

*La distance n'y fait rien; il n'y a que le premier pas qui coûte.
[The distance is nothing; it is only the first step that costs.]*

— Marie Anne de Vichy-Chamrond, marquise du Deffand,
letter to Jean le Rond d'Alembert, July 7, 1763

Cecil Graham: *What is a cynic?*

Lord Darlington: *A man who knows the price of everything, and the value of nothing.*

Cecil Graham: *And a sentimentalist, my dear Darlington, is a man who sees an absurd value in everything and doesn't know the market price of any single thing.*

— Oscar Wilde, *Lady Windermere's Fan, A Play About a Good Woman* (1892)

*25½ Minimum Cost Flows

Now we consider a significant generalization of the maximum-flow problem that cannot be solved by simply modifying the graph and applying a standard flow algorithm. For this problem, every edge e in the input graph has both a capacity $c(e)$ and a cost $\$(e)$. The cost function describes the cost of sending a unit of flow through the edges; thus, the cost any flow f is defined as follows:

$$\$(f) = \sum_{e \in E} \$(e) \cdot f(e).$$

Costs can either be positive, negative, or zero. The *minimum-cost flow* problem is to compute a feasible flow with minimum cost, instead of a feasible flow with maximum value.



Standard reductions:

- Reduce maximum flow to min-cost circulation: add $t \rightarrow s$ with cost -1 and infinite capacity, and give all other edges cost 0 .
- Reduce min-cost flow (with arbitrary vertex balances) to minimum-cost circulation: Find a feasible flow and work in the residual graph. Most natural formulation for cycle canceling.
- Reduce min-cost flow to **transportation**—non-trivial vertex balances, non-negative costs, infinite capacities, zero lower bounds. Replace each edge e with two edges; move upper and lower bounds to balances on vertices. Most natural formulation for successive shortest path.

25½.1 Cycle Cancellation

Without loss of generality, it suffices to consider the *minimum-cost circulation* problem, where the network has no specified source or target, and we seek a minimum-cost flow with value 0 . For flow networks with a source s and target t , we can reduce to the minimum-cost circulation problem by adding a single edge $t \rightarrow s$ with infinite capacity and cost 0 . For flow networks with nontrivial balance constraints at the vertices, we first find a feasible flow f using the algorithm in the previous section and then seek a minimum-cost circulation in the residual graph G_f .

Fix a flow network $G = (V, E)$. Without loss of generality, assume that whenever G contains an edge $u \rightarrow v$, it does not contain the reversed edge $v \rightarrow u$. If G contains both $u \rightarrow v$ and $v \rightarrow u$, we can introduce a new vertex x and then replace $v \rightarrow u$ with two edges $v \rightarrow x$ and $x \rightarrow u$, each with capacity $c(v \rightarrow u)$ and cost $\$(v \rightarrow u)/2$.



Figure!

Fix a circulation f in G . Each edge in G and its reversal has both a *residual capacity* and a *residual cost*:

$$c_f(u \rightarrow v) := \begin{cases} c(u \rightarrow v) - f(u \rightarrow v) & \text{if } u \rightarrow v \in E \\ f(v \rightarrow u) & \text{if } v \rightarrow u \in E \end{cases}$$

$$\$(u \rightarrow v) := \begin{cases} \$(u \rightarrow v) & \text{if } u \rightarrow v \in E \\ -\$(v \rightarrow u) & \text{if } v \rightarrow u \in E \end{cases}$$

As before, the *residual network* G_f consists of all edges with non-zero residual capacity.



