
Combinatorial Optimization

lecture notes, winter term 2019/20
Prof. Dr. Yann Disser



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Contents

1	Introduction	3
2	Shortest Paths	5
2.1	Generic Search	5
2.1.1	Fibonacci Heaps	6
2.2	All-Pairs Shortest Paths	9
2.2.1	All-Pairs Dijkstra	12
2.2.2	A Duality Approach to Vertex Potentials	14
2.3	Goal-Directed Search	15
3	Maximum Flows	17
3.1	Dinic Algorithm	17
3.1.1	Blocking Flows	19
3.2	Push-Relabel Algorithm	21
4	Minimum-Cost Flows	27
4.1	Minimum-Mean-Cycle Cancelling	30
4.1.1	Finding a minimum-mean cycle	34
4.2	Successive Shortest Paths	36
4.2.1	Capacity Scaling	40
4.3	Orlin's Algorithm	42
5	Maximum Cardinality Matchings	49
5.1	Bipartite Graphs Revisited	49
5.2	Characterization of Matchings	51
5.3	Edmonds' Algorithm	55



1 Introduction

This lecture focuses on advanced combinatorial algorithms. We will revisit most topics of the lecture *Algorithmic Discrete Mathematics* and provide a deeper and more complete picture of the state of the art in each of them. In particular, we will introduce more sophisticated algorithmic techniques that will allow us to design algorithms that come close to achieving the best known running times for various problems. Some of these techniques are based on structural insights in the theory of mathematical optimization that were discussed in the lectures *Introduction to Optimization* and *Discrete Optimization*. Specifically, we will first consider shortest path problems (Chapter 2), then maximum flow problems (Chapters 3 and 4), and finally matching problems (Chapters 5 and ??). This advanced lecture relies on fundamental concepts and notations of computational complexity theory and graph theory. We refrain from reintroducing these notions here and refer to the lecture notes of the course *Algorithmic Discrete Mathematics* instead.



2 Shortest Paths

In this chapter, we consider the problem of finding shortest paths in a given directed graph $G = (V, E)$ with arc weights $c: E \rightarrow \mathbb{R}$. For convenience, we consider an arbitrary but fixed directed graph G with weights c throughout this section. As usual, we use the notations $c(u, v) := c((u, v))$ for $(u, v) \in E$, $c(u, u) := 0$ for $u \in V$, and $c(u, v) := \infty$ for $(u, v) \notin E$ and $u \neq v$. We recall the following fundamental definition of distances in G with respect to c .

Definition 2.1. We let \mathcal{W}_{uv} denote the set of all walks from u to v in G and define^a

$$d_c^{(k)}(u, v) := \min\{c(W) \mid W = (v_0, e_1, \dots, v_k) \in \mathcal{W}_{uv} \text{ with } k' \leq k\},$$

$$d_c(u, v) := \lim_{k \rightarrow \infty} d_c^{(k)}(u, v).$$

If P is a u - v -path in G with $c(P) = d_c(u, v)$, we say that P is a *shortest path* in G with respect to c .

^a We use the convention that $d_c^{(k)}(u, v) = \infty$ if $\mathcal{W}_{uv} = \emptyset$.

Throughout this chapter, we will be interested in computing the values $d_c(u, v)$ for $u, v \in V$ with the understanding that we can obtain shortest paths from these values by backtracking using depth-first search.

2.1 Generic Search

In the lecture *Algorithmic Discrete Mathematics* we saw several *greedy* algorithms that compute spanning trees and/or shortest paths by growing a component from the starting vertex s . Abstractly, these greedy strategies can be formulated as follows, where the different algorithms differ in their choices of $g(u, v)$.

Algorithm: GENERICSEARCH(G, c, s)

input: graph $G = (V, E)$, weights $c: E \rightarrow \mathbb{R}$, vertex $s \in V$
output: distances $(d_{sv})_{v \in V}$, parents $(p_v)_{v \in V}$

$$d_{sv} \leftarrow \begin{cases} 0, & v = s, \\ \infty, & \text{otherwise.} \end{cases}$$

$$p_v \leftarrow v \quad \forall v \in V$$

$$R \leftarrow V \quad \text{(INSERT)}$$

while $R \neq \emptyset$:

$$u \leftarrow \arg \min_{v \in R} \{d_{sv}\}$$

$$R \leftarrow R \setminus \{u\} \quad \text{(EXTRACTMIN)}$$

for $v \in \Gamma(u) \cap R$:

if $g(u, v) < d_{sv}$:

$$d_{sv} \leftarrow g(u, v) \quad \text{(DECREASEVALUE)}$$

$$p_v \leftarrow u$$

For example, if we set $g(u, v) = d_{su} + 1$ we recover the simple graph traversal algorithm breadth-first search (BFS) that can be used to compute shortest paths in an unweighted graph. If we set $g(u, v) = c(u, v)$, we obtain Prim's algorithm to compute a minimum spanning tree in an undirected graph. Finally, if we set $g(u, v) = d_{su} + c(u, v)$, then `GENERICSEARCH` turns into Dijkstra's algorithm for computing shortest paths in directed graphs. We will see another parameterization of the greedy strategy in Section 2.3 below.

We recall the following result for `PRIM` and `DIJKSTRA`. Note that BFS can be implemented to run in linear time.

Theorem 2.2. `GENERICSEARCH` can be implemented with a running time of $\mathcal{O}(m + n \log n)$.

This result relies on an efficient data structure, called a *priority queue*, that supports the operations `INSERT` to insert elements, `EXTRACTMIN` to extract the element associated with the smallest value, and `DECREASEVALUE` to decrease the value associated with some element. We will now show in detail how such a data structure can be obtained.

2.1.1 Fibonacci Heaps

We describe a data structure called a *Fibonacci heap*. A *heap* is a rooted tree T with the additional property that the root of every subtree is associated with the smallest value amongst all values associated to nodes of the subtree. For convenience, in the following we identify nodes with the values associated to them. With this, a heap is a rooted tree where every subtree has its smallest value as the root.

A Fibonacci heap stores its values in a forest of rooted trees, each of which is a heap. Additionally, it maintains a list R of all roots. The crucial feature of a Fibonacci heap is that it ensures that the heaps it contains have relatively few children but contain a relatively large number of nodes. This structure makes extraction of the smallest element easy: Since the smallest element is guaranteed to be a root and the number of roots is relatively small, we can afford to search R to identify the smallest element.

We now describe how a Fibonacci heap maintains its special structure by defining its operations. The operation `MAKEROOT` and `MAKECHILD` are auxiliary routines that are used by the main operations. In the following, we write T_x to denote the subtree rooted at node x and $|T_x|$ to denote the number of nodes in T_x .

Algorithm: `MAKEROOT(x)`

$R \leftarrow R \cup \{x\}$
 $p_x \leftarrow x$, **unmark** x

Algorithm: `MAKECHILD(c, x)`

$R \leftarrow R \setminus \{c\}$
 $p_c \leftarrow x$

Algorithm: `INSERT(x)`

`MAKEROOT`(x)
 $x_{\min} \leftarrow \min\{x, x_{\min}\}$

Algorithm: `EXTRACTMIN`

$r \leftarrow x_{\min}$
for each *child* c of r :
 `MAKEROOT`(c)
 $R \leftarrow R \setminus \{r\}$
while R contains vertices $x < y$ of the same degree :
 `MAKECHILD`(y, x)
 $x_{\min} \leftarrow \min\{x \in R\}$
return r

Algorithm: `DECREASEVALUE(x)`

$x_{\min} \leftarrow \min\{x, x_{\min}\}$
if $x < p_x$:
 mark x
 while x is marked :
 $y \leftarrow p_x$
 `MAKEROOT`(x)
 $x \leftarrow y$
 if $x \neq p_x$:
 mark x

Observe that we allow nodes to lose two children before moving the corresponding subtree to R . Additionally, we defer reducing the length of R until the next `EXTRACTMIN` operation. We will see that this “laziness” is crucial for the performance of the Fibonacci heap. While this laziness introduces some variability in the shape of the rooted trees, we will show that the shape is closely related to the famous *Fibonacci numbers*.

Definition 2.3. The *Fibonacci numbers* $(F_k)_{k \in \mathbb{N} \cup \{0\}}$ are defined by

$$F_k := \begin{cases} 0 & \text{for } k = 0, \\ 1 & \text{for } k = 1, \\ F_{k-1} + F_{k-2} & \text{for } k \geq 2. \end{cases}$$

The Fibonacci numbers have many beautiful properties. In particular, the following holds (the proof is left as an exercise).

Proposition 2.4. The Fibonacci numbers satisfy

- (i) $F_{k+2} = 1 + \sum_{i=0}^k F_i$ for all $k \in \mathbb{N} \cup \{0\}$,
- (ii) $F_{k+2} \geq \phi^k$ for all $k \in \mathbb{N} \cup \{0\}$, where $\phi := (1 + \sqrt{5})/2$ is the golden ratio.

With this, we can establish that the trees of a Fibonacci heap must be relatively large, compared to the degrees of their roots. Here and throughout we assume that Fibonacci heaps start empty and are built up using the above operations.

Lemma 2.5. For every node x of a Fibonacci heap it holds that $|T_x| \geq F_{k+2}$, where k is the number of children of x .

Proof. Let s_k be the smallest possible size of a subtree whose root has k children. Observe that s_k is monotonically increasing in k : We can always reduce both the degree and number of vertices of the root by invoking `DECREASEVALUE` on a child, which implies $s_{k-1} \leq s_k$. We show that $s_k \geq F_{k+2}$ by induction on k . The induction basis for $k \leq 1$ holds, since $F_3 = 2$ and $F_2 = 1$.

Now let T be the smallest possible tree whose root x has $k \geq 2$ children. Let c_1, \dots, c_k denote these children in the order in which they were attached to x . Observe that every vertex that is not in the root list R can have lost at most one child during `DECREASEVALUE` operations. This means that c_i has $k' \geq i - 2$ children. By induction, and since s_i and F_i grow monotonically in i , for $i \geq 2$, we have $|T_{c_i}| \geq s_{k'} \geq s_{i-2} \geq F_i$. Overall, by Proposition 2.4 (i), we obtain

$$s_k = 1 + |T_{c_1}| + \sum_{i=2}^k |T_{c_i}| \geq 2 + \sum_{i=2}^k F_i = 1 + \sum_{i=0}^k F_i \stackrel{\text{Prop. 2.4}}{=} F_{k+2}. \quad \square$$

This immediately implies that Fibonacci trees maintain small degrees.

Corollary 2.6. Every node in a Fibonacci heap with n nodes has at most $\log_\phi n$ children.

Proof. If there was a node with $k > \log_\phi n$ children, by Proposition 2.4 (ii) and Lemma 2.5, the corresponding subtree would have to have size at least $F_{k+2} \geq \phi^k > n$. \square

We are now ready to show that Fibonacci heaps support arbitrary sequences of operations efficiently. Note that the bound below immediately yields the running time of Theorem 2.2.

Theorem 2.7. Fibonacci heaps can be implemented to take time $\mathcal{O}(m + n \log n)$ for inserting $\mathcal{O}(n)$ elements and performing $\mathcal{O}(m)$ EXTRACTMIN and DECREASEVALUE operations.

Proof. We define a potential function

$$\Phi := \kappa \cdot (|R| + 2n_{\text{marked}}),$$

where n_{marked} is the number of marked nodes, and $\kappa \in \mathbb{N}$ is a suitably large constant. Let $\Phi_0 := 0$ and let Φ_i denote the value of Φ at the end of the i -th INSERT, EXTRACTMIN or DECREASEVALUE operation. Let further t_i denote the running time needed for the i -th operation and $a_i := t_i + \Phi_i - \Phi_{i-1}$ denote the *amortized running time* in step i . Because the potential is always non-negative, the total running time for N operations can be bounded via

$$\sum_{i=0}^N t_i = \sum_{i=0}^N (a_i - \Phi_i + \Phi_{i-1}) = \sum_{i=0}^N a_i - \Phi_N + \Phi_0 \leq \sum_{i=0}^N a_i.$$

It thus suffices to bound the amortized running time in each step.

First note that the INSERT operation has a constant running time and only changes the potential by a constant amount, hence its amortized running time is $\mathcal{O}(1)$.

The operation EXTRACTMIN consists of three phases. First, all children of the smallest node in the root list R are moved to the root list themselves. By Corollary 2.6, this phase takes time at most $\mathcal{O}(\log n)$ and the increase in potential is also at most $\mathcal{O}(\log n)$. The amortized running time of this phase thus is $\mathcal{O}(\log n)$. In a second phase, we merge trees until no two roots have the same degrees. We can accomplish this going over all roots and inserting them into a list of length $\lfloor \log_{\phi} n \rfloor + 1$ at a position corresponding to their degree. Whenever this position is already occupied, we merge the two trees and put the resulting tree into the next position of the list, etc. Since, in every step, we either find a free position or reduce the number of trees, the running time of the second phase is linear in $|R|$, i.e., it is bounded by $\kappa' \cdot |R|$ for some constant $\kappa' \in \mathbb{N}$. After the second phase, the root list is R' and has length $|R'| \leq \log_{\phi} n + 1$, thus the potential is decreased by $\kappa \cdot (|R| - |R'|)$. For $\kappa \geq \kappa'$, we obtain an amortized running time of $\kappa' \cdot |R| - \kappa \cdot (|R| - |R'|) \leq \kappa \cdot |R'| = \mathcal{O}(\log n)$ for the second phase. In the third phase, the new minimum is determined by going over all elements in R' . This takes time linear in $|R'|$ and does not change Φ , hence the amortized running time of this phase is $\mathcal{O}(\log n)$. Overall, EXTRACTMIN has an amortized running time of $\mathcal{O}(\log n)$.

The running time of the DECREASEVALUE(x) operation is constant if the value of x does not get smaller than the value of its parent (and the potential does not change in this case). Otherwise, the running time is linear in the number $n_{\text{marked}}(x)$ of (consecutive) marked ancestors of x , i.e., it is bounded by $\kappa'' \cdot n_{\text{marked}}(x)$ for some constant κ'' . These ancestors are unmarked and added to the root list R . Since at most one vertex is marked by DECREASEVALUE(x), the change in potential is at most $\kappa \cdot (n_{\text{marked}}(x) + 1) - \kappa \cdot 2n_{\text{marked}}(x) + \kappa \cdot 2$. For $\kappa \geq \kappa''$, we obtain an amortized running time of $\kappa'' \cdot n_{\text{marked}}(x) - \kappa \cdot n_{\text{marked}}(x) + 3\kappa < 3\kappa = \mathcal{O}(1)$.

Finally, observe that there can be at most $\mathcal{O}(n)$ EXTRACTMIN operations. Thus, the total amortized cost for all INSERT, EXTRACTMIN and DECREASEVALUE operations is

$$\mathcal{O}(n) + \mathcal{O}(n \log n) + \mathcal{O}(m) = \mathcal{O}(m + n \log n).$$

As discussed above, this also bounds the total running time of these operations. Note, however, that we cannot guarantee the individual operations to be efficient. \square

2.2 All-Pairs Shortest Paths

We move on to the problem of computing *all* shortest paths in G , or, more precisely, we consider the following problem.

All-Pairs Shortest Paths Problem

input: graph $G = (V, E)$; edge weights $c: E \rightarrow \mathbb{R}$
problem: find shortest u - v -paths in G for all $u, v \in V$

Of course, we can solve this problem by solving the single-source shortest path problem starting from every vertex $s \in V$. However, optimum substructure (see lecture *Algorithmic Discrete Mathematics*) suggests that we can do better by reusing optimal subpaths instead of recomputing all paths from scratch n times.

