w) = \pi(v) + c(v, w)$ 
     $\text{pred}(w) = v$ 
}
    
```



Dijkstra's Algorithm


Dijkstra's algorithm. Maintain a valid set of weights $\pi(v)$ and a set of **explored vertices** S for which $\pi(v)$ is the length shortest s - v path.

- Initialize: $S = \{s\}$, $\pi(s) = 0$.
- Repeatedly choose unexplored node w which minimizes:

$$\pi(w) = \min_{(v,w): v \in S} \pi(v) + c(v,w)$$

- set $\text{pred}(w) = v$
- add w to S , and set $\pi(w) = \pi(v) + c(v, w)$

← shortest path to some v in explored part, followed by a single edge (v, w)



Dijkstra's Algorithm


Dijkstra's algorithm. Maintain a valid set of weights $\pi(v)$ and a set of **explored vertices** S for which $\pi(v)$ is the length shortest s - v path.

- Initialize: $S = \{s\}$, $\pi(s) = 0$.
- Repeatedly choose unexplored node w which minimizes:

$$\pi(w) = \min_{(v,w): v \in S} \pi(v) + c(v,w)$$

- set $\text{pred}(w) = v$

- add w to S , and set $\pi(w) = \pi(v) + c(v, w)$



shortest path to some v in explored part, followed by a single edge (v, w)


9

Dijkstra's Algorithm: Proof of Correctness

Invariant. For each vertex v in S , $\pi(v)$ is the length of shortest s - v path.



Pf. (by induction on $|S|$)

- Let w be next vertex added to S .
- $\pi(w) = \pi(v) + c(v, w)$ is length of some s - v path.
- Consider any s - v path P , and let x be first node on path outside S .
- P is already too long as soon as it reaches x by greedy choice.





10

Dijkstra's Algorithm



11

Shortest Path Tree



12

Dijkstra's Algorithm: Implementation

Critical step. Choose unexplored node w which minimizes:

$$\pi(w) = \min_{(v,w): v \in S} \pi(v) + c(v, w)$$

Brute force implementation. Test all edges.

Efficient implementation. Maintain a priority queue of unexplored vertices, prioritized by $\pi(w)$.

Q. How to maintain π ?

A. When exploring v , for each (v, w) leaving v , update

$$\pi(w) = \min \{ \pi(w), \pi(v) + c(v, w) \}.$$

13

Dijkstra's Algorithm: Java Implementation

Dijkstra's algorithm.

- Initialize $\pi[v] = \infty$ and $\pi[s] = 0$.

```
private void dijkstra(WeightedDigraph G, int s) {
    IndexMinPQ<Double> pq = new IndexMinPQ<Double>(G.V());
    pq.insert(s, pi[s]);
    while (!pq.isEmpty()) {
        int v = pq.delMin();
        for (Edge e : G.adj(v)) {
            int w = e.other(v);
            if (pi[w] > pi[v] + e.weight()) {
                pi[w] = pi[v] + e.weight();
                pred[w] = e;
                if (pq.contains(w)) pq.decrease(w, pi[w]);
                else pq.insert(w, pi[w]);
            }
        }
    }
}
```

14

Indexed Priority Queue

Indexed PQ.

- Assume elements k are named 0 to $N-1$.
- Insert, delete min, test if empty. [PQ ops]
- Decrease key, contains. [ST-like ops]

Return Type	Method	Action
	<code>IndexMinPQ(int N)</code>	create empty pq on N elements
void	<code>insert(int k, Value val)</code>	add element k with given value
void	<code>decrease(int k, Value val)</code>	decrease value of element k
int	<code>delMin()</code>	delete and return the smallest element
boolean	<code>isEmpty()</code>	is the PQ empty?
boolean	<code>contains(int k)</code>	does the PQ contain element k?

15

Indexed Priority Queue: Array Implementation

Indexed PQ: array implementation.

- Maintain vertex indexed array `vals[k]`.
- Insert key: change `vals[k]`.
- Decrease key: change `vals[k]`.
- Delete min: scan through `vals[k]` for each vertex v .
- Maintain a boolean array `marked[k]` to mark vertices in the PQ.


Operation	Array	Dijkstra
insert	1	$\times V$
delete-min	V	$\times V$
decrease-key	1	$\times E$
is-empty	1	$\times V$
contains	1	$\times V$
total	V^2	

16

Indexed Priority Queue

Indexed PQ: binary heap implementation.

- Assume elements k are named 0 to $N-1$.
- Store priorities in a binary heap.



k	pq	qp	pri
0	-	6	47
1	8	5	18
2	4	9	91
3	3	3	42
4	6	2	14
5	1	8	83
6	0	4	78
7	10	11	77
8	5	1	7
9	2	10	81
10	9	7	45
11	7	-	-

How to decrease key of element i ? Bubble it up.

How to know which heap node to bubble up? Maintains an extra array $qp[i]$ that stores the heap index of element i .

Indexed Binary Heap: Java Implementation