Figure showing flow f and residual graph G_f

Now let γ be a directed cycle in the residual graph G_f . Let $F = \min_{e \in \gamma} c_f(e)$ be the minimum residual capacity of edges in γ , and let $\Phi = \sum_{e \in \gamma} \(e) be the *sum* of the residual costs of edges in γ . Just as in the Ford-Fulkerson algorithm, we can **augment** the circulation f , by pushing F units of flow around γ , to obtain a new circulation f' :

$$f'(u \rightarrow v) = \begin{cases} f(u \rightarrow v) + F & \text{if } u \rightarrow v \in \gamma \\ f(u \rightarrow v) - F & \text{if } v \rightarrow u \in \gamma \\ f(u \rightarrow v) & \text{otherwise} \end{cases}$$

Straightforward calculation now implies that

$$\$(f') = \$(f) + \Phi \cdot F.$$

In particular, if $\Phi < 0$, the new circulation f' has lower cost than the original circulation f . We immediately conclude that **f is a minimum-cost circulation if and only if G_f contains no negative cycles.**



Figure showing negative cycle and updated flow f'

This observation implies that we can compute minimum-cost circulations using a **cycle cancelling** algorithm first proposed by Morton Klein in 1967, which is a straightforward variant of Ford and Fulkerson's augmenting path algorithm. We start by setting f to the all-zero circulation. Then we repeatedly augment f along an arbitrary negative-cost cycle in the residual graph G_f , until there are no more negative residual cycles.

In each iteration of the cycle canceling algorithm, we can find a negative-cost cycle in $O(VE)$ time using a straightforward modification of the Shimbel-Bellman-Ford shortest path algorithm. To bound the number of iterations, we assume that both the capacity and the cost of each edge is an integer, and we define

$$C = \max_{e \in E} c(e) \quad \text{and} \quad D = \max_{e \in E} |$(e)|.$$

The total cost of any feasible circulation clearly lies between $-ECD$ and ECD , and each augmentation step decreases the cost of the flow by a positive integer, and therefore by at least 1. Since we start with the zero circulation, we conclude that the algorithm requires at most ECD iterations, and therefore runs in $O(VE^2CD)$ time. As with the raw Ford-Fulkerson algorithm, this running time is exponential in the complexity of the input, and the algorithm may never terminate if capacities and/or costs are irrational.



For the network produced by our earlier reduction from maximum flows to min-cost flows, cycle cancelling **IS** Ford-Fulkerson, because a cycle in G_f is negative iff it contains $t \rightarrow s$.

Like Ford-Fulkerson, more careful choices of *which* cycle to cancel can lead to more efficient algorithms. Unfortunately, some obvious choices are NP-hard to compute, including the cycle with most negative cost and the negative cycle with the fewest edges. In the late 1980s, Andrew Goldberg and Bob Tarjan developed a minimum-cost flow algorithm that repeatedly cancels the so-called *minimum-mean cycle*, which is the cycle whose *average cost per edge* is smallest. By combining an algorithm of Karp to compute minimum-mean cycles in $O(VE)$ time, efficient dynamic tree data structures, and other sophisticated techniques that are (unfortunately) beyond the scope of this class, their algorithm achieves a running time of $O(VE^2 \log^2 V)$.

25½.2 Successive Shortest Paths



- Proposed by Ford and Fulkerson.
- Instead of trivial balance, assume non-negative costs, by saturating all negative-cost edges and considering the residual graph.
- Feasible flow f is optimal iff G_f has no negative cycles.
- Maintain pseudoflow f such that G_f has no negative cycles. If f is not a flow, augment f along the minimum-cost path in G_f from any node s with excess to any node t with deficit, and recurse.
- Lemma: Pushing flow along a shortest excess-to-deficit path cannot introduce negative cycles.
- Integer balances \implies at most $B = \sum_v |b(v)|$ iterations $\implies O(BVE)$ time via Bellman-Ford.
- For the network produced by our earlier reduction from maximum flows to min-cost flows, successive shortest paths is Ford-Fulkerson in disguise, because every path from s to t has cost 0 and is therefore a shortest path!