Recall that the algorithm of Bellman-Ford(-Moore) is based precisely on optimum substructure in that it extends shortest paths/walks into longer shortest paths/walks in each iteration. The final formulation of `BELLMANFORD(MOORE)` is a refinement of the following dynamic program (the final algorithm does not explicitly maintain $d_{sv}^{(k)}$ for all $k \in \{1, \dots, n-1\}$ and only iterates over actual arcs in the inner loop).

Algorithm: `DYNAMICSP(G, c, s)`

input: directed graph $G = (V, E)$, weights $c: E \rightarrow \mathbb{R}$, vertex $s \in V$

output: distances $d_{sv} = d_c(s, v)$ for all $v \in V$

$$d_{sv}^{(0)} \leftarrow \begin{cases} 0, & \text{if } v = s, \\ \infty, & \text{otherwise.} \end{cases}$$

for $k \leftarrow 1, 2, \dots, n-1$:

for $v \in V$:

$d_{sv}^{(k)} \leftarrow \min_{u \in V} \{d_{su}^{(k-1)} + c(u, v)\}$

return $(d_{sv}^{(n-1)})_{v \in V}$

Correctness of this algorithm is based on the following insights (for proofs see lecture *Algorithmic Discrete Mathematics*).

Proposition 2.8. Let $G = (V, E)$ be a directed graph, let $u \in V$, and let $c: E \rightarrow \mathbb{R}$ not induce negative cycles reachable from u . Then, for all $v \in V$ reachable from u , there exists a shortest u - v -path in G with respect to c and it holds that $d_c(u, v) = d_c^{(n-1)}(u, v)$.

Lemma 2.9. The following recurrence holds:

$$d_c^{(k)}(u, v) = \begin{cases} c(u, v) & \text{if } k = 1, \\ \min_{w \in V} \{d_c^{(k-1)}(u, w) + c(w, v)\} & \text{if } k \geq 2. \end{cases}$$

In particular, `DYNAMICSP` correctly computes $d_c^{(n-1)}(s, v)$ for all $v \in V$.

These insights immediately imply the following result.

Theorem 2.10. `DYNAMICSP` solves the single-source shortest path problem in time $\Theta(n^3)$ on directed graphs with arcs weights that do not induce negative cycles reachable from s .

Note that DYNAMICSP only computes shortest paths for graphs without negative cycles. Importantly, however, Lemma 2.9 does allow negative cycles, which means that the values $d_c^{(n-1)}(s, v)$ for all $v \in V$ computed by DYNAMICSP are still correct in this case. This yields an easy way of verifying whether c induces negative cycles (the proof is left as an exercise):

Proposition 2.11. The graph G contains a negative cycle reachable from s if and only if there are $u, v \in V$ with $d_c^{(n-1)}(s, v) > d_c^{(n-1)}(s, u) + c(u, v)$.

With this, we can solve the all-pairs shortest paths problem naively by repeating DYNAMICSP from all starting vertices, which yields the algorithm REPEATBFM. A more systematic, but equivalent, approach is algorithm PARALLELBFM that computes the values $d_{uv}^{(k)}$ in order of increasing k (i.e., it changes the order of computation of the dynamic program). For convenience, we let $V = \{1, \dots, n\}$ in the following and introduce the cost matrix $C = (c_{uv})_{u, v \in V} \in \mathbb{R}^{n \times n}$ with $c_{uv} := c(u, v)$. Similarly, we organize the values $d_c^{(k)}(u, v)$ and $d_{uv}^{(k)}$ into matrices $D_c^{(k)}$ and $D^{(k)}$.

Algorithm: REPEATBFM(V, C)

input: vertices $V = \{1, \dots, n\}$,
weights $C \in (\mathbb{R} \cup \{\infty\})^{n \times n}$

output: distance matrix $D_c^{(n-1)}$

$D^{(1)} \leftarrow C$

for $u \leftarrow 1, \dots, n$:

for $k \leftarrow 2, \dots, n-1$:

for $v \leftarrow 1, \dots, n$:

$d_{uv}^{(k)} \leftarrow \min_{v' \in V} \{d_{uv'}^{(k-1)} + c(v', v)\}$

return $D^{(n-1)}$

Algorithm: PARALLELBFM(G, C)

input: vertices $V = \{1, \dots, n\}$,
weights $C \in (\mathbb{R} \cup \{\infty\})^{n \times n}$

output: distance matrix $D_c^{(n-1)}$

$D^{(1)} \leftarrow C$

for $k \leftarrow 2, \dots, n-1$:

for $u \leftarrow 1, \dots, n$:

for $v \leftarrow 1, \dots, n$:

$d_{uv}^{(k)} \leftarrow \min_{v' \in V} \{d_{uv'}^{(k-1)} + c(v', v)\}$

return $D^{(n-1)}$

We can further refine this naive algorithm by expressing it using matrix multiplications in the *min tropical semiring* (or *min-plus-algebra*), where '+' plays the role of the product and 'min' plays the role of addition (see algorithm MATRIXBFM). We denote this matrix multiplication using the $\otimes_{\min,+}$ -operator.

Of course, the running time of MATRIXBFM is still $\Theta(n^4)$. However, this reformulation allows for an elegant trick to improve the running time. Consider the computation of a conventional matrix product $C \times C \times \dots \times C = C^n$. Assuming that n is a power of 2, we can significantly reduce our computational effort by computing the equivalent product $(C \times \dots \times C) \times (C \times \dots \times C) = \dots = (\dots((C^2)^2)\dots)^2 = C^n$. This product only requires $\log_2 n$ matrix multiplications instead of $n-1$. To apply this reformulation, we only used that the matrix product is associative. It turns out that this is still the case in the tropical semiring (the proof is left as an exercise).

Lemma 2.12. The matrix product $\otimes_{\min,+}$ in the tropical semiring is associative.

We can therefore reformulate MATRIXBFM using repeated matrix squaring in order to improve overall performance to $\Theta(n^3 \log n)$ (see algorithm MATRIXSQUARING). This is significantly better than naively computing single-source shortest paths from every starting location via REPEATBFM!

Algorithm: MATRIXBFM(G, C)

input: vertices $V = \{1, \dots, n\}$,
 weights $C \in (\mathbb{R} \cup \{\infty\})^{n \times n}$
output: distance matrix D_c

 $D^{(1)} \leftarrow C$
for $k \leftarrow 2, \dots, n-1$:

 $\quad \lfloor D^{(k)} \leftarrow D^{(k-1)} \otimes_{\min,+} C$
return $D^{(n-1)}$

Algorithm: MATRIXSQUARING(G, C)

input: vertices $V = \{1, \dots, n\}$,
 weights $C \in (\mathbb{R} \cup \{\infty\})^{n \times n}$
output: distance matrix D_c

 $Q^{(1)} \leftarrow C$
for $k \leftarrow 2, \dots, \lceil \log_2(n-1) \rceil$:

 $\quad \lfloor Q^{(k)} \leftarrow Q^{(k-1)} \otimes_{\min,+} Q^{(k-1)}$
return $Q^{(\lceil \log_2(n-1) \rceil)}$

In terms of the underlying dynamic program, we achieved the above improvement by shrinking the dynamic programming table and changing dependencies. We now improve the dynamic program further by identifying a more suitable form of optimum substructure, based on induced subgraphs of the form $G[\{1, 2, \dots, w\}]$ of G for increasing values of $w \in \{0, \dots, n\}$.

Definition 2.13. For $u, v, w \in V = \{1, \dots, n\}$ we define

$$d_{uv}^{(\leq 0)} := c(u, v),$$

$$d_{uv}^{(\leq w)} := \min\{d_{uv}^{(\leq w-1)}, d_{uw}^{(\leq w-1)} + d_{wv}^{(\leq w-1)}\}.$$

We formalize the intuition behind this definition:

Lemma 2.14. Let $u, v \in V = \{1, \dots, n\}$ and $w \in V \cup \{0\}$ such that $d_{uv}^{(\leq w)} < \infty$. Then, there exists a u - v -walk P in $G[V_{uvw}]$ of length $d_{uv}^{(\leq w)}$ that is not longer than any u - v -path in $G[V_{uvw}]$, where $V_{uvw} := \{v' \in V \mid v' \leq w\} \cup \{u, v\}$.

Proof. We prove the statement via induction on w . The induction basis for $w = 0$ is trivial.

By Definition 2.13, $d_{uv}^{(\leq w)} < \infty$ implies that $d_{uv}^{(\leq w-1)} < \infty$ or $d_{uw}^{(\leq w-1)} + d_{wv}^{(\leq w-1)} < \infty$. In the former case, the u - v -walk $P_{uv}^{(\leq w-1)}$ of the induction hypothesis is a u - v -walk in $G[V_{uvw}]$ of the required length. In the latter case, the u - v -walk $P_{uw}^{(\leq w-1)} \oplus P_{wv}^{(\leq w-1)}$ of the induction hypothesis is a u - v -walk in $G[V_{uvw}]$ of the required length. It remains to show that $d_{uv}^{(\leq w)} \leq c(P_{uv})$.

If $w \notin P_{uv}$, by the induction hypothesis, we have $d_{uv}^{(\leq w)} \leq d_{uv}^{(\leq w-1)} = c(P_{uv}^{(\leq w-1)}) \leq c(P_{uv})$. Otherwise, let $P_{uw} \oplus P_{wv} := P_{uv}$ be such that $P_{uw} \cap P_{wv} = \{w\}$. We have $P_{uw} \subseteq \{1, \dots, w-1\} \cup \{u, w\}$ and $P_{wv} \subseteq \{1, \dots, w-1\} \cup \{w, v\}$. By induction, this implies $d_{uv}^{(\leq w)} \leq c(P_{uw}^{(\leq w-1)}) + c(P_{wv}^{(\leq w-1)}) \leq c(P_{uw}) + c(P_{wv}) = c(P_{uv})$. \square

This interpretation of Definition 2.13 suggests a different dynamic program called the *Floyd-Warshall* algorithm.

Algorithm: FLOYDWARSHALL(G, c)

input: vertices $V = \{1, \dots, n\}$, weights $C \in (\mathbb{R} \cup \{\infty\})^{n \times n}$
output: distance matrix D_c

 $D^{(\leq 0)} \leftarrow C$
for $v' \leftarrow 1, \dots, n$:

 $\quad \lfloor$ **for** $u \leftarrow 1, \dots, n$:

 $\quad \quad \lfloor$ **for** $v \leftarrow 1, \dots, n$:

 $\quad \quad \quad \lfloor d_{uv}^{(\leq v')} \leftarrow \min\{d_{uv}^{(\leq v'-1)}, d_{uv'}^{(\leq v'-1)} + d_{v'v}^{(\leq v'-1)}\}$
 $\quad \quad \quad \rfloor$
 $\quad \quad \rfloor$
 $\quad \rfloor$
 \rfloor
return $D^{(\leq n)}$

It is easy to see that this algorithm is correct and yields another (minor) improvement to the running time.

Theorem 2.15. If $c : V \rightarrow \mathbb{R}$ does not induce negative cycles in $G = (V, E)$, then FLOYDWARSHALL solves the all pairs shortest paths problem in time $\Theta(n^3)$.

Proof. By Lemma 2.14, for every $u, v \in V$, there is a u - v -walk W of length $c(W) = d_{uv}^{(\leq n)} \leq d_c(u, v)$ in G . We can decompose W into a u - v -path P and a set \mathcal{C} of cycles (see lecture *Algorithmic Discrete Mathematics*). Since c does not induce negative cycles, we have $d_c(u, v) \geq d_{uv}^{(\leq n)} = c(W) = c(P) + \sum_{C \in \mathcal{C}} c(C) \geq c(P)$. Since P cannot be shorter than a shortest u - v -path in G , equality must hold, and, in particular, $d_{uv}^{(\leq n)} = d_c(u, v)$. This implies that $D^{(\leq n)} = D_c$ and FLOYDWARSHALL computes the correct solution. The algorithm obviously has cubic running time. \square

In fact, the all-pairs shortest paths problem is a very important problem in algorithm theory and many more efficient algorithms have been developed. However, all of these algorithms have a running time of the form $\mathcal{O}(n^3)/n^{o(1)}$, which means that their running times are better than $\mathcal{O}(n^3)$ only by a subpolynomial factor (for example by $\log n / \log \log n$). Because of our inability to obtain better algorithms, many scientists believe that the following conjecture holds. This conjecture is sometimes used as a basis for conditional proofs of hardness for other problems.

Conjecture 2.16 (APSP Conjecture). *No algorithm can solve the All-Pairs Shortest Paths problem in time $\mathcal{O}(n^{3-\varepsilon})$ for any fixed $\varepsilon > 0$.*

2.2.1 All-Pairs Dijkstra

If all arc weights are non-negative, we can obtain a running time of $\mathcal{O}(mn + n^2 \log n)$ simply by repeating DIJKSTRA from all start vertices $s \in V$. This is never worse than the running time of FLOYDWARSHALL and better on sparse graphs, i.e., on graphs with $m \in o(n^2)$. We will now see a way of adapting this algorithm for the case where arc weights can be negative. The idea is to modify arc weights in such a way that they become non-negative, while preserving shortest paths. We can achieve this if we can find a potential function in the following sense.

Definition 2.17. The *reduced costs* c_π with respect to $\pi : V \rightarrow \mathbb{R}$ are defined by

$$c_\pi(u, v) := c(u, v) + \pi(u) - \pi(v),$$

for all $u, v \in V$. If $c_\pi(u, v) \geq 0$ for all $u, v \in V$, we say that π is a *vertex potential*.

Crucially, vertex potentials have the property that they do not alter the set of shortest paths.

Proposition 2.18. Let $\pi : V \rightarrow \mathbb{R}$ and P be a u - v -path in G with $u, v \in V$. Then, P is a shortest path with respect to c if and only if it is a shortest path with respect to c_π .

Proof. For every u - v -walk W we have

$$c_\pi(W) = \sum_{(a,b) \in W} c_\pi(a, b) = \sum_{(a,b) \in W} (c(a, b) + \pi(a) - \pi(b)) = \pi(u) - \pi(v) + \sum_{(a,b) \in W} c(a, b) = \delta_{uv} + c(W),$$

where $\delta_{uv} := \pi(u) - \pi(v)$ is a constant that only depends on u and v . This means that the costs of all u - v -walks are shifted by the same additive constant. Thus, the set of shortest paths is preserved. \square

Proposition 2.18 implies that, if we have a vertex potential $\pi: V \rightarrow \mathbb{R}$, we can solve the all-pairs shortest paths problem for arbitrary weights by applying DIJKSTRA using the reduced costs c_π .

Algorithm: ALLPAIRSDIJKSTRA(G, c)

input: vertices $V = \{1, \dots, n\}$, weights $c: E \rightarrow \mathbb{R}$

output: distances $(d_c(u, v))_{u, v \in V}$

$\pi \leftarrow \text{POTENTIAL}(G, c)$

for $u \in V$:

⌊ $(d_{uv})_{v \in V} \leftarrow \text{DIJKSTRA}(G, c_\pi, u)$

return $(d_c(u, v) - \pi(u) + \pi(v))_{u, v \in V}$

We still need a way of finding vertex potentials. This can be accomplished as follows.