```
public class IndexMinPQ<Value extends Comparable> {
    private int N;
    private int[] pq, qp;
    private Comparable[] vals;

    public IndexMinPQ(int MAXN) {
        vals = new Comparable[MAXN + 1];
        pq = new int[MAXN + 1];
        qp = new int[MAXN + 1];
        for (int i = 0; i <= MAXN; i++) qp[i] = -1;
    }

    private boolean greater(int i, int j) {
        return vals[pq[i]].compareTo(vals[pq[j]]) > 0;
    }

    private void exch(int i, int j) {
        int swap = pq[i]; pq[i] = pq[j]; pq[j] = swap;
        qp[pq[i]] = i; qp[pq[j]] = j;
    }
}
```

17

18

Indexed Binary Heap: Java Implementation

```
public void insert(int k, Value val) {
    N++;
    qp[k] = N;
    pq[N] = k;
    vals[k] = val;
    swim(N);
}

public int delMin() {
    int min = pq[1];
    qp[min] = -1;
    exch(1, N--);
    sink(1);
    return min;
}

public void decrease(int k, Value val) {
    vals[k] = val;
    swim(qp[k]);
}

public boolean contains(int k) { return qp[k] != -1; }
```

19

Dijkstra's Algorithm: Priority Queue Choice

The choice of priority queue matters in Dijkstra's implementation.

- Array: $\Theta(V^2)$.
- Binary heap: $O(E \log V)$.
- Fibonacci heap: $O(E + V \log V)$.

Operation	Array	Binary heap	Fib heap	Dijkstra
insert	1	$\log V$	1^\dagger	$\times V$
delete-min	V	$\log V$	$\log V^\dagger$	$\times V$
decrease-key	1	$\log V$	1^\dagger	$\times E$
is-empty	1	1	1	$\times V$
contains	1	1	1	$\times V$
total	V^2	$E \log V$	$E + V \log V$	

† amortized

20

Dijkstra's Algorithm: Priority Queue Choice

The choice of priority queue matters in Dijkstra's implementation.

- Array: $\Theta(V^2)$.
- Binary heap: $O(E \log V)$.
- Fibonacci heap: $O(E + V \log V)$.

Best choice depends on whether graph is sparse or dense.

- 2,000 vertices, 1 million edges. Heap: 2-3x slower.
- 100,000 vertices, 1 million edges. Heap: 500x faster.
- 1 million vertices, 2 million edges. Heap: 10,000x faster.

Bottom line.

- Array implementation optimal for dense graphs.
- Binary heap far better for sparse graphs.
- Fibonacci heap best in theory, but not in practice.

Priority First Search


Priority first search. Maintain a set of explored vertices S , and grow S by exploring edges with exactly one endpoint leaving S .

DFS. Edge from vertex which was discovered most recently.

BFS. Edge from vertex which was discovered least recently.

Prim. Edge of minimum weight.

Dijkstra. Edge to vertex which is closest to s .



21

22

Edsger W. Dijkstra

The question of whether computers can think is like the question of whether submarines can swim.

Do only what only you can do.

In their capacity as a tool, computers will be but a ripple on the surface of our culture. In their capacity as intellectual challenge, they are without precedent in the cultural history of mankind.

The use of *COBOL* cripples the mind; its teaching should, therefore, be regarded as a criminal offence.

APL is a mistake, carried through to perfection. It is the language of the future for the programming techniques of the past: it creates a new generation of coding bums.



Edsger Dijkstra
Turing award 1972

Bellman-Ford-Moore

23

Application: Currency Conversion

Currency conversion. Given currencies and exchange rates, what is best way to convert one ounce of gold to US dollars?


- 1 oz. gold \Rightarrow \$327.25.
- 1 oz. gold \Rightarrow £208.10 \Rightarrow 208.10 (1.5714) \Rightarrow \$327.00.
- 1 oz. gold \Rightarrow 455.2 Francs \Rightarrow 304.39 Euros \Rightarrow \$327.28.

Currency	£	Euro	¥	Franc	\$	Gold
UK Pound	1.0000	0.6853	0.005290	0.4569	0.6368	208.100
Euro	1.4599	1.0000	0.007721	0.6677	0.9303	304.028
Japanese Yen	189.050	129.520	1.0000	85.4694	120.400	39346.7
Swiss Franc	2.1904	1.4978	0.011574	1.0000	1.3929	455.200
US Dollar	1.5714	1.0752	0.008309	0.7182	1.0000	327.250
Gold (oz.)	0.004816	0.003295	0.0000255	0.002201	0.003065	1.0000

Application: Currency Conversion

Graph formulation.

- Vertex = currency.
- Edge = transaction, with weight equal to exchange rate.
- Find path that maximizes **product** of weights.




25

26

Application: Currency Conversion

Reduction to shortest path problem.

- Let $\gamma(v, w)$ be exchange rate from currency v to w .
- Let $c(v, w) = -\lg \gamma(v, w)$.
- Shortest path with costs c corresponds to best exchange sequence.




Difficulty. Shortest path problem with **negative weights**.

Shortest Paths with Negative Weights: Failed Attempts

Dijkstra. Can fail if negative edge costs.



Re-weighting. Adding a constant to every edge weight can fail.




27


28

Shortest Paths: Negative Cost Cycles

Negative cycle. Directed cycle whose sum of edge costs is negative.



Observation. If negative cycle C on path from s to t , then shortest path can be made arbitrarily negative by spinning around cycle; otherwise, there exists a shortest s - t path that is simple.



29

Dynamic Programming

Dynamic programming.

- Initialize $pi[v] = \infty$, $pi[s] = 0$.
- Repeat V times: relax each edge.