25½.2.1 Node Potentials and Reduced Costs



- Improvement due to Edmonds and Karp in 1969 (published 1972) and independently Nobuaki Tomizawa in 1970 (published 1971).
- Define node potentials and reduced costs. Shortest paths don't change. Cycle lengths don't change.
- No negative residual cycles \iff for some potential function, reduced residual costs are all non-negative
- Fix any node s . Let $dist_\pi(s, v)$ denote the shortest-path distance in G_f from s to v , using reduced costs with respect to π as distances. If reduced costs are non-negative with respect to π , then reduced costs are also non-negative with respect to π' , where $\pi'(v) = \pi(v) - dist(s, v)$ for all nodes v .
- Moreover, reduced costs with respect to π' are zero for all shortest-path edges in G_f . (Complementary slackness!)
- Modify successive shortest-path algorithm to maintain node potentials. At each iteration, set $\pi(v) \leftarrow \pi(v) - dist(s, v)$. (Alternatively, just let $\pi(v) = -dist(s, v)$?)
- Non-negative reduced costs \implies shortest path via Dijkstra in $O(E \log V)$ time $\implies O(BE \log V)$ time overall.

The fastest minimum-cost circulation algorithm currently known,¹ due to James Orlin in the

¹among algorithms whose running times do not depend on C and D

early 1990s, reduces the problem to $O(E \log V)$ iterations of Dijkstra's shortest-path algorithm and therefore runs in $O(E^2 \log^2 V)$ time.

25½.3 Maximum-Weight Matchings




This is an immediate application of successive shortest paths with node potentials. Do it really need/deserve a separate presentation?

Recall from the previous lecture note that we can find a maximum-cardinality matching in any bipartite graph in $O(VE)$ time by reduction to the standard maximum flow problem.

Now suppose the input graph has *weighted* edges, and we want to find the matching with maximum total *weight*. Given a bipartite graph $G = (U \times W, E)$ and a non-negative weight function $w: E \rightarrow \mathbb{R}$, the goal is to compute a matching M whose total weight $w(M) = \sum_{uw \in M} w(uw)$ is as large as possible. Maximum-weight matchings can't be found directly using standard maximum-flow algorithms², but we can modify the algorithm for maximum-cardinality matchings described in the previous note.

It will be helpful to reinterpret our earlier maximum-matching algorithm directly in terms of the original bipartite graph instead of the derived flow network. Our algorithm maintains a matching M , which is initially empty. We say that a vertex is *matched* if it is an endpoint of an edge in M . At each iteration, we find an *alternating path* π that starts and ends at unmatched vertices and alternates between edges in $E \setminus M$ and edges in M . Equivalently, let G_M be the directed graph obtained by orienting every edge in M from W to U , and every edge in $E \setminus M$ from U to W . An alternating path is just a directed path in G_M between two unmatched vertices. Any alternating path has odd length and has exactly one more edge in $E \setminus M$ than in M . The iteration ends by setting $M \leftarrow M \oplus \pi$, thereby increasing the number of edges in M by one. The max-flow/min-cut theorem implies that when there are no more alternating paths, M is a maximum matching.



A matching M with 5 edges, an alternating path π , and the augmented matching $M \oplus \pi$ with 6 edges.

To find the maximum-weight matching in a bipartite graph with weighted edges, we need to make only two changes to this algorithm.

- First, instead of looking for an arbitrary alternating path at each iteration, we look for the alternating path π such that $M \oplus \pi$ has largest weight. Suppose we weight the edges in

²However, maximum-flow algorithms can be modified to compute maximum *weighted* flows, where every edge has both a capacity and a weight, and the goal is to maximize $\sum_{u \rightarrow v} w(u \rightarrow v) f(u \rightarrow v)$.


the residual graph G_M as follows:

$$w'(u \rightarrow w) = -w(uw) \quad \text{for all } uw \notin M$$

$$w'(w \rightarrow u) = w(uw) \quad \text{for all } uw \in M$$

We now have $w(M \oplus \pi) = w(M) - w'(\pi)$. Thus, the correct augmenting path π must be the directed path in G_M with minimum total residual weight $w'(\pi)$. In other words, we are implicitly applying the successive shortest-path algorithm!

- Second, because the matching with maximum weight may not be the matching with the maximum cardinality, we return the heaviest matching considered in any iteration of the algorithm, not just in the last iteration.



A maximum-weight matching is not necessarily a maximum-cardinality matching.

Before we determine the running time of the algorithm, we need to check that it actually finds the maximum-weight matching. After all, it's a greedy algorithm, and greedy algorithms don't work unless you prove them into submission! Let M_i denote the maximum-weight matching in G with exactly i edges. In particular, $M_0 = \emptyset$, and the global maximum-weight matching is equal to M_i for some i . (The figure on the previous page shows M_1 and M_2 for the same graph.) Let G_i denote the directed residual graph for M_i , let w_i denote the residual weight function for M_i as defined above, and let π_i denote the directed path in G_i such that $w_i(\pi_i)$ is minimized. To simplify the proof, I will assume that there is a unique maximum-weight matching M_i of any particular size; this assumption can be enforced by applying a consistent tie-breaking rule. With this assumption in place, the correctness of our algorithm follows inductively from the following lemma.

Lemma 1. *If G contains a matching with $i + 1$ edges, then $M_{i+1} = M_i \oplus \pi_i$.*

Proof: I will prove the equivalent statement $M_{i+1} \oplus M_i = \pi_{i-1}$. To simplify notation, call an edge in $M_{i+1} \oplus M_i$ red if it is an edge in M_{i+1} , and blue if it is an edge in M_i .

The graph $M_{i+1} \oplus M_i$ has maximum degree 2, and therefore consists of pairwise disjoint paths and cycles, each of which alternates between red and blue edges. Since G is bipartite, every cycle must have even length. The number of edges in $M_{i+1} \oplus M_i$ is odd; specifically, $M_{i+1} \oplus M_i$ has $2i + 1 - 2k$ edges, where k is the number of edges that are in both matchings. Thus, $M_{i+1} \oplus M_i$ contains an odd number of paths of odd length, some number of paths of even length, and some number of cycles of even length.

Let γ be a cycle in $M_{i+1} \oplus M_i$. Because γ has an equal number of edges from each matching, $M_i \oplus \gamma$ is another matching with i edges. The total weight of this matching is exactly $w(M_i) - w_i(\gamma)$, which must be less than $w(M_i)$, so $w_i(\gamma)$ must be positive. On the other hand, $M_{i+1} \oplus \gamma$ is a matching with $i + 1$ edges whose total weight is $w(M_{i+1}) + w_i(\gamma) < w(M_{i+1})$, so $w_i(\gamma)$ must be negative! We conclude that no such cycle γ exists; $M_{i+1} \oplus M_i$ consists entirely of disjoint paths.

Exactly the same reasoning implies that no path in $M_{i+1} \oplus M_i$ has an even number of edges.

Finally, since the number of red edges in $M_{i+1} \oplus M_i$ is one more than the number of blue edges, the number of paths that start with a red edge is exactly one more than the number of paths that start with a blue edge. The same reasoning as above implies that $M_{i+1} \oplus M_i$ does not contain a blue-first path, because we can pair it up with a red-first path.

We conclude that $M_{i+1} \oplus M_i$ consists of a single alternating path π whose first edge is red. Since $w(M_{i+1}) = w(M_i) - w_i(\pi)$, the path π must be the one with minimum weight $w_i(\pi)$. \square

We can find the alternating path π_i using a single-source shortest path algorithm. (See, I told you we were using successive shortest paths!) Modify the residual graph G_i by adding zero-weight edges from a new source vertex s to every *unmatched* node in U , and from every *unmatched* node in W to a new target vertex t , exactly as in our unweighted matching algorithm. Then π_i is the shortest path from s to t in this modified graph. Since M_i is the maximum-weight matching with i vertices, G_i has no negative cycles, so this shortest path is well-defined. We can compute the shortest path in G_i in $O(VE)$ time using Shimbel-Bellman-Ford, so the overall running time of our algorithm is $O(V^2E)$.

The residual graph G_i has negative-weight edges, so we can't immediately speed up the algorithm by replacing Shimbel-Bellman-Ford with Dijkstra's algorithm. However, we can use the same repricing trick as for the more general successive shortest-path algorithm, or equivalently, a variant of Johnson's all-pairs shortest path algorithm, to improve the running time to $O(VE \log V)$. Let $d_i(v)$ denote the distance from s to v in the residual graph G_i , using the distance function w_i . Let \tilde{w}_i denote the modified distance function $\tilde{w}_i(u \rightarrow v) := d_{i-1}(u) + w_i(u \rightarrow v) - d_{i-1}(v)$. As we argued above (and in the discussion of Johnson's algorithm), shortest paths with respect to w_i are still shortest paths with respect to \tilde{w}_i . Moreover, $\tilde{w}_i(u \rightarrow v) > 0$ for every edge $u \rightarrow v$ in G_i :

- If $u \rightarrow v$ is an edge in G_{i-1} , then $w_i(u \rightarrow v) = w_{i-1}(u \rightarrow v)$ and $d_{i-1}(v) \leq d_{i-1}(u) + w_{i-1}(u \rightarrow v)$.
- If $u \rightarrow v$ is not in G_{i-1} , then $w_i(u \rightarrow v) = -w_{i-1}(v \rightarrow u)$ and $v \rightarrow u$ is an edge in the shortest path π_{i-1} , so $d_{i-1}(u) = d_{i-1}(v) + w_{i-1}(v \rightarrow u)$.

Let $\tilde{d}_i(v)$ denote the shortest path distance from s to v with respect to the distance function \tilde{w}_i . Because \tilde{w}_i is positive everywhere, we can quickly compute $\tilde{d}_i(v)$ for all v using Dijkstra's algorithm. This gives us both the shortest alternating path π_i and the distances $d_i(v) = \tilde{d}_i(v) + d_{i-1}(v)$ needed for the next iteration.

Exercises

1. Describe and analyze an algorithm for the following problem, first posed and solved by the German mathematician Carl Jacobi in the early 1800s.³

Disponantur nn quantitates $h_k^{(i)}$ quaecunque in schema Quadrati, ita ut k habeantur n series horizontales et n series verticales, quarum quaeque est n terminorum. Ex illis quantitatibus eligantur n transversales, i.e. in seriebus horizontalibus simul atque verticalibus diversis positae, quod fieri potest $1.2 \dots n$ modis; ex omnibus illis modis quaerendum est is, qui summam n numerorum electorum suppeditet maximam.

For the few students who are not fluent in mid-19th century academic Latin, here is a modern English translation of Jacobi's problem. Suppose we are given an $n \times n$ matrix M .

³Carl Gustav Jacob Jacobi. De investigando ordine systematis aequationum differentialium vulgarium cujuscunque. *J. Reine Angew. Math.* 64(4):297–320, 1865. Posthumously published by Carl Borchardt.

Describe and analyze an algorithm that computes a permutation σ that maximizes the sum $\sum_{i=1}^n M_{i,\sigma(i)}$, or equivalently, permutes the columns of M so that the sum of the elements along the diagonal is as large as possible.

2. Let G be a directed flow network whose edges have costs, but which contains no negative-cost cycles. Prove that one can compute a minimum-cost maximum flow in G using a variant of Ford-Fulkerson that repeatedly augments the (s, t) -path of *minimum total cost* in the current residual graph. What is the running time of this algorithm?
3. An (s, t) -*series-parallel* graph is an directed acyclic graph with two designated vertices s (the *source*) and t (the *target* or *sink*) and with one of the following structures:
 - **Base case:** A single directed edge from s to t .
 - **Series:** The union of an (s, u) -series-parallel graph and a (u, t) -series-parallel graph that share a common vertex u but no other vertices or edges.
 - **Parallel:** The union of two smaller (s, t) -series-parallel graphs with the same source s and target t , but with no other vertices or edges in common.
 - (a) Describe an efficient algorithm to compute a maximum flow from s to t in an (s, t) -series-parallel graph with arbitrary edge capacities.
 - (b) Describe an efficient algorithm to compute a *minimum-cost* maximum flow from s to t in an (s, t) -series-parallel graph whose edges have *unit* capacity and arbitrary costs.
 - * (c) Describe an efficient algorithm to compute a *minimum-cost* maximum flow from s to t in an (s, t) -series-parallel graph whose edges have *arbitrary* capacities and costs.
4. Every year, Professor Dumbledore assigns the instructors at Hogwarts to various faculty committees. There are n faculty members and c committees. Each committee member has submitted a list of their *prices* for serving on each committee; each price could be positive, negative, zero, or even infinite. For example, Professor Snape might declare that he would serve on the Student Recruiting Committee for 1000 Galleons, that he would *pay* 10000 Galleons to serve on the Defense Against the Dark Arts Course Revision Committee, and that he would not serve on the Muggle Relations committee for any price.

Conversely, Dumbledore knows how many instructors are needed for each committee, as well as a list of instructors who would be suitable members for each committee. (For example: “Dark Arts Revision: 5 members, anyone but Snape.”) If Dumbledore assigns an instructor to a committee, he must pay that instructor’s price from the Hogwarts treasury.

Dumbledore needs to assign instructors to committees so that (1) each committee is full, (2) no instructor is assigned to more than three committees, (3) only suitable and willing instructors are assigned to each committee, and (4) the total cost of the assignment is as small as possible. Describe and analyze an efficient algorithm that either solves Dumbledore’s problem, or correctly reports that there is no valid assignment whose total cost is finite.

5. Vince wants to borrow a certain amount of money from his friends as cheaply as possible, possibly after first arranging a sequence of intermediate loans. Each of Vince’s friends have

a different amount of money that they can lend (possibly zero). For any two people x and y , there is a maximum amount of money (possibly zero or infinite) that x is willing to lend to y and a certain profit (possibly zero or even negative) that x expects from any loan to y .

For example, suppose Vince wants to borrow \$100 from his friends Ben and Naomi, who have the following constraints:

- Ben has \$500 available to lend.
- Ben is willing to lend up to \$150 to Vince at a profit of 20¢ per dollar.
- Ben is willing to lend up to \$50 to Naomi, at a loss of 10¢ per dollar.
- Naomi has \$50 available to lend.
- Naomi is willing to lend any amount of money to Vince, at a profit of 10¢ per dollar.
- Naomi is not willing to lend money to Ben.

If Vince borrows \$100 directly from Ben, he needs \$120 to pay off the loan. If Vince borrows \$50 from Ben and \$50 from Naomi, he needs \$115 to pay off the loan: \$60 for Ben and \$55 for Naomi. But if Vince asks Naomi to borrow \$50 from Ben and then borrows the entire \$100 from Naomi, then he needs only \$110 to pay off Naomi, who can then pay off Ben with just \$45. With the same constraints, the maximum amount of money that Vince can borrow is \$250.

Describe and analyze an algorithm that finds a sequence of loans that minimizes the amount Vince needs to pay everyone off, or correctly reports that Vince cannot borrow his desired amount. The input has the following components:

- An array $Money[1..n]$, where $Money[i]$ is the amount of money that friend i has.
- An array $MaxLoan[1..n, 0..n]$, where $MaxLoan[i, j]$ is the amount of money that friend i is willing to lend to friend j . “Friend 0” is Vince.
- An array $Profit[1..n, 0..n]$, where $Profit[i, j]$ is the profit per dollar that friend i expects from any loan to friend j . Again, “friend 0” is Vince.
- The total amount T that Vince wants to borrow.