Algorithm: POTENTIAL(G, c)

input: vertices $V = \{1, \dots, n\}$, weights $c: E \rightarrow \mathbb{R}$

output: vertex potential

$s \leftarrow n + 1$; $n' \leftarrow n + 1$

$G' \leftarrow (V \cup \{s\}, E \cup (\{s\} \times V))$

$$c'(u, v) \leftarrow \begin{cases} 0, & \text{if } u = s \text{ and } v \in V \cup \{s\}, \\ \infty, & \text{if } u \in V \text{ and } v = s, \\ c(u, v), & \text{if } u, v \in V. \end{cases}$$

$(d_{sv}^{(n'-1)})_{v \in V \cup \{s\}} \leftarrow \text{BELLMANFORD(MOORE)}(G', c', s)$

return $(d_{sv}^{(n'-1)})_{v \in V}$

We show that POTENTIAL finds a vertex potential if there are no negative cycles. Conversely, we cannot hope for a potential if there are negative cycles, since otherwise correctness of DIJKSTRA would guarantee existence of shortest paths, leading to a contradiction.

Proposition 2.19. If c does not induce negative cycles, POTENTIAL computes a vertex potential in time $\mathcal{O}(mn)$.

Proof. Since c does not induce negative cycles in G , Proposition 2.11 yields that $d_c^{(n-1)}(s, v) \leq d_c^{(n-1)}(s, u) + c(u, v)$ for all $u, v \in V$. It follows that $d_c^{(n-1)}(s, u) + c(u, v) - d_c^{(n-1)}(s, v) \geq 0$ for all $u, v \in V$. This means that $\pi(v) := d_c^{(n-1)}(s, v)$ defines a vertex potential. The running time of POTENTIAL is dominated by the running time of BELLMAN-FORD(-MOORE) (see lecture *Algorithmic Discrete Mathematics*). \square

With this, we obtain the following result.

Theorem 2.20. If c does not induce negative cycles in G , then ALLPAIRSDIJKSTRA solves the all pairs shortest paths problem in time $\mathcal{O}(n(m + n \log n))$.

Proof. Correctness follows from Propositions 2.18 and 2.19. The running time follows from Proposition 2.19 together with the fact that DIJKSTRA can be implemented with a running time of $\mathcal{O}(m + n \log n)$ (Theorem 2.2). \square

Remark 2.21. Note that this result does not contradict the APSP Conjecture, since it only provides an improved running time for sparse graphs. In dense graphs, ALLPAIRSDIJKSTRA and FLOYDWARSHALL provide the same theoretical performance guarantee. For practical purposes, FLOYDWARSHALL is often superior in dense graphs, since it is much leaner in terms of the hidden constant in the \mathcal{O} -notation and is much easier to implement.

2.2.2 A Duality Approach to Vertex Potentials

We now expose a tight relationship between negative cycles and vertex potentials in terms of linear programming theory. We start with an integer linear programming (ILP) formulation for finding negative cycles:

$$(NC) \quad \begin{array}{ll} \min & \sum_{(u,v) \in E} c_{uv} x_{uv} \\ \text{s.t.} & \sum_{(v,u) \in E} x_{vu} = \sum_{(u,v) \in E} x_{uv}, \quad \forall v \in V \\ & x_{uv} \in \mathbb{N} \cup \{0\} \quad \forall (u,v) \in E \end{array}$$

Clearly, this ILP is feasible, since $x_{uv} = 0$ for all $u, v \in V$ trivially is a feasible solution. If there are no negative cycles, then (NC) is also bounded by 0, otherwise it is unbounded. Furthermore, the polytope defined by the constraints of (NC) is integral, which means that we do not need to require x_{uv} to be integral to have an integral optimum.

Proposition 2.22. The LP-relaxation of (NC) is integral.

Proof. First observe that the constraint matrix of (NC) is the incidence matrix of the underlying graph G . We know that the incidence matrix of a graph is totally unimodular, which implies that the corresponding polytope is integral (see lecture *Discrete Optimization*). Thus, the LP-relaxation of (NC) is integral, as claimed. \square

We can therefore consider the LP-relaxation (NCP) instead of (NC).

$$(NCP) \quad \begin{array}{ll} \min & \sum_{(u,v) \in E} c_{uv} x_{uv} \\ \text{s.t.} & \sum_{(u,v) \in E} x_{uv} - \sum_{(v,u) \in E} x_{vu} = 0, \quad \forall v \in V \\ & x_{uv} \geq 0 \quad \forall (u,v) \in E \end{array} \quad (NCD) \quad \begin{array}{ll} \max & 0 \\ \text{s.t.} & y_v - y_u \leq c_{uv}, \quad \forall (u,v) \in E \end{array}$$

Observation 2.23. The linear program (NCP) is always feasible, and it is bounded if and only if c does not induce negative cycles.

We will now consider the dual linear program (NCD) and use weak duality to relate the existence of vertex potentials to the existence of negative cycles. Note that the following proposition also follows from Proposition 2.19 and the correctness of DIJKSTRA, as discussed above.

Proposition 2.24. A directed graph $G = (V, E)$ with arc weights $c: E \rightarrow \mathbb{R}$ admits a vertex potential if and only if c does not induce negative cycles.

Proof. Consider the dual LP (NCD) for (NCP). Trivially, (NCD) is always bounded. By identifying $\pi(v) \equiv y(v)$, we see that (NCD) is feasible (and bounded) if and only if we can find a vertex potential.

On the other hand, by strong duality (see lecture *Introduction to Optimization*), the dual (NCD) is feasible and bounded if and only if the primal (NCP) is feasible and bounded. From Observation 2.23 we know that this is the case precisely if c does not induce negative cycles. \square

2.3 Goal-Directed Search

In many applications, in particular routing, we are not interested in all shortest paths in a graph but rather a shortest s - t -path between specific vertices s and t . Intuitively, from a theoretical worst-case perspective, this problem is equivalent to the single-source shortest path problem, since the shortest s - t -path may contain all other shortest paths from s . However, in practice, many heuristic techniques have been developed that are tailored to finding a shortest path to the destination as quickly as possible. While these algorithms do not perform better than general shortest path algorithms in the worst case, they are often superior on real graphs like road networks. We present one of the most important heuristics and establish its correctness.

As a motivation, consider a densely connected graph on points of the plane with arc weights being Euclidean distances (see Figure 2.1). In such a graph, DIJKSTRA explores a ball of radius $d_c(s, t)$ before finding a shortest s - t -path. We can significantly reduce the search space by invoking DIJKSTRA simultaneously from s (in $G = (V, E)$) and from t (in the reverse graph (V, \bar{E})). The corresponding search space consists of two balls of radius $\frac{1}{2} \cdot d_c(s, t)$, which, intuitively, should contain about half the number of vertices as before. We can further reduce the search space if we have some reasonable estimate $h_t(v)$ for the distance from v to t for all $v \in V$. Such an estimate might, for example, be the length of the straight-line connection between v and t in the plane. With such an estimate, we can direct DIJKSTRA towards t by favoring vertices with smaller estimates when choosing where to proceed. One way of doing this is the algorithm A^* below.

Algorithm: $A^*(G, c, s, t)$

input: graph $G = (V, E)$,
weights $c: E \rightarrow \mathbb{R}_{\geq 0}$,
vertices $s, t \in V$

output: distance $d_c(s, t)$

$$d_{sv} \leftarrow \begin{cases} 0, & v = s, \\ \infty, & \text{otherwise.} \end{cases}$$

$R \leftarrow V$

while $t \in R$:

$u \leftarrow \arg \min_{v \in R} \{d_{sv} + h_t(v)\}$

$R \leftarrow R \setminus \{u\}$

for $v \in \Gamma(u) \cap R$:

$d_{sv} \leftarrow \min\{d_{sv}, d_{su} + c(u, v)\}$

return d_{st}

Algorithm: A^* -REFORMULATION(G, c, s, t)

input: graph $G = (V, E)$,
weights $c: E \rightarrow \mathbb{R}_{\geq 0}$,
vertices $s, t \in V$

output: distance $d_c(s, t)$

$$d'_{sv} \leftarrow \begin{cases} h_t(s), & v = s, \\ \infty, & \text{otherwise.} \end{cases}$$

$R \leftarrow V$

while $t \in R$:

$u \leftarrow \arg \min_{v \in R} \{d'_{sv}\}$

$R \leftarrow R \setminus \{u\}$

for $v \in \Gamma(u) \cap R$:

$d'_{sv} \leftarrow \min\{d'_{sv}, d'_{su} - h_t(u) + c(u, v) + h_t(v)\}$

return $d'_{st} - h_t(t)$

This algorithm can be expressed as a parametrization of GENERICSEARCH with $g(u, v) = d'_{su} - h_t(u) + c(u, v) + h_t(v)$, see A^* -REFORMULATION. We have to be careful not to be too generous with our estimation $h_t(v)$, otherwise we may lose correctness of the algorithm. The following theorem gives a relatively weak condition for h_t under which A^* is guaranteed to compute a shortest s - t -path. Whether or not our choice of h_t actually improves the practical performance entirely depends on the structure of G and c .




Figure 2.1: An intuitive illustration of the search spaces of DIJKSTRA (red), bidirectional DIJKSTRA (blue), and A* (green) in the plane with Euclidean distances.

Theorem 2.25. If h_t is monotone, i.e., if $h_t(u) \leq c(u, v) + h_t(v)$ for all $(u, v) \in E$, then A* is correct.

Proof. Consider algorithm A*-REFORMULATION. This algorithm is equivalent to A* with $d'_{sv} = d_{sv} + h_t(v)$. On the other hand, the algorithm is nothing else than DIJKSTRA(G, c_π, s) with vertex potential $\pi(v) = -h_t(v)$ for all $v \in V$. Observe that π defines a potential, since $c_\pi(u, v) = c(u, v) + \pi(u) - \pi(v) = c(u, v) - h_t(u) + h_t(v) \geq 0$ by monotonicity of h_t . Correctness thus follows from correctness of DIJKSTRA and Proposition 2.18. \square

3 Maximum Flows

In this chapter we revisit the maximum flow problem defined as follows. We refer to the lecture *Algorithmic Discrete Mathematics* for all basic definitions including residual graphs, augmenting paths and flow decompositions, as well as fundamental theorems like the max-flow-min-cut theorem, Menger's theorem, etc. For the remainder of the chapter we will consider a fixed but arbitrary network $(G = (V, E), \mu, s, t)$.

Maximum Flow Problem

input: network $(G = (V, E), \mu, s, t)$
problem: find s - t -flow in G of maximum value

3.1 Dinic Algorithm

Recall the Edmonds-Karp algorithm that implements the Ford-Fulkerson method for computing a maximum s - t -flow by increasing flow along augmenting paths until we obtain a maximum flow:

Algorithm: EDMONDSKARP(G, μ, s, t)

input: network $(G = (V, E), \mu, s, t)$
output: maximum s - t -flow

$f \leftarrow 0$
while $\exists s$ - t -path in G_f :
 $P \leftarrow$ shortest s - t -path in G_f (unweighted)
 $\delta \leftarrow \min_{e \in P} \mu_f(e)$
 AUGMENT(P, δ)
return f

Recall that the length of the shortest path P is monotonically increasing during an execution of EDMONDSKARP. This means that, conceptually, we can subdivide an execution into phases of constant path lengths. We know that there can be at most $n - 1$ such phases, since this is the maximum number of arcs any s - t -path can have. We showed that EDMONDSKARP executes each phase in time $\mathcal{O}(m^2)$ on average, which yields an overall running time of $\mathcal{O}(m^2n)$. We now present the algorithm of Dinic, which is based on an idea how to perform each phase more efficiently.

Intuitively, we aim to execute a phase in a single step, i.e., to simultaneously saturate all shortest s - t -paths in G_f . To do this, we consider the level graph of G_f , which is defined as follows.

Definition 3.1. The *level graph* $G_L = (V, E_L)$ of a graph $G = (V, E)$ with respect to a source vertex $s \in V$ is a subgraph of G defined by $E_L := \{(u, v) \in E \mid d(s, v) = d(s, u) + 1\}$.

We observe crucial properties of the level graph.

Observation 3.2. The level graph G_L is acyclic and contains all (unweighted) shortest s - v -paths of G for all $v \in V$.

The fact that the level graph $G_{f,L}$ of G_f contains all shortest augmenting paths implies that EDMONDSKARP increases the flow only along arcs of $G_{f,L}$ in each iteration. The idea of the Dinic algorithm is to increase flow along more than one path in $G_{f,L}$ in order to simultaneously saturate all s - t -paths in $G_{f,L}$. The algorithm achieves this by computing a blocking flow in $G_{f,L}$, which is defined as follows.

Definition 3.3. An s - t -flow f in a network (G, μ, s, t) is called a *blocking flow* if there is no augmenting path in $G \cap G_f$.

With this, we can state Dinic's algorithm, based on a subroutine for computing blocking flows.

Algorithm: DINIC(G, μ, s, t)

input: network $(G = (V, E), \mu, s, t)$

output: maximum s - t -flow

$f \leftarrow 0$

while \exists s - t -path in G_f :

$G_{f,L} \leftarrow (V, \{(v, w) \in G_f \mid d_{G_f}(s, w) = d_{G_f}(s, v) + 1\})$

$f' \leftarrow \text{BLOCKINGFLOW}(G_{f,L}, \mu_f, s, t)$

$f \leftarrow f + f'$

return f

We will now show that the length of a shortest path in G_f is strictly increasing in each iteration of DINIC. In addition, because $G_{f,L}$ is acyclic, we will be able to efficiently compute a blocking flow.

Proposition 3.4. DINIC computes a maximum s - t -flow in at most $n - 1$ iterations. In each iteration $d_{G_f}(s, t)$ strictly increases.

Proof. By definition of the algorithm, once DINIC terminates, there is no augmenting path in G_f , which implies that f is a maximum s - t -flow (see lecture *Algorithmic Discrete Mathematics*). It remains to show that this happens after $n - 1$ iterations at the latest. To prove this, we show that the distance $d_{G_f}(s, t)$ between s and t in G_f increases in each iteration.

Let $f, f'' = f + f'$ be the flows before and after some iteration of DINIC. We show that $d_{G_f}(s, t) < d_{G_{f''}}(s, t)$, i.e., that every shortest s - t -path in $G_{f''}$ (if it exists) is longer than every shortest s - t -path in G_f . Let $G_{f,L}$ be the level graph of G_f , and let $G_{f,\bar{L}} := G_f \cup \bar{G}_{f,L}$ be the graph that consists of the arcs of G_f and the reverse arcs of $G_{f,L}$. It holds that $G_{f''} \subseteq G_{f,\bar{L}}$, since all new arcs in the residual graph after the iteration must be reverse arcs of $G_{f,L}$. Moreover, $G_{f,\bar{L}}$ contains no arcs (u, v) with $d_{G_f}(s, v) > d_{G_f}(s, u) + 1$: obviously such arcs cannot be part of G_f and for $(u, v) \in \bar{G}_{f,L}$ it even holds that $d_{G_f}(s, v) = d_{G_f}(s, u) - 1$. By definition, the arcs $(u, v) \in G_{f,\bar{L}}$ with $d_{G_f}(s, v) = d_{G_f}(s, u) + 1$ are part of $G_{f,L}$. Since f' is a blocking flow for $G_{f,L}$, there is no s - t -path in $G_{f''} \cap G_{f,L}$. Hence, every shortest s - t -path P in $G_{f''} \subseteq G_{f,\bar{L}}$ must use at least one arc $(u, v) \in G_{f''} \setminus G_{f,L}$, which must therefore satisfy $d_{G_f}(s, v) \leq d_{G_f}(s, u)$. Since P is an s - t -path, we obtain

$$d_{G_{f''}}(s, t) = d_{G_f}(s, t) - d_{G_f}(s, s) = \sum_{(u,v) \in P} (d_{G_f}(s, v) - d_{G_f}(s, u)) \leq |P| - 1 = d_{G_{f''}}(s, t) - 1,$$

i.e., $d_{G_{f''}}(s, t) < d_{G_f}(s, t)$, as claimed. □

3.1.1 Blocking Flows

We still need a way of efficiently computing a blocking flow, i.e., we need to solve the following problem.

Blocking Flow Problem

input: acyclic network $(G = (V, E), \mu, s, t)$

problem: find blocking s - t -flow

Crucially, we restrict ourselves to acyclic inputs, since the level graph for which we need a blocking flow in DINIT is acyclic (Observation 3.2). This gives us the existence of a *topological ordering* of the graph.

Definition 3.5. A topological ordering of a graph $G = (V, E)$ is an ordering (v_1, \dots, v_n) of the vertices of G , i.e., $V = \{v_1, \dots, v_n\}$, with $i < j$ for all $(v_i, v_j) \in E$.

A topological ordering can be computed efficiently (proof left as an exercise).

Proposition 3.6. A directed graph is acyclic if and only if it has a topological ordering. In that case a topological ordering can be computed in time $\mathcal{O}(n + m)$.

A natural approach to finding a blocking flow is to “push” as much flow forward, in topological order, as possible. If we never decrease flow values, the last vertex, before t , that is still overflowing will never be able to push more flow. So we may as well send flow back until it is not overflowing anymore and afterwards keep all flow entering and leaving this vertex unchanged. We can continue alternating push and balancing phases, until no vertex (other than t) has excess left. We obtain the following algorithm.

Algorithm: BLOCKINGFLOW(G, μ, s, t)

input: acyclic network $(G = (V, E), \mu, s, t)$

output: blocking flow

$(s = v_1, \dots, v_n = t) \leftarrow \text{TOPOLOGICALSORT}(G, s)$ (we set $\delta^-(s) = \delta^+(t) = \emptyset$)

$L_{v \in V} \leftarrow \emptyset, f \leftarrow 0, \text{ex}_f(s) \equiv \infty$

for $n - 1$ **times:**

 /* push step */

for $v \leftarrow v_1, \dots, v_n$ **with** $\text{ex}_f(v) > 0$:

for $(v, w) \in \delta^+(v)$ **with** w **not marked:**

$e \leftarrow (v, w)$

$\gamma \leftarrow \text{PUSH}(v, e, \infty)$ (send as much flow as possible)

$L_w \leftarrow L_w \oplus (e, \gamma)$ (remember push)

 /* balancing step */

$i \leftarrow \max\{i < n \mid \text{ex}_f(v_i) > 0\}$

while $\text{ex}_f(v_i) > 0$ **and** $v_i \neq s$:

$(e, \gamma) \leftarrow \text{EXTRACTLAST}(L_{v_i})$

$\text{PUSH}(v_i, \bar{e}, \gamma)$ (send as much flow as possible back)

mark v_i

return f

Algorithm: PUSH(v, e, γ_{\max} [default = ∞])

$\gamma \leftarrow \min\{\text{ex}_f(v), \mu_f(e), \gamma_{\max}\}$

AUGMENT(e, γ)

return γ

Note that the function $f : E \rightarrow \mathbb{R}_{\geq 0}$ that the above algorithm is maintaining is not a valid flow until the very end, since it does not satisfy flow conservation. Observe however, that no vertex other than s may underflow, i.e., have negative excess. Functions with relaxed flow conservation in this sense will become important later, which warrants the following definition.

Definition 3.7. A function $f : E \rightarrow \mathbb{R}_{\geq 0}$ is an s - t -preflow in the network $(G = (V, E), \mu, s, t)$ if it satisfies:

- (i) (capacity constraints) $f(e) \leq \mu(e)$ for all arcs $e \in E$,
- (ii) (overflow) $\text{ex}_f(v) \geq 0$ for all vertices $v \in V \setminus \{s\}$.

We establish the key invariant of the BLOCKINGFLOW algorithm.

Lemma 3.8. After BLOCKINGFLOW performs a balancing step with vertex v_i , the algorithm maintains $\text{ex}_f(v_i) = 0$, unless $v_i = s$, and the fact that there is no v_i - t -path in $G_f \cap G$.

Proof. When starting a balancing step with v_i , there cannot be a vertex $v_j \neq t$ with $j \in \{i + 1, \dots, n - 1\}$ behind v_i in the topological order that overflows. By definition of a topological ordering, flow is only sent to predecessors of v_i in the topological order during the balancing step. This means that the balancing step maintains $\text{ex}_f(v_j) = 0$ for all $j \in \{i + 1, \dots, n - 1\}$. We say that all these vertices are in a balanced state. Since the incoming flow at each vertex is undone in reverse order of when it was created, none of the vertices v_j , $j \in \{i + 1, \dots, n\}$ will send flow back to v_i in future balancing steps (provided they do not receive additional flow from v_i). In addition, since v_i is marked after the balancing step, it will never receive new flow in future push steps. If $v_i \neq s$, then the balancing step continues until v_i is in a balanced state, therefore $\text{ex}_f(v_i) = 0$ is maintained from this time onward.

For the second part of the claim we use induction on the number of performed balancing steps. The claim trivially holds at the beginning. At the time when v_i is selected for the balancing step it overflows. This means that v_i can only have outgoing arcs to marked vertices in $G_f \cap G$ otherwise it could have pushed more flow. All these vertices were previously selected for balancing steps. This means that, by induction, there are no paths from v_i 's neighbors to t in $G_f \cap G$ now and in the future. As discussed above, the flow to successors of v_i in the topological order will never be undone and v_i will never receive new flow, which means that the set of neighbors of v_i in $G_f \cap G$ cannot grow. Thus, there will not be a v_i - t -path in $G_f \cap G$ from this time onward. \square

With this, we are ready to show correctness and running time of BLOCKINGFLOW.

Theorem 3.9. BLOCKINGFLOW computes a blocking flow in time $\mathcal{O}(n^2)$.

Proof. By Lemma 3.8, BLOCKINGFLOW computes a blocking flow: The lemma implies that no vertex is selected in more than one balancing step, and after all vertices have been selected once, at the latest, no vertex, except for s , has any overflow. Also, since s is always eventually selected in a balancing step, we have that there is no s - t -path in $G_f \cap G$ in the end, which means that f is a blocking flow.

For the running time, we first analyze the total number of PUSH operations. First observe that the number of PUSH operations in balancing steps is bounded by the number of PUSH operations in push steps. Therefore, it suffices to bound the number of PUSH operations in push steps. Since every vertex only pushes flow forward until some arc is not saturated by a push, and v_n does not perform any pushes, the number of unsaturating pushes is bounded by $n - 1$ in each iteration. For every saturating push we observe that the capacity along the arc only becomes available again in the future once a balancing step is performed with the other endpoint of the arc. Therefore this arc will not be used again in a push step. Hence, the total number of saturating pushes over all iterations is thus at most m . Overall, the

total number of PUSH operations can be bounded by $\mathcal{O}(n^2 + m) = \mathcal{O}(n^2)$. Our claim for the running time follows, since each PUSH operation can be performed in constant time. \square

Finally, using BLOCKINGFLOW as a subroutine for DINIC yields the following running time. Note the improvement over EDMONDSKARP for dense graphs, where we can expect many shortest augmenting paths in each phase of the algorithm.

Theorem 3.10. DINIC can be implemented with a running time of $\mathcal{O}(n^3)$.

Proof. We can compute the level graph in time $\mathcal{O}(m)$ using BFS. The statement therefore follows from Proposition 3.4 and Theorem 3.9. \square

3.2 Push-Relabel Algorithm

We now turn to a quite different approach to computing maximum flows. Intuitively, this approach is motivated by physical water networks, where flow is naturally maximized simply because water flows downhill whenever possible. With this in mind, we can try to obtain a maximum flow simply by emulating a system where the source vertex is at a higher elevation than the sink, and by simulating the water flow. The following definition serves as a mathematical basis for implementing this idea.

Definition 3.11. A function $h: V \rightarrow \mathbb{N} \cup \{0\}$ is a *height function* with respect to an s - t -preflow f if the following hold:

- (i) $h(s) = n$,
- (ii) $h(t) = 0$,
- (iii) $h(u) \leq h(v) + 1$ for all $(u, v) \in G_f$.

Note that we forbid arcs to go downhill by more than one unit of height. This will turn out to be crucial, in particular, it ensures that the following properties hold.

Lemma 3.12. Let h be a height function with respect to an s - t -preflow f . Then,

- (i) s is reachable in G_f from every vertex in $\{v \in V \mid \text{ex}_f(v) > 0\}$,
- (ii) for every u - v -path P in G_f it holds that $h(u) \leq h(v) + |P|$.

Proof. The second property holds by definition of height functions. For the first property, let $v' \in \{v \in V \mid \text{ex}_f(v) > 0\}$ be an overflowing vertex, and let R denote the set of all vertices reachable from v' in G_f . Then, $\mu_f(e) = 0$ for all $e \in \delta_{G \cup \bar{G}}^+(R)$ and, in particular, $f(e) = 0$ for all $e \in \delta_G^-(R)$. By definition of excess, we thus have

$$\begin{aligned}
 \sum_{v \in R} \text{ex}_f(v) &= \sum_{v \in R} \left(\sum_{(u,v) \in E} f(u,v) - \sum_{(v,w) \in E} f(v,w) \right) \\
 &= \sum_{v \in R} \left(\sum_{(u,v) \in E \setminus R^2} f(u,v) - \sum_{(v,w) \in E \setminus R^2} f(v,w) \right) \\
 &= \sum_{e \in \delta_G^-(R)} f(e) - \sum_{e \in \delta^+(R)} f(e) \\
 &= - \sum_{e \in \delta^+(R)} f(e) \leq 0.
 \end{aligned}$$

But since $\text{ex}_f(v') > 0$, there must be a vertex $u \in R$ with $\text{ex}_f(u) < 0$. By definition of a preflow, this can only be s . Hence s is reachable from v' in G_f , as claimed. \square

We can initially set up a preflow and height function by giving s a large height and by completely saturating all arcs in $\delta^+(s)$ (here and throughout we assume $\delta^-(s) = \delta^+(t) = \emptyset$, otherwise, we remove these arcs from the graph without changing the maximum flow value). At this point, the vertices in $\Gamma(s)$ have positive excess. Now the idea is to gradually increase the height of overflowing vertices and to simulate water flowing downhill whenever possible. The definition of a height function requires that we raise vertices one unit of height at a time, immediately simulate all water flow, and stop increasing the height once a vertex is not overflowing anymore. Ideally, we repeat this process until no vertex other than t has positive excess. We now show that if we reach this point, we are guaranteed to have computed a maximum flow!

Lemma 3.13. Let f be an s - t -flow in the network (G, μ, s, t) . If any height function with respect to f exists, then f is a maximum flow.

Proof. By Lemma 3.12, every s - t -path in G_f must contain at least $h(s) - h(t)$ arcs. But, by definition, we have $h(s) - h(t) = n$ and an s - t -path cannot consist of more than $n - 1$ arcs. Hence, no augmenting path can exist in G_f , which implies that f is a maximum s - t -flow (see lecture *Algorithmic Discrete Mathematics*). \square

We are now ready to formulate our approach as an algorithm. This algorithm is called the Goldberg-Tarjan algorithm, or the push-relabel algorithm.

Algorithm: GOLDBERG-TARJAN (G, μ, s, t)

input: network $(G = (V, E), \mu, s, t)$
output: maximum s - t -flow

$$f(u, v) \leftarrow \begin{cases} \mu(u, v), & \text{for } (u, v) \in \delta^+(s), \\ 0, & \text{for } (u, v) \in E \setminus \delta^+(s). \end{cases}$$

$$h(v) \leftarrow \begin{cases} n, & \text{for } v = s, \\ 0, & \text{for } v \in V \setminus \{s\}. \end{cases}$$

while $\exists v \in V \setminus \{t\}$ with $\text{ex}_f(v) > 0$:

if $\exists e = (v, w) \in \delta_{G_f}^+(v)$ with $h(v) = h(w) + 1$:	PUSH(v, e)	(push)
else	$h(v) \leftarrow \min_{w \in \Gamma_{G_f}(v)} \{h(w) + 1\}$ (relabel)	

return f

Algorithm: RELABEL-TO-FRONT (G, μ, s, t)

input: network $(G = (V, E), \mu, s, t)$
output: maximum s - t -flow

$$f(u, v) \leftarrow \begin{cases} \mu(u, v), & \text{for } (u, v) \in \delta^+(s), \\ 0, & \text{for } (u, v) \in E \setminus \delta^+(s). \end{cases}$$

$$h(v) \leftarrow \begin{cases} n, & \text{for } v = s, \\ 0, & \text{for } v \in V \setminus \{s\}. \end{cases}$$

while $\exists v \in \arg \max_{v \in V \setminus \{t\}} \{h(v) \mid \text{ex}_f(v) > 0\}$:

if $\exists e = (v, w) \in \delta_{G_f}^+(v)$ with $h(v) = h(w) + 1$:	PUSH(v, e)	(push)
else	$h(v) \leftarrow \min_{w \in \Gamma_{G_f}(v)} \{h(w) + 1\}$ (relabel)	

return f

We show that we never need to raise vertices above a height of $2n$ in order to eliminate all excess.

Lemma 3.14. GOLDBERG-TARJAN maintains a height function h with respect to the preflow f with $h(v) < 2n$ for all $v \in V$.

Proof. At the beginning of the algorithm we have $h(s) = n$ and $h(v) = 0$ for all $v \in V \setminus \{s\}$. Moreover, all arcs in $\delta^+(s)$ are saturated, hence G_f does not have any “downhill” arcs and h is a height function.

During a relabel step, h remains a height function by definition. Furthermore, we only apply relabeling steps to vertices v with $\text{ex}_f(v) > 0$. At this point, by Lemma 3.12 (i), there exists a v - s -path P in G_f . By Lemma 3.12 (ii) it further follows that $h(v) \leq h(s) + |P| \leq h(s) + n - 1 = 2n - 1$.

During a PUSH operation along arc $e = (v, w)$ the height function h is not changed, but the reverse arc (w, v) may be created in G_f . That is no problem, since $h(w) = h(v) - 1 < h(v) + 1$ ensures that h remains a height function with respect to G_f . \square

We have bounded the maximum height of all vertices and it remains to show that GOLDBERG-TARJAN keeps increasing heights, which then implies that the algorithm eventually terminates with f being an s - t -flow.

Lemma 3.15. In every relabel step GOLDBERG-TARJAN strictly increases the height of some vertex, and heights never decrease.

Proof. The height of a vertex v can only change during a relabel step, which is only applied once no push step is possible anymore. By Lemma 3.14, h is a height function, hence all arcs $(v, w) \in \delta_{G_f}^+(v)$ satisfy $h(w) \geq h(v)$ before a relabel step for v . By Lemma 3.12, $\text{ex}_f(v) > 0$ implies that a v - s -path exists in G_f , and, in particular, $\delta_{G_f}^+(v) \neq \emptyset$. Thus, the height of v increases during the relabel step. \square

From Lemmas 3.13 and 3.14 it follows that GOLDBERG-TARJAN is correct. We proceed to bound the number of operations needed by the algorithm to obtain a bound on its running time.

Proposition 3.16. GOLDBERG-TARJAN computes a maximum s - t -flow using $\mathcal{O}(n^2m)$ PUSH and $\mathcal{O}(n^2)$ relabel operations.

Proof. By Lemma 3.14, the algorithm maintains a height function. When GOLDBERG-TARJAN terminates, no vertex except t may overflow, thus f is an s - t -flow. By Lemma 3.13, f thus is a maximum flow.

We analyze the number of relabel steps. By Lemma 3.15, we increase the height by at least one during each relabel step. By Lemma 3.14, the number of relabel steps is thus bounded by $\mathcal{O}(n^2)$.

Now consider the number of push steps in which the corresponding arc (v, w) is saturated. Since the heights of vertices can never decrease (Lemma 3.15), the height of w must increase by at least two before a push step can be performed along the reverse arc, after which a push step with (v, w) can occur again. Since the height is bounded by $2n - 1$ (Lemma 3.14), the number of saturating push steps along the same arc is thus bounded by $\mathcal{O}(n)$. The total number of saturating push steps thus is in $\mathcal{O}(mn)$.

To bound the number of non-saturating push steps, we consider the potential $\Phi = \sum_{v: \text{ex}_f(v) > 0} h(v)$. During a relabel step, the potential can increase by at most $2n - 1$ (Lemma 3.14). During each push step at most one new overflowing vertex is created. During a saturating push step this may increase the potential again by at most $2n - 1$ (Lemma 3.14). During a non-saturating push step along arc (v, w) , however, v loses its excess completely. Since $h(v) = h(w) + 1$, the potential therefore decreases by one (at least – since no new overflowing vertex may have been created). We already know that the number of relabel and saturating push steps is $\mathcal{O}(n^2 + mn) = \mathcal{O}(mn)$. During these steps the potential increases by $\mathcal{O}(n)$, hence the potential can only decrease at most $\mathcal{O}(n^2m)$ times overall (every decrease is by at least one, we initially have $\Phi = 0$, and $\Phi \geq 0$ throughout). Hence, the total number of non-saturating push steps is at most $\mathcal{O}(n^2m)$. \square

Proposition 3.16 suggests that GOLDBERG-TARJAN does not perform better than DINIC. However, we can further improve its running time by choosing the vertex more carefully that we consider in each iteration of the algorithm. The version of GOLDBERG-TARJAN that always selects an overflowing vertex of largest height is often called relabel-to-front algorithm. We show that this algorithm requires significantly fewer operations.

Proposition 3.17. RELABELTOFRONT only requires $\mathcal{O}(n^2\sqrt{m})$ push steps.

Proof. Recall from the proof of Proposition 3.16 that the number of saturating push steps is bounded by $\mathcal{O}(mn)$, hence it suffices to bound the number of non-saturating push steps.

Let $h_{\max} := \max_{v \in \{v' \in V \mid \text{ex}_f(v') > 0\}} h(v)$ be the largest height among all overflowing vertices. We divide the execution of the algorithm into maximal phases in such a way that h_{\max} remains constant during each phase. Since h can only change during a relabel step and can only increase (Lemma 3.15), and since $h < 2n$ (Lemma 3.14), the total increase of $\sum_{v \in V} h(v)$ and thus of h_{\max} is at most $2n^2$. Since $h_{\max} \geq 0$ and initially $h_{\max} = 0$, we obtain that h_{\max} can only decrease at most $2n^2$ times (each time by at least one). Hence, the number of phases is at most $4n^2 = \mathcal{O}(n^2)$.

First consider “cheap” phases with at most \sqrt{m} non-saturating push steps. The total number of non-saturating push steps in all these phases together is $\mathcal{O}(n^2\sqrt{m})$. It remains to consider “expensive” phases.

We define the potential

$$\Phi = \sum_{v \in \{v' \in V \mid \text{ex}_f(v') > 0\}} \Phi(v) := \sum_{v \in \{v' \in V \mid \text{ex}_f(v') > 0\}} |\{w \in V \mid h(w) \leq h(v)\}|.$$

Since $\Phi(v) \leq n$ holds for all $v \in V$, we initially have $\Phi \leq n^2$ and every relabel or saturating push step can increase Φ by at most $n - 1$. A non-saturating push step reduces Φ by at least one. Since the number of relabel and saturating push steps is at most $\mathcal{O}(mn)$ (Proposition 3.16) and each such step can increase Φ by at most n , the amount by which Φ can be reduced over all non-saturating push steps is at most $\mathcal{O}(n^2m)$.

Now consider an expensive phase and only the vertices of height h_{\max} during that phase. Note that heights remain unchanged during a phase: Only vertices with excess and of maximum height are candidates for relabel steps. Thus, once a relabel step occurs (or when no vertex other than t has excess) the phase ends. In particular, the set $V_{\max} := \{v \in V \mid h(v) = h_{\max}\}$ remains unchanged during the phase. With every non-saturating push step a vertex loses its excess. Since there are more than \sqrt{m} such steps in an expensive phase, we have $|V_{\max}| > \sqrt{m}$. It follows that $\Phi(v) > \sqrt{m}$ for all $v \in V_{\max}$. Every time when a vertex loses its excess, Φ thus decreases by more than \sqrt{m} (note that a new vertex u with excess may be created, but for this vertex we have $\Phi(v) - \Phi(u) > \sqrt{m}$, since $h(u) < h(v)$). Since Φ can only be reduced by $\mathcal{O}(n^2m)$ overall, the total number of non-saturating push steps in expensive phases is bounded by $\mathcal{O}(n^2\sqrt{m})$. \square

In order to implement RELABELTOFRONT efficiently, we need to ensure that each iteration of the algorithm needs constant time. We can achieve this by explicitly maintaining $h_{\max} := \max\{h(v) \mid v \in V \wedge \text{ex}_f(v) > 0\}$ as well as the lists $L(h) = \{v \in V \setminus \{t\} \mid h(v) = h \wedge \text{ex}_f(v) > 0\}$ for each height $h \in \{0, \dots, 2n - 1\}$. Additionally, we store lists $E_v = \{(v, w) \in \delta_{G_f}^+(v) \mid h(v) = h(w) + 1\}$ for every vertex $v \in V$ that contain the arcs in $\delta_{G_f}^+(v)$ going downhill. Note that we only ever need to access or modify the first or last elements of these lists, which can be accomplished in constant time. The following gives an efficient implementation using these lists.

Algorithm: RELABELTOFRONT (implementation)

input: network $(G = (V, E), \mu, s, t)$

output: maximum s - t -flow

$$f(u, v) \leftarrow \begin{cases} \mu(u, v), & \text{for } (u, v) \in \delta^+(s), \\ 0, & \text{for } (u, v) \in E \setminus \delta^+(s). \end{cases}$$

$$h(v) \leftarrow \begin{cases} n, & \text{for } v = s, \\ 0, & \text{for } v \in V \setminus \{s\}. \end{cases}$$

$$L(h) \leftarrow \emptyset \text{ for all } h \in \{0, \dots, 2n-1\}$$

$$h_{\max} \leftarrow 0, L(h_{\max}) \leftarrow \Gamma(s) \setminus \{t\}$$

$$E_v \leftarrow \emptyset \text{ for all } v \in V$$

while $\exists v \in L(h_{\max})$: (take first element)

if $\exists e \in E_v$: (take first element)

 | PUSH(v, e)

else

 | RELABEL(v)

return f

Algorithm: PUSH($v, e \equiv (v, w)$) (impl.)

if $\text{ex}_f(w) = 0$ **and** $w \notin \{s, t\}$:

 | $L(h(w)) \leftarrow L(h(w)) \oplus (w)$

AUGMENT($e, \min\{\text{ex}_f(v), \mu_f(e)\}$)

if $\text{ex}_f(v) = 0$:

 | $L(h_{\max}) \leftarrow L(h_{\max}) \setminus \{v\}$ (remove first)

while $h_{\max} > 0$ **and** $L(h_{\max}) = \emptyset$:

 | $h_{\max} \leftarrow h_{\max} - 1$

if $\mu_f(e) = 0$:

 | $E_v \leftarrow E_v \setminus \{e\}$ (remove first)

Algorithm: RELABEL(v) (implementation)

$L(h(v)) \leftarrow L(h(v)) \setminus \{v\}$ (remove first)

$h(v) \leftarrow \min\{h(w) + 1 \mid w \in \Gamma_{G_f}(v)\}$

$E_v \leftarrow \{(v, w) \in \delta_{G_f}^+(v) \mid h(v) = h(w) + 1\}$

$h_{\max} \leftarrow h(v)$

$L(h_{\max}) \leftarrow \{v\}$

We analyze the running times of this implementation.

Observation 3.18. The running time for a RELABEL operation for vertex v is $\Theta(\deg_G(v))$.

Lemma 3.19. The running time for k PUSH operations is $\mathcal{O}(k + n^2)$.

Proof. Except for updating h_{\max} , the running time of a PUSH operation is constant. Since h is never reduced (Lemma 3.15) and $h < 2n$ (Lemma 3.14), the total increase of $\sum_{v \in V} h(v)$ is bounded by $\mathcal{O}(n^2)$. Thus, the total decrease of h_{\max} is also bounded by $\mathcal{O}(n^2)$, which proves the claim. \square

With this, we can bound the running time of RELABELTOFRONT. Observe that the running time is never worse than that of DINIC and better on sparse graphs.

Theorem 3.20. RELABELTOFRONT can be implemented with a running time of $\mathcal{O}(n^2 \sqrt{m})$.

Proof. As shown above (Proposition 3.17), the number of PUSH operations is $\mathcal{O}(n^2 \sqrt{m})$. By Lemma 3.19, the running time for all these operations together is thus $\mathcal{O}(n^2 \sqrt{m} + n^2) = \mathcal{O}(n^2 \sqrt{m})$. We have $\mathcal{O}(n)$ RELABEL operations for each vertex (Proposition 3.16), with a total running time of $\mathcal{O}(mn) = \mathcal{O}(n^2 \sqrt{m})$ (Observation 3.18). \square



4 Minimum-Cost Flows

In the previous chapter, we focused on the maximum flow problem as a framework for algorithmic problems where a throughput through a network needs to be maximized. In practice, we are often not only interested in maximizing throughput, but also on simultaneously minimizing the cost of our solution with respect to some cost function. For example, in a logistics network, we are not only interested in maximizing the amount of goods that we can deliver, but we also want to minimize (among other factors) the total distance we need to travel. To account for this, we now consider the problem of finding a maximum flow that has minimum cost among all maximum flows. Since we want to restrict ourselves to maximum flows, it will be useful to introduce a notion of networks with prescribed demands associated with all vertices. In this setting, we implicitly fix the total desired flow value, instead of having source and sink vertices with unbounded demands and maximizing the flow.

Definition 4.1. A *circulation network* is a triple (G, μ, b) where $G = (V, E)$ is a directed graph, $\mu: E \rightarrow \mathbb{R}_{\geq 0}$ are capacities, and $b: V \rightarrow \mathbb{R}$ are demands satisfying $\sum_{v \in V} b(v) = 0$. As for networks, to simplify notation, we additionally require that $\{(u, v), (v, u)\} \not\subseteq E$ for all $u, v \in V$.

We need to adapt our definition of a flow accordingly. Note that we no longer need flow conservation.

Definition 4.2. A function $f: E \rightarrow \mathbb{R}_{\geq 0}$ is a *b-flow* in a circulation network (G, μ, b) if $f(e) \leq \mu(e)$ for all $e \in E$ and $\text{ex}_f(v) = -b(v)$ for all $v \in V$.

Note that while the values $b(v)$ are generally referred to as *demands*, they might better be thought of as negative demands or supplies.

The following observation summarizes how network flows and *b*-flows in circulation networks relate. In particular, observe that a circulation network fixes the flow value of the *s-t*-flow in the corresponding network. This allows us to easily restrict ourselves to maximum flows (e.g., by setting $b(s) = |f^*| = -b(t)$) to transition from a network with maximum flow value $|f^*|$ to a circulation network with prescribed flow value $|f^*|$.

Observation 4.3. Let (G, μ, b) be a circulation network with $G = (V, E)$ and (G', μ', s, t) be a network with $G' = (V \cup \{s, t\}, E \cup (\{s\} \times V) \cup (V \times \{t\}))$. Let further $\mu'(s, v) = \max\{0, b(v)\}$ and $\mu'(v, t) = \max\{0, -b(v)\}$ for all $v \in V$, and $\mu'|_E = \mu$. Then, there is a *b*-flow f in (G, μ, b) if and only if there is an *s-t*-flow f' in (G', μ', s, t) with $|f'| = \frac{1}{2} \sum_{v \in V} |b(v)|$, and we may choose f and f' such that $f'|_E = f$.

This observation allows us to formulate the minimum-cost flow problem without explicitly demanding maximum flows:

Minimum-Cost Flow Problem

input: circulation network $(G = (V, E), \mu, b)$; arc costs $c: E \rightarrow \mathbb{R}$
problem: find *b*-flow f of minimum cost $c(f) := \sum_{e \in E} f(e) \cdot c(e)$

For convenience, we consider an arbitrary but fixed circulation network $(G = (V, E), \mu, b)$ with fixed costs $c: E \rightarrow \mathbb{R}$ throughout this chapter. We will often operate on residual graphs of G and we set $c(\bar{e}) := -c(e)$ for the cost of reverse arcs $\bar{e} \in \bar{E}$.

We begin by reformulating the max-flow min-cut theorem (see lecture *Algorithmic Discrete Mathematics*) for *b*-flows.

Proposition 4.4. A b -flow exists if and only if for all $X \subseteq V$, it holds that

$$\sum_{e \in \delta^+(X)} \mu(e) \geq \sum_{v \in X} b(v).$$

Proof. Let G' be the graph of Observation (4.3). By the max-flow min-cut theorem, there is an s - t -flow of value (at least) $\frac{1}{2} \sum_{v \in V} |b(v)|$ in (G', μ', s, t) , and therefore a b -flow in (G, μ, b) , if and only if

$$\sum_{e \in \delta_{G', \{s\} \cup X}^+} \mu(e) \geq \frac{1}{2} \sum_{v \in V} |b(v)| \quad (4.1)$$

for all $X \subseteq V$. In addition, since $\sum_{v \in V} b(v) = 0$ by definition,

$$\begin{aligned} \sum_{e \in \delta_{G', \{s\} \cup X}^+} \mu'(e) &= \sum_{e \in \delta_G^+(X)} \mu(e) + \sum_{v \in V \setminus X} \max\{0, b(v)\} + \sum_{v \in X} \max\{0, -b(v)\} \\ &= \sum_{e \in \delta_G^+(X)} \mu(e) + \sum_{v \in V} \max\{0, b(v)\} - \sum_{v \in X} \max\{0, b(v)\} + \sum_{v \in X} \max\{0, -b(v)\} \\ &= \sum_{e \in \delta_G^+(X)} \mu(e) + \frac{1}{2} \sum_{v \in V} |b(v)| + \sum_{v \in X} (\max\{0, -b(v)\} - \max\{0, b(v)\}) \\ &= \sum_{e \in \delta_G^+(X)} \mu(e) + \frac{1}{2} \sum_{v \in V} |b(v)| - \sum_{v \in X} b(v). \end{aligned}$$

Together with (4.1) this concludes the proof. \square

As discussed above, finding a valid b -flow is equivalent to finding a maximum network flow. This can be accomplished as follows:

Algorithm: BFLOW(G, μ, b)

input: circulation network $(G = (V, E), \mu, b)$

output: b -flow if one exists

$G' \leftarrow (V \cup \{s, t\}, E \cup (\{s\} \times V) \cup (V \times \{t\}))$

$$\mu'(u, v) \leftarrow \begin{cases} \max\{0, b(v)\}, & \text{for } u = s, v \in V \\ \max\{0, -b(u)\}, & \text{for } u \in V, v = t, \\ \mu(u, v), & \text{for } u \in V, v \in V. \end{cases}$$

$f \leftarrow \text{RELABELTOFRONT}(G', \mu', s, t)$

if $2|f| = \sum_{v \in V} |b(v)|$:

return $f|_E$

else

fail

Proposition 4.5. If a b -flow exists, BFLOW finds one in time $\mathcal{O}(n^2 \sqrt{m})$.

Proof. This follows from Observation 4.3 and the running time of RELABELTOFRONT (Theorem 3.20). \square

Intuitively, every pair of b -flows differ by a set of cycles. The following notion will be important to formalize this intuition.

Definition 4.6. A *circulation* is a function that is a b -flow in the circulation network (G, μ, b) with $b = 0$.

We move on by defining the difference between two b -flows.

Definition 4.7. For two b -flows f and f' we define $f' \ominus f := g$ with $g(e) = \max\{0, f'(e) - f(e)\}$ and $g(\bar{e}) = \max\{0, f(e) - f'(e)\}$ for all $e \in E$.

We can now formally establish that b -flows always differ by a set of cycles.

Lemma 4.8. For two b -flows f and f' it holds that $g = f' \ominus f$ is a circulation in $G \cup \bar{G}$ (!) with $g(e) = 0$ for all $e \in (E \cup \bar{E}) \setminus G_f$ and with $c(g) = c(f') - c(f)$.

Proof. For every vertex $v \in V$ of $G \cup \bar{G}$ it holds that

$$\begin{aligned}
 -\text{ex}_g(v) &= \sum_{e \in \delta_{G \cup \bar{G}}^+(v)} g(e) - \sum_{e \in \delta_{G \cup \bar{G}}^-(v)} g(e) \\
 &= \sum_{e \in \delta_G^+(v)} g(e) - \sum_{e \in \delta_{\bar{G}}^-(v)} g(e) - \sum_{e \in \delta_{\bar{G}}^-(v)} g(e) + \sum_{e \in \delta_{\bar{G}}^+(v)} g(e) \\
 &= \sum_{e \in \delta_G^+(v)} (g(e) - g(\bar{e})) - \sum_{e \in \delta_{\bar{G}}^-(v)} (g(e) - g(\bar{e})) \\
 &= \sum_{e \in \delta_G^+(v)} (f'(e) - f(e)) - \sum_{e \in \delta_{\bar{G}}^-(v)} (f'(e) - f(e)) \\
 &= \left(\sum_{e \in \delta_G^+(v)} f'(e) - \sum_{e \in \delta_{\bar{G}}^-(v)} f'(e) \right) - \left(\sum_{e \in \delta_G^+(v)} f(e) - \sum_{e \in \delta_{\bar{G}}^-(v)} f(e) \right) \\
 &= b(v) - b(v) \\
 &= 0,
 \end{aligned}$$

hence, g is a circulation in $G \cup \bar{G}$.

For $e \in E \setminus G_f$ we have $f(e) = \mu(e)$ and thus $f'(e) \leq f(e)$, which implies $g(e) = 0$. For $\bar{e} \in \bar{E} \setminus G_f$ we have $f(e) = 0$ and thus $g(\bar{e}) = 0$.

Lastly, in $G \cup \bar{G}$, it holds that

$$\begin{aligned}
 c(g) &= \sum_{e \in E \cup \bar{E}} c(e) \cdot g(e) \\
 &= \sum_{e \in E} (c(e) \cdot g(e) + c(\bar{e}) \cdot g(\bar{e})) \\
 &= \sum_{e \in E} c(e) \cdot (g(e) - g(\bar{e})) \\
 &= \sum_{e \in E} c(e) \cdot (f'(e) - f(e)) \\
 &= c(f') - c(f).
 \end{aligned}$$

□

Finally, we can translate the fact that flows can be decomposed into paths and cycles (see lecture *Algorithmic Discrete Mathematics*) to circulations.

Corollary 4.9. For every circulation g there exists a family of cycles \mathcal{K} with $|\mathcal{K}| \leq |E|$ and weights $w: \mathcal{K} \rightarrow \mathbb{R}_{>0}$ with $g(e) = \sum_{K \in \mathcal{K}: e \in K} w(K)$ for all $e \in E$.

Proof. This follows immediately by existence of a path decomposition for s - t -flows, since g is an s - t -flow (of value 0) for arbitrary $s, t \in V$. \square

We are now ready to state the first fundamental characterization of minimum-cost b -flows. Note how negative(-cost) cycles play a similar role as augmenting paths for network flows.

Theorem 4.10. A b -flow f has minimum cost if and only if G_f does not contain a negative-cost cycle.

Proof. If there is a negative-cost cycle, we can increase the flow along this cycle and obtain a b -flow of lower cost, hence f cannot have minimum cost.

On the other hand, let f' be a b -flow with smaller cost than f . We consider $g = f' \ominus f$. By Lemma 4.8, g is a circulation in $G \cup \bar{G}$ with $c(g) = c(f') - c(f) < 0$. By Corollary 4.9, we can decompose g into cycles \mathcal{K} with weights $w: \mathcal{K} \rightarrow \mathbb{R}_{>0}$, such that $g(e) = \sum_{K \in \mathcal{K}: e \in K} w(K)$. From $g(e) = 0$ for all $e \in (E \cup \bar{E}) \setminus G_f$ (Lemma 4.8) and $w > 0$ it follows that every cycle in \mathcal{K} is a subgraph of G_f . We obtain

$$\begin{aligned} c(g) &= \sum_{e \in E \cup \bar{E}} c(e) \cdot g(e) \\ &= \sum_{e \in E \cup \bar{E}} c(e) \cdot \sum_{K \in \mathcal{K}: e \in K} w(K) \\ &= \sum_{K \in \mathcal{K}} w(K) \cdot \sum_{e \in K} c(e) \\ &= \sum_{K \in \mathcal{K}} w(K) \cdot c(K). \end{aligned}$$

From $c(g) < 0$ and $w > 0$ it follows that \mathcal{K} must contain a negative-cost cycle. \square

4.1 Minimum-Mean-Cycle Cancellation

Theorem 4.10 immediately suggests an algorithm that repeatedly augments flow along negative cycles until all such cycles are eliminated, very similarly to the FORDFULKERSON method. As with the FORDFULKERSON method, it turns out that it is important how we select negative cycles in each step. Much like EDMONDSKARP selects a shortest augmenting path in each iteration, it turns out that selecting short (i.e., very negative) cycles first is beneficial. However, we additionally need to account for the number of arcs in the cycle. We do this by selecting the cycle first whose arcs have minimum cost on average. More precisely, we use the following notion to compare cycles.

Definition 4.11. Let f be a b -flow and let \mathcal{K}_f be the set of all cycles in G_f . A cycle K in G_f is a *minimum-mean cycle* with respect to c if $c(K)/|K| = \text{mc}(f) := \min_{K \in \mathcal{K}_f} c(K)/|K|$.

Assuming we can find minimum-mean cycles, we can compute a minimum-cost b -flow as follows.

Algorithm: MINMEANCYCLECANCELLING(G, μ, b, c)

input: circulation network (G, μ, b) , costs $c: E \rightarrow \mathbb{R}$

output: minimum cost b -flow

$f \leftarrow \text{BFLOW}(G, \mu, b)$

$K \leftarrow \text{MINMEANCYCLE}(G_f, c)$

while $K \neq \emptyset$ **and** $c(K) < 0$:

$\gamma \leftarrow \min_{e \in K} \mu_f(e)$

AUGMENT(K, γ)

$K \leftarrow \text{MINMEANCYCLE}(G_f, c)$

return f

To prove that MINMEANCYCLECANCELLING always terminates and to estimate its running time, we proceed analogously to the analysis of EDMONDSKARP. We first establish that the algorithm makes progress in terms of the value of $\text{mc}(f)$.

Lemma 4.12. Let f_i denote the flow and K_i the chosen cycle at the start of each iteration i of MINMEANCYCLECANCELLING. For any two such cycles K_j, K_ℓ with $j < \ell$, we have

$$|\text{mc}(f_j)| \geq \frac{n}{n - 2 \min\{1, m_{j\ell}\}} \cdot |\text{mc}(f_\ell)|,$$

where $m_{j\ell} := |\{e \in E \mid e, \bar{e} \in K_j \cup K_\ell\}|$.

Proof. We may assume that there is no cycle K_i , $j < i < \ell$ that contains an arc e with $\bar{e} \in K_\ell$. Otherwise, by induction on $\ell - j$, we immediately have $|\text{mc}(f_j)| \geq |\text{mc}(f_i)| \geq \frac{n}{n-2} |\text{mc}(f_\ell)| \geq \frac{n}{n-2 \min\{1, m_{j\ell}\}} |\text{mc}(f_\ell)|$.

Let $\bar{E}_{j\ell} := \{e \in E \cup \bar{E} \mid e, \bar{e} \in K_j \cup K_\ell\}$. We consider the graphs $H_j = (V, E_j) := (V, K_j \setminus \bar{E}_{j\ell})$, $H_\ell = (V, E_\ell) := (V, K_\ell \setminus \bar{E}_{j\ell})$, and $H_{j\ell} = (V, E_{j\ell}) := (V, (K_j \cup K_\ell) \setminus \bar{E}_{j\ell})$. Since the cycles $K_{j+1}, \dots, K_{\ell-1}$ do not contain any reverse arcs to arcs in K_ℓ , we have $K_\ell \setminus \bar{E}_{j\ell} \subseteq G_{f_j}$ and hence $H_j, H_\ell, H_{j\ell} \subseteq G_{f_j}$. Moreover, we have $|\delta_{H_j}^+(v)| + |\delta_{H_\ell}^+(v)| = |\delta_{H_j}^-(v)| + |\delta_{H_\ell}^-(v)|$ for all $v \in V$, which means that $g: E_{j\ell} \rightarrow \{1, 2\}$ with $g(e) = |\{H \in \{H_j, H_\ell\} \mid e \in H\}|$ defines a circulation in $H_{j\ell}$. By Corollary 4.9 we can decompose g into cycles, each of which are subgraphs of $H_{j\ell} \subseteq G_{f_j}$ and thus have average cost at least $\text{mc}(f_j)$, by definition. Since g encodes all arcs of H_j and H_ℓ , we obtain

$$c(E_j) + c(E_\ell) \geq \text{mc}(f_j) \cdot (|E_j| + |E_\ell|). \quad (4.2)$$

Since the costs of opposing arcs cancel out, we further have

$$c(E_j) + c(E_\ell) = c(K_j) + c(K_\ell) = \text{mc}(f_j) |K_j| + \text{mc}(f_\ell) |K_\ell|. \quad (4.3)$$

Finally, it holds that

$$\begin{aligned} |E_j| + |E_\ell| &= |K_j| + |K_\ell| - 2m_{j\ell} \\ &\leq |K_j| + |K_\ell| - \frac{|K_\ell|}{n} 2m_{j\ell} \\ &= |K_j| + \frac{n - 2m_{j\ell}}{n} |K_\ell|. \end{aligned} \quad (4.4)$$

Together with $\text{mc}(f_j) < 0$, this implies

$$\begin{aligned} \text{mc}(f_j) \left(|K_j| + \frac{n - 2m_{j\ell}}{n} |K_\ell| \right) &\stackrel{(4.4)}{\leq} \text{mc}(f_j) (|E_j| + |E_\ell|) \\ &\stackrel{(4.2)}{\leq} c(E_j) + c(E_\ell) \\ &\stackrel{(4.3)}{=} \text{mc}(f_j) |K_j| + \text{mc}(f_\ell) |K_\ell|, \end{aligned}$$

and thus $\text{mc}(f_j) \leq \frac{n}{n-2m_{j\ell}} \cdot \text{mc}(f_\ell)$. This completes the proof, since $\text{mc}(f_j), \text{mc}(f_\ell) < 0$. \square

We can conclude that the value of $\text{mc}(f)$ must approach 0 (from below) significantly in each batch of $n \cdot m$ iterations.

Lemma 4.13. `MINMEANCYCLECANCELLING` reduces $|\text{mc}(f)|$ at least by a factor of 2 in $n \cdot m$ consecutive iterations (if it does not terminate).

Proof. Let K_i, \dots, K_{i+m} be the augmenting cycles chosen in $m+1$ consecutive iterations, and let f_i, \dots, f_{i+m} be the b -flows at the start of the corresponding iterations. Since every augmentation removes an arc of the residual graph, there are two cycles K_j, K_ℓ , $i \leq j < \ell \leq i+m$ who contain opposing arcs. According to Lemma 4.12, we have

$$\text{mc}(f_i) \leq \text{mc}(f_j) \leq \frac{n}{n-2} \text{mc}(f_\ell) \leq \frac{n}{n-2} \text{mc}(f_{i+m}).$$

Thus, after $n \cdot m$ iterations, the reduction of $|\text{mc}(f)|$ is at least $\left(\frac{n}{n-2}\right)^n = \left(\left(\frac{x}{x-1}\right)^x\right)^2 > e^2 \approx (2.72)^2 > 2$, where $x := n/2$. \square

From this, it is easy to infer a bound on the overall number of iterations. Note, however, that the following bound is not strongly polynomial in n and m , since it depends on the number of bits (and not only the number of numbers) needed to encode the input.

Corollary 4.14. If $c(e) \in \mathbb{Z}$ for all $e \in E$, `MINMEANCYCLECANCELLING` computes a min-cost b -flow in $\mathcal{O}(mn(\log n + \log |c_{\min}|))$ iterations, where $c_{\min} := \min_{e \in E} c(e)$.

Proof. At the beginning, we have $|\text{mc}(f)| \leq |c_{\min}|$. By Lemma (4.13), the algorithm needs $\mathcal{O}(nm \log |c_{\min}|)$ iterations until $\text{mc}(f) > -1$, and, hence, $\mathcal{O}(nm(\log n + \log |c_{\min}|))$ iterations until $\text{mc}(f) > -1/n$. Since all costs in the residual graph remain integral, and augmenting cycles have at most n arcs, we have $\text{mc}(f) \notin (-1/n, 0)$. Consequently, we have $\text{mc}(f) \geq 0$ after $\mathcal{O}(nm(\log n + \log |c_{\min}|))$ iterations and the algorithm terminates. By Theorem 4.10, the resulting b -flow is optimal. \square

We now give a more careful estimation that yields a polynomial bound.

Theorem 4.15. `MINMEANCYCLECANCELLING` computes a min-cost b -flow in $\mathcal{O}(m^2 n \log n)$ iterations.

Proof. Correctness follows from Theorem (4.10). We show that after each additional $mn(\lceil \log_2 n \rceil + 1)$ iterations an additional arc will remain unchanged in future iterations. Let f, f' be the flows before and after such a sequence of iterations, and let K_f be the augmenting cycle chosen for f . The proof proceeds as follows:

- (i) We show that the costs of arcs in $G_{f'}$ are bounded from below.

- (ii) We show that there is a very cheap arc $e_0 \in G_f \setminus G_{f'}$, so that \bar{e}_0 is very expensive in $G_{f'}$.
- (iii) We claim that the residual capacity of e_0 stays 0 in future iterations of the algorithm.
- (iv) By Lemma 4.12, to show this it suffices to show that, for all b -flows f'' with $e_0 \in G_{f''}$ it holds that $\text{mc}(f'') < \text{mc}(f')$.
- (v) We define $g := f'' \ominus f'$ and consider the decomposition of g into a set of cycles \mathcal{K} .
- (vi) We show there must be a cycle $K_0 \in \mathcal{K}$ with $\bar{e}_0 \in K_0 \subseteq G_{f'}$ and, because of 1. and 2., that K_0 is very expensive.
- (vii) It follows that $\bar{K}_0 = \{\bar{e} \mid e \in K_0\} \subseteq G_{f''}$ is too cheap in $G_{f''}$ (i.e., $\text{mc}(f'') < \text{mc}(f')$).

We now address these steps in detail.

- (i) By Lemma 4.13 we have

$$\text{mc}(f) \leq 2^{\lceil \log_2 n \rceil + 1} \text{mc}(f') \leq 2n \text{mc}(f'), \quad (4.5)$$

since $\text{mc}(f), \text{mc}(f') < 0$.

We define $c'(e) = c(e) + |\text{mc}(f')| = c(e) - \text{mc}(f')$ for all $e \in G_{f'}$. With this, c' does not induce any negative cycles in $G_{f'}$. By Proposition 2.24 there exists a potential $\pi: V \rightarrow \mathbb{R}$ with $0 \leq c'_\pi(e) = c_\pi(e) - \text{mc}(f') - \text{recall that } c_\pi(u, v) := \pi(u) - \pi(v) + c(u, v)$. Hence,

$$c_\pi(e) \geq \text{mc}(f'), \quad (4.6)$$

for all $e \in G_{f'}$.

- (ii) We obtain

$$\sum_{e \in K_f} c_\pi(e) = \sum_{e \in K_f} c(e) = \text{mc}(f) \cdot |K_f| \stackrel{(4.5)}{\leq} 2n \text{mc}(f') |K_f|.$$

This means we can choose $e_0 \in K_f \subseteq G_f$ with

$$c_\pi(e_0) \leq 2n \text{mc}(f') < \text{mc}(f'). \quad (4.7)$$

By (4.6) we have $e_0 \notin G_{f'}$.

- (iii) We claim that the flow along e_0 remains unchanged in future iterations.
- (iv) Let f'' be a b -flow with $e_0 \in G_{f''}$. We show that $\text{mc}(f'') < \text{mc}(f')$. Since $\text{mc}(f)$ is growing monotonically during the course of the algorithm (Lemma 4.12), this proves 3.
- (v) Let $g := f'' \ominus f'$. We show that $g(e) = \max\{0, \mu_{f''}(\bar{e}) - \mu_{f'}(\bar{e})\}$ for all $e \in E \cup \bar{E}$: For $e \in E$ we have $g(e) = \max\{0, f''(e) - f'(e)\} = \max\{0, \mu_{f''}(\bar{e}) - \mu_{f'}(\bar{e})\}$. For $e \in \bar{E}$ we have $g(e) = \max\{0, f'(\bar{e}) - f''(\bar{e})\} = \max\{0, (\mu(\bar{e}) - \mu_{f'}(\bar{e})) - (\mu(\bar{e}) - \mu_{f''}(\bar{e}))\} = \max\{0, \mu_{f''}(\bar{e}) - \mu_{f'}(\bar{e})\}$. In particular, it follows from $e_0 \in G_{f''} \setminus G_{f'}$ that $\mu_{f'}(e_0) = 0 < \mu_{f''}(e_0)$, hence $g(\bar{e}_0) = \mu_{f''}(e_0) > 0$. By Lemma 4.8 and Corollary 4.9, we may consider the decomposition of g into a set of cycles \mathcal{K} in $G \cup \bar{G}$.
- (vi) Let $K_0 \in \mathcal{K}$ with $\bar{e}_0 \in K_0$. By Lemma 4.8, we have $g(e) = 0$ for $e \in (E \cup \bar{E}) \setminus G_{f'}$, which implies $K_0 \subseteq G_{f'}$. We obtain (using $\text{mc}(f') < 0$, otherwise, the statement is trivial)

$$\begin{aligned} c(K_0) &= \sum_{e \in K_0} c_\pi(e) \\ &= c_\pi(\bar{e}_0) + \sum_{e \in K_0 \setminus \{\bar{e}_0\}} c_\pi(e) \\ &\stackrel{(4.6)}{\geq} -c_\pi(e_0) + \sum_{e \in K_0 \setminus \{\bar{e}_0\}} \text{mc}(f') \\ &\stackrel{(4.7)}{\geq} -2n \cdot \text{mc}(f') + (n-1) \cdot \text{mc}(f') \\ &> -n \cdot \text{mc}(f'). \end{aligned}$$

(vii) Let \bar{K}_0 be the cycle consisting of the opposing arcs to the arcs in K_0 . For all $e \in K_0$ it holds that $0 < g(e) = \max\{0, \mu_{f''}(\bar{e}) - \mu_{f'}(\bar{e})\}$, hence $\mu_{f''}(\bar{e}) > 0$. It follows that $\bar{K}_0 \subseteq G_{f''}$. But $c(\bar{K}_0) = -c(K_0) < n \cdot \text{mc}(f')$. This completes the proof (by (iv)), since

$$\text{mc}(f'') \leq \frac{c(\bar{K}_0)}{|\bar{K}_0|} < \frac{n \text{mc}(f')}{|\bar{K}_0|} \stackrel{\text{mc}(f') < 0}{\leq} \text{mc}(f'). \quad \square$$

4.1.1 Finding a minimum-mean cycle

For a bound on the overall running time of `MINMEANCYCLECANCELLING` we first need to describe how to find a minimum-mean cycle in each iteration. In other words, we consider the following problem.

Minimum-Mean Cycle Problem

input: graph $G = (V, E)$; arc costs $c: E \rightarrow \mathbb{R}$

problem: find a minimum-mean cycle K in G with respect to c

We will need the following notation.

Definition 4.16. We let $d_c^{(=k)}(u, v)$ denote the minimum cost among all walks from u to v with exactly k arcs.

We first characterize the cases in which the minimum-mean cycle has cost exactly 0 and show how to find such a cycle.

Lemma 4.17. Let c not induce negative cycles in G and let all vertices be reachable from $s \in V$. Then G has a cycle of cost 0 if and only if there is $v \in V$ with

$$\max_{0 \leq k < n} (d_c^{(=n)}(s, v) - d_c^{(=k)}(s, v)) = 0.$$

In this case, every shortest walk with n arcs from s to v contains such a cycle.

Proof. It holds that $d_c^{(=n)}(s, v) \geq d_c(s, v) = \min_{0 \leq k < n} d_c^{(=k)}(s, v)$, since c does not induce negative cycles in G . Hence, we have $\max_{0 \leq k < n} (d_c^{(=n)}(s, v) - d_c^{(=k)}(s, v)) \geq 0$ for all vertices $v \in V$. Now let K be a cycle of cost 0 and $u \in K$ be a vertex on K . We consider a shortest s - u -path and extend it by adding n copies of K . The resulting walk is a shortest s - u -walk, thus, by optimum substructure (see lecture *Algorithmic Discrete Mathematics*), the prefix W with exactly n arcs must be a shortest walk from s to some vertex $v \in K$. This implies $d_c^{(=n)}(s, v) = c(W) = d_c(s, v) = \min_{0 \leq k < n} d_c^{(=k)}(s, v)$ and thus $\max_{0 \leq k < n} (d_c^{(=n)}(s, v) - d_c^{(=k)}(s, v)) = 0$.

Conversely, if there is a vertex $v \in V$ with $\max_{0 \leq k < n} (d_c^{(=n)}(s, v) - d_c^{(=k)}(s, v)) = 0$, then we have $d_c^{(=n)}(s, v) \leq d_c^{(=k)}(s, v)$ for all $k \in \{0, \dots, n-1\}$, and hence $d_c^{(=n)}(s, v) = d_c(s, v)$. It follows that every s - v -walk with n arcs and length $d_c^{(=n)}(s, v)$ is a shortest s - v -walk and must therefore contain a cycle of cost 0 (note that every walk with n arcs must contain a cycle). \square

We can now apply Lemma 4.17 to the general case by shifting costs so that the minimum cycle has cost 0.

Lemma 4.18. Let all vertices of G be reachable from $s \in V$. Then the minimum mean cost mc among all cycles in G satisfies

$$mc = \min_{v \in V} \max_{0 \leq k < n} \frac{d_c^{(=n)}(s, v) - d_c^{(=k)}(s, v)}{n - k}.$$

Let v be the vertex for which the minimum is attained. Then every shortest s - v -walk with n arcs contains a cycle of minimum mean cost.

Proof. Let $mc' := \min_{v \in V} \max_{0 \leq k < n} \frac{d_c^{(=n)}(s, v) - d_c^{(=k)}(s, v)}{n - k}$. We first show that $mc' \leq mc$.

Consider the cost function given by $c'(e) := c(e) - mc$ for all $e \in E$. Compared to c , the mean cost of every walk is shifted by exactly mc with respect to c' . In particular, c' does not induce negative cycles and it induces a cycle of cost exactly 0. Then $d_{c'}^{(=k)}(u, v) = d_c^{(=k)}(u, v) - k \cdot mc$ is the cost of a shortest u - v -walk with exactly k arcs with respect to c' . By Lemma 4.17, there exists $v \in V$ with $\max_{0 \leq k < n} (d_{c'}^{(=n)}(s, v) - d_{c'}^{(=k)}(s, v)) = 0$. We obtain

$$\begin{aligned} mc' &\leq \max_{0 \leq k < n} \frac{d_c^{(=n)}(s, v) - d_c^{(=k)}(s, v)}{n - k} \\ &= \max_{0 \leq k < n} \frac{d_{c'}^{(=n)}(s, v) - d_{c'}^{(=k)}(s, v) + (n - k) \cdot mc}{n - k} \\ &= \max_{0 \leq k < n} \frac{d_{c'}^{(=n)}(s, v) - d_{c'}^{(=k)}(s, v)}{n - k} + mc \\ &= mc. \end{aligned}$$

It remains to show that $mc' \geq mc$. Let $v \in V$ be a vertex for which the minimum mc' is attained. We define $c''(e) := c(e) - mc'$ for all $e \in E$. This shifts the mean cost of every walk by mc' . Since $mc' \leq mc$, we know that c'' does not induce negative cycles in G . Then $d_{c''}^{(=k)}(s, v) := d_c^{(=k)}(s, v) - k \cdot mc'$ is the cost of a shortest s - v -walk with exactly k arcs with respect to c'' . For v it holds that

$$\begin{aligned} mc' &= \max_{0 \leq k < n} \frac{d_c^{(=n)}(s, v) - d_c^{(=k)}(s, v)}{n - k} \\ &= \max_{0 \leq k < n} \frac{d_{c''}^{(=n)}(s, v) - d_{c''}^{(=k)}(s, v) + (n - k) \cdot mc'}{n - k} \\ &= \max_{0 \leq k < n} \frac{d_{c''}^{(=n)}(s, v) - d_{c''}^{(=k)}(s, v)}{n - k} + mc', \end{aligned}$$

hence $\max_{0 \leq k < n} \frac{d_{c''}^{(=n)}(s, v) - d_{c''}^{(=k)}(s, v)}{n - k} = 0$. Since c'' does not induce negative cycles, Lemma 4.17 implies that c'' induces a cycle of cost 0. This cycle can be found as claimed and has mean cost mc' with respect to the original costs c . By definition of mc , it follows that $mc' \geq mc$. \square

With this, we efficiently find a minimum-mean cycle.

Algorithm: MINMEANCYCLE(G, c)

input: graph $G = (V, E)$, costs $c : E \rightarrow \mathbb{R}$

output: cycle of minimum mean cost

$G' \leftarrow (V \cup \{s\}, E \cup (\{s\} \times V))$

$$c'(u, v) \leftarrow \begin{cases} 0, & \text{for } u = s, v \in V, \\ c(u, v), & \text{for } u \in V, v \in V. \end{cases}$$

$d_{sv}^{(=k)} \leftarrow \text{BELLMANFORD}(-\text{MOORE})'(G', c')$

$v \in \arg \min_{u \in V} \max_{0 \leq k < n} (d_{su}^{(=n)} - d_{su}^{(=k)}) / (n - k)$

$W \leftarrow \text{BACKTRACK}(d_{sv}^{(=n)}, s, v)$ (shortest s - v -walk with n arcs)

return any cycle in W (or \emptyset if none exists)

Theorem 4.19. We can find a cycle of minimum-mean cost in time $\mathcal{O}(nm)$.

Proof. We can introduce a vertex s and connect it to all other vertices, and use an algorithm very similar to BELLMANFORD(-MOORE) to compute $d_c^{(=k)}(s, v)$ for all $0 \leq k \leq n$ and $v \in V$. Lemma 4.18 gives us a way to obtain a cycle of minimum-mean cost (see algorithm MINMEANCYCLE). \square

Using this together with Theorem 4.15 yields the following running time.

Theorem 4.20. MINMEANCYCLECANCELLING can be implemented to compute a min-cost b -flow in time $\mathcal{O}(m^3 n^2 \log n)$.

4.2 Successive Shortest Paths

In the previous section, we started with a b -flow that may be too expensive and repeatedly decreased its cost by augmenting along negative cycles in the residual graph. A different approach is to start with some b' -flow that is guaranteed to have minimum cost among all b' -flows, but may have $b' \neq b$, and to then repeatedly send flows along paths in order to transition to a b -flow. Such a transition is made possible by the following insight.

Lemma 4.21. Let f be a minimum-cost b -flow and f' be a b' -flow that results from augmenting f along a shortest (with respect to c) path P between some vertices s and t in G_f . Then, f' is a minimum-cost b' -flow.

Proof. For the sake of contradiction, assume that f' did not have minimum cost. Then, by Theorem 4.10, c induces a negative cycle K in $G_{f'}$.

Let $\bar{E}_{PK} := \{e \in E \cup \bar{E} \mid e, \bar{e} \in P \cup K\}$. We consider the graphs $H_P = (V, E_{PK}) := (V, P \setminus \bar{E}_{PK})$, $H_K := (V, K \setminus \bar{E}_{PK})$, and $H_{PK} := (V, (P \cup K) \setminus \bar{E}_{PK})$. No arc $e \in G_{f'} \setminus G_f$ can be contained in H_{PK} , since $\bar{e} \in P$. Hence, $H_{PK} \subseteq G_f$. Therefore, as f has minimum cost, by Theorem 4.10, H_{PK} cannot contain negative cycles. We have

$$c(H_P) + c(H_K) = c(P) + c(K) < c(P).$$

Moreover, we have $|\delta_{H_P}^+(v)| + |\delta_{H_K}^+(v)| = |\delta_{H_P}^-(v)| + |\delta_{H_K}^-(v)|$ for all $v \in V \setminus \{s, t\}$, and $|\delta_{H_P}^+(s)| + |\delta_{H_K}^+(s)| = |\delta_{H_P}^-(s)| + |\delta_{H_K}^-(s)| + 1$. This means that $g : E_{PK} \rightarrow \{1, 2\}$ with $g(e) = |\{H \in \{H_P, H_K\} \mid e \in H\}|$ defines an s - t -flow of value $|g| = 1$ in $H_{PK} \subseteq G_f$. We can therefore decompose g into an s - t -path P' and a set of cycles (see lecture *Algorithmic Discrete Mathematics*). Since these cycles cannot have negative costs and $c(g) = c(H_P) + c(H_K)$, we have $c(P') \leq c(H_P) + c(H_K) < c(P)$, which contradicts our choice of P . \square

Lemma 4.21 immediately suggests another algorithm similar to EDMONDSKARP. To start our algorithm, we need some minimum-cost b' -flow for arbitrary b' . By Theorem 4.10, such a flow can be obtained by eliminating all negative cycles. This can be accomplished easily, by fully saturating all arcs of negative costs.

Algorithm: SUCCESSIVESHORTESTPATH(G, μ, b, c)

input: circ. network (G, μ, b) , costs $c: E \rightarrow \mathbb{R}$

output: minimum cost b -flow

$b, f \leftarrow \text{MINCOSTSTARTFLOW}(G, \mu, b, c)$

while $b \neq 0$:

if $\exists s$ - t -path in G_f with $b(s) > 0$ and $b(t) < 0$:

$P \leftarrow \text{SHORTESTPATH}(G_f, c, s, t)$

 PUSH(P)

else

fail

return f

Algorithm: MINCOSTSTARTFLOW(G, μ, b, c)

input: circ. network (G, μ, b) , costs $c: E \rightarrow \mathbb{R}$

output: demands $b': V \rightarrow \mathbb{R}$,

min-cost $(b - b')$ -flow f

$f \leftarrow 0$

for $e \in E$:

if $c(e) < 0$:

$f(e) \leftarrow \mu(e)$

for $v \in V$:

$b'(v) \leftarrow b(v) + \text{ex}_f(v)$

return b', f

Algorithm: PUSH(P)

$\gamma \leftarrow \min\{\min_{e \in P} \mu_f(e), b(s), -b(t)\}$

$b(s) \leftarrow b(s) - \gamma, b(t) \leftarrow b(t) + \gamma$

AUGMENT(P, γ)

To show that SUCCESSIVESHORTESTPATH is correct we need to show that it only fails if no b -flow exists.

Theorem 4.22. SUCCESSIVESHORTESTPATH is correct, i.e., if it terminates without failing, it returns a minimum-cost b -flow.

Proof. Let b^* be the original value of b at the start of the algorithm, and let b_0 be the value after computation of the starting flow f_0 . We first show that the algorithm only fails if no b^* -flow exists. A b^* -flow exists in G if and only if a b_0 -flow exists in G_{f_0} . Without loss of generality, we may assume that $f_0 = 0$, $b_0 = b^*$, and $G_{f_0} = G$, otherwise we could consider $f \ominus f_0$ instead of f in the following.

Let (G', μ', s, t) be the network of Observation 4.3 with respect to b_0 and G_{f_0} , let b' be the value of b when the algorithm fails, and let f be the $(b_0 - b')$ -flow at this time. By definition of the algorithm and the PUSH operation, we have $|b'(v)| \leq |b_0(v)|$ and $b'(v) \cdot b_0(v) \geq 0$ for all $v \in V$. We can therefore define the s - t -flow f' in (G', μ', s, t) with $f'(u, v) = f(u, v)$ for all $(u, v) \in E$, with $f'(s, v) = b_0(v) - b'(v) \in [0, b_0(v)]$ for all $(s, v) \in \delta_G^+(s)$, and with $f'(v, t) = -b_0(v) + b'(v) \in [0, -b_0(v)]$ for all $(v, t) \in \delta_G^-(t)$. When SUCCESSIVESHORTESTPATH does not find an augmenting path anymore, there is no s - t -path in $G'_{f'}$. But then f' is a maximum s - t -flow of value $|f'| = \frac{1}{2} \sum_{v \in V} |b_0(v) - b'(v)| < \frac{1}{2} \sum_{v \in V} |b_0(v)|$ in (G', μ', s, t) , since $b' \neq 0$. By Observation 4.3, this means that no b_0 -flow exists in G_{f_0} , and hence no b^* -flow in G .

Since f_0 saturates all arcs of negative costs and does not use other arcs, f_0 is a minimum-cost b_0 -flow. By Lemma 4.21, f remains a minimum-cost flow (for changing demands). Since the algorithm only properly terminates when f is a b^* -flow, it follows that SUCCESSIVESHORTESTPATH returns a minimum-cost b^* -flow if it terminates. \square

For integral demands and capacities, we can bound the total number of iterations of SUCCESSIVESHORTESTPATH. In case of irrational data, similarly to FORDFULKERSON, the algorithm may not terminate.




Figure 4.1: Recursive construction of the network N_i with $2(i + 1)$ vertices on which `SUCCESSIVESHORTESTPATH` needs exponentially many iterations. Each arc is labeled by its cost and its capacity in this order, and we set $b(s_i) = 2^i$ and $b(t_i) = -2^i$. The cost of the shortest s_i - t_i -path in iteration $j \in \{1, \dots, 2^i\}$ is exactly $j - 1$.

Corollary 4.23. If all demands and all capacities are integral, `SUCCESSIVESHORTESTPATH` computes a minimum-cost b -flow in $\mathcal{O}(B)$ augmentations, where $B := \sum_{v \in V} |b(v)|$.

Proof. If demands and capacities are integral, in every augmentation from some vertex $s \in V$ to some vertex $t \in V$, the absolute values of $b(s)$ and $b(t)$ are reduced by at least one, while their signs remain unchanged. \square

Observe that the above bound is not polynomial in n and m , since demands can have exponential values. The bound is tight, i.e., there are instances on which `SUCCESSIVESHORTESTPATH` requires exponentially many iterations (see Figures 4.1 and 4.2) (note that the network N_i has $2(i + 1)$ vertices).




Figure 4.2: Illustration of the iterations performed by `SUCCESSIVESHORTESTPATH` on the network N_2 of Figure 4.1. The shortest path in each iteration is marked in red, and arcs are oriented in the direction in which they are used next. Note that, after $2^i = 4$ iterations, the configuration of the residual graph is the same as in the beginning if we switch the roles of s_2 and t_2 . This allows to recursively embed the network into N_3 and so on.

Proposition 4.24. `SUCCESSIVESHORTESTPATH` needs $2^i = \Theta(B)$ augmentations to compute a minimum-cost flow in the network N_i of Figure 4.1 for any $i \in \mathbb{N}$.

In order to bound the running time of `SUCCESSIVESHORTESTPATH` beyond bounding the number of iterations, we need to specify how shortest paths are computed. As shown in Section 2.2.1, we can use `DIJKSTRA`, provided that we have a vertex potential at hand. We cannot recompute such a potential in every iteration, since this would dominate the running time of `DIJKSTRA`. However, we can compute a vertex potential once and update it cheaply in every iteration.

Algorithm: `SUCCESSIVESHORTESTPATH`(G, μ, b, c) (potential)

input: circ. network $(G = (V, E), \mu, b)$, costs $c : E \rightarrow \mathbb{R}$

output: minimum cost b -flow

$b, f \leftarrow \text{MINCOSTSTARTFLOW}(G, \mu, c)$

$\pi \leftarrow \text{POTENTIAL}(G_f, c)$

while $b \neq 0$:

if $\exists s$ - t -path in G_f **with** $b(s) > 0$ **and** $b(t) < 0$:

$E' \leftarrow ((\{t\} \times V) \setminus G_f)$

$c'(E) \leftarrow c_\pi(E), c'(E') \leftarrow 0$

$G' \leftarrow (V, E_f \cup E')$

$(d_{sv})_{v \in V} \leftarrow \text{DIJKSTRA}(G', c', s)$

$P \leftarrow \text{BACKTRACK}((d_{sv})_{v \in V}, s, t)$

`PUSH`(P)

for $v \in V$:

$\pi(v) \leftarrow \pi(v) + d_{sv}$

else

fail

return f

We now show that vertex potentials are maintained correctly.

Theorem 4.25. `SUCCESSIVESHORTESTPATH` with potential is correct, i.e., if it terminates, then it returns a minimum-cost b -flow.

Proof. By Theorem 4.22, it suffices to show that π is always a vertex potential and we choose a shortest augmenting path in G_f in every step. By correctness of `POTENTIAL` (Proposition 2.19), π is a potential at the start of the algorithm. Let f_i be the b -flow at the start of iteration i , let π_i be the potential at the start of iteration i , and let P_i be the s - t -path chosen in this iteration. We show that the algorithm maintains the invariant that π_i is a potential for G_{f_i} at the start of iteration i , and that P_i is a shortest path (with respect to c) in G_{f_i} .

Since, by our invariant, it holds that $c_{\pi_i} \geq 0$, we have $c' \geq 0$ and we choose P_i as a shortest path in G' with respect to c' . We may thus assume that P_i does not use arcs from E' . It follows that P_i is a shortest s - t -path in G_{f_i} with respect to c_{π_i} , and thus also with respect to c (Proposition 2.18).

Since, by our invariant, π_i is a potential at the start of iteration i , `DIJKSTRA` computes the values $(d_{sv})_{v \in V}$ such that $d_{sv} = d_{c', G'}(s, v)$. We have to show that $c_{\pi_{i+1}}$ is a potential, i.e., that $c_{\pi_{i+1}}(u, v) = c(u, v) +$

$\pi_{i+1}(u) - \pi_{i+1}(v) \geq 0$ for all $(u, v) \in G_{f_{i+1}}$ with $\pi_{i+1}(v) = \pi_i(v) + d_{sv}$ for all $v \in V$. First consider arcs $(u, v) \in G_{f_i} \cap G_{f_{i+1}}$. We have $c_{\pi_i}(u, v) + d_{su} \geq d_{sv}$ and thus

$$\begin{aligned} c_{\pi_{i+1}}(u, v) &= c(u, v) + \pi_{i+1}(u) - \pi_{i+1}(v) \\ &= c(u, v) + \pi_i(u) + d_{su} - \pi_i(v) - d_{sv} \\ &= c_{\pi_i}(u, v) + d_{su} - d_{sv} \\ &\geq 0. \end{aligned}$$

Now consider arcs $(u, v) \in G_{f_{i+1}} \setminus G_{f_i}$. We have $(v, u) \in P_i$, since only the flow along P_i was changed in iteration i . Because P_i is a shortest path, by optimum substructure (see lecture *Algorithmic Discrete Mathematics*), we have $d_{su} = d_{sv} + c_{\pi_i}(v, u) = d_{sv} - c_{\pi_i}(u, v)$, and thus the inequality above holds even with equality. \square

This allows us to bound the overall running time. Note that, exponential worst-case running time notwithstanding, SUCCESSIVESHORTESTPATH is the best known algorithm for the case $B = \mathcal{O}(n)$ and $\mu \in \mathbb{N}$, $b \in \mathbb{Z}$.

Corollary 4.26. If all demands and all capacities are integral, SUCCESSIVESHORTESTPATH with potential computes a minimum-cost b -flow in time $\mathcal{O}(nm + B(m + n \log n))$, where $B := \sum_{v \in V} |b(v)|$.

Proof. This follows from Theorem 4.25 and Corollary 4.23, together with the running times of POTENTIAL (Proposition 2.19) and DIJKSTRA (Theorem 2.2). \square

Remark 4.27. Note that we can slightly improve the running time of SUCCESSIVESHORTESTPATH by initializing the potential to 0 instead of invoking POTENTIAL. We can do this, since we eliminate all arcs of negative costs from the initial residual graph. However, this approach no longer works for the algorithm of the next section.

4.2.1 Capacity Scaling

Intuitively, in the example of Figure 4.1, the problem is that SUCCESSIVESHORTESTPATH augments the flow by very small amounts in each step, even though paths of larger capacity would be available. It turns out that we can significantly improve performance by prioritizing augmenting paths that allow for large changes in flow values. We need to balance this carefully with trying to augment along cheap paths first. This can be accomplished by operating in phases, where each phase only considers paths that change the flow by at least some threshold value. Once a phase is complete, we lower the threshold and repeat the process. This general method is called *capacity scaling*.

For convenience, from now on, we assume that $\mu = \infty$, i.e., capacities are unbounded. We can transform every instance to have this property, as shown in the following lemma. Note, however, that this transformation affects the running time of all algorithms, since we introduce an additional vertex for each arc.

Lemma 4.28. Every instance $(G = (V, E), \mu, b, c)$ of the minimum-cost flow problem can be transformed into an equivalent, bipartite instance $(G' = (U' \cup V, E'), \infty, b', c')$ with $n + m$ vertices and $2m$ arcs.

Proof. We set $U' = E$, so that G' has a vertex for every vertex and every arc of G . For every arc $e = (u, v) \in E$ we introduce arcs $(e, u) \in E'$ and $(e, v) \in E'$, with $c'(e, u) = 0$ and $c'(e, v) = c(e)$. We further define the demand by $b'(e) = \mu(e)$ for all $e \in E$ and $b'(v) = b(v) - \sum_{e \in \delta_G^+(v)} \mu(e)$ for all $v \in V$.

Let f be a b -flow in (G, μ, b, c) . For $e = (u, v) \in E$, we set $f'(e, v) = f(e)$ and $f'(e, u) = \mu(e) - f(e)$. Then, f' is a b' -flow in (G', ∞, b', c') of cost $c'(f') = c(f)$. Conversely, let f' be a b' -flow in (G', ∞, b', c') . We set $f(u, v) = f'((u, v), v)$ for all $(u, v) \in E$ and obtain a b -flow f with $c(f) = c'(f')$. \square

Eliminating capacities allows us to cleanly state a capacity-scaling algorithm, where the scaling aspect only reflects in our selection of source and sink in each step.

Algorithm: CAPACITYSCALING(G, b, c)

input: circ. network $(G = (V, E), \infty, b)$, costs $c: E \rightarrow \mathbb{R}$
output: minimum cost b -flow

$f \leftarrow 0$
 $\gamma \leftarrow 2^{\lfloor \log(\max_{v \in V} b(v)) \rfloor}$
while $b \neq 0$:
 while $\exists s$ - t -path in G_f **with** $b(s) \geq \gamma$ **and** $b(t) \leq -\gamma$:
 $P \leftarrow \text{SHORTESTPATH}(G_f, c, s, t)$
 PUSH(P, γ)
 $\gamma \leftarrow \gamma/2$
 if $\gamma < 1$ **and** $b \neq 0$:
 fail
 return f

Algorithm: PUSH(P, γ)

if $\gamma < 0$:
 PUSH($\bar{P}, -\gamma$)
else
 $b(s) \leftarrow b(s) - \gamma$
 $b(t) \leftarrow b(t) + \gamma$
 AUGMENT(P, γ)

We now show that this algorithm is indeed still correct and achieves a better (weakly polynomial) running time than SUCCESSIVESHORTESTPATH.

Theorem 4.29. If $\mu = \infty$, $b(v) \in \mathbb{Z}$ for all $v \in V$, and c does not induce negative cycles, CAPACITYSCALING computes a minimum-cost b -flow in time $\mathcal{O}(n(m + n \log n) \log b_{\max})$, where $b_{\max} := \max_{v \in V} b(v)$.

Proof. First observe that all residual capacities are either ∞ or integer multiples of γ throughout the algorithm, hence all PUSH operations are well-defined. Since b is integral throughout, the algorithm only fails if there exists no s - t -path with $b(s) > 0$ and $b(t) < 0$. As in the proof of Theorem 4.22, this implies that no solution exists. With this, correctness follows again by Lemma 4.21 together with the fact that the initial flow ($f = 0$) is a minimum-cost circulation (