```

for (int i = 1; i <= V; i++) { ← phase i
  for (int v = 0; v < G.V(); v++) {
    for (Edge e : G.adj(v)) {
      int w = e.other(v);
      if (pi[w] > pi[v] + e.weight()) {
        pi[w] = pi[v] + e.weight();
        pred[w] = v;
      }
    }
  }
}
    
```

31

Review: Edge Relaxation


Valid weights. For all v , $pi[v]$ is length of some path from s to v .

Edge relaxation.

- Consider edge $e = v \rightarrow w$.
- If current path from s to v plus edge $v \rightarrow w$ is better than current path to w , then update current path to w .

```

if (pi[w] > pi[v] + e.weight()) {
  pi[w] = pi[v] + e.weight();
  pred[w] = v;
}
    
```



30

Dynamic Programming: Analysis

Running time. $\Theta(EV)$.

Invariant. At end of phase i , $pi[v] \leq$ length of any path from s to v using at most i edges.

Theorem. Assuming no negative cycles, upon termination $pi[v]$ is the length of the shortest path from s to v .

and $pred[v]$ are the shortest paths

32

Bellman-Ford-Moore

Observation. If $pi[v]$ doesn't change during phase i , no need to relax any edges leaving v in phase $i+1$.

FIFO implementation. Maintain **queue** of vertices whose distance changed.

be careful to keep at most one copy of each vertex on queue

Running time. Still $\Omega(EV)$ in worst case, but much faster in practice.

Bellman-Ford-Moore Algorithm

Bellman-Ford-Moore. Initialize $pi[v] = \infty$, $pi[s] = 0$.

```

Queue<Integer> q = new Queue<Integer>();
q.enqueue(s);

while (!q.isEmpty()) {
    int v = q.dequeue();
    marked[v] = false;
    for (Edge e : G.adj(v)) {
        int w = e.other(v);
        if (pi[w] > pi[v] + e.weight()) {
            pi[w] = pi[v] + e.weight();
            pred[w] = e;
            if (!marked[w]) {
                marked[w] = true;
                q.enqueue(w);
            }
        }
    }
}
    
```

33

34

Single Source Shortest Paths Implementation: Cost Summary

Algorithm	Worst Case	Best Case	Space
Dijkstra (classic) †	V^2	V^2	$E + V$
Dijkstra (heap) †	$E \log V$	$E + V$	$E + V$
Dynamic programming ‡	$E V$	$E V$	$E + V$
Bellman-Ford ‡	$E V$	$E + V$	$E + V$

† nonnegative costs
‡ no negative cycles


Remark 1. Negative weights makes the problem harder.

Remark 2. Negative cycles makes the problem intractable.

Arbitrage

Arbitrage. Is there an arbitrage opportunity in currency graph?

- Ex: $\$1 \Rightarrow 1.3941 \text{ Francs} \Rightarrow 0.9308 \text{ Euros} \Rightarrow \1.00084 .
- Is there a negative cost cycle?
- Fastest algorithm very valuable!




$$-0.4793 + 0.5827 - 0.1046 < 0$$

35

36

Negative Cycles

If negative cycle reachable from s , Bellman-Ford gets stuck in infinite loop, updating vertices in a cycle.




Finding a negative cycle. If any vertex v is updated in phase v , there must be a negative cycle, and we can trace back $\text{pred}[v]$ to find it.

37

Negative Cycle Detection

Goal. Identify a negative cycle (reachable from any vertex).

Solution. Add 0-cost edge from artificial source s to each vertex v . Run Bellman-Ford from vertex s .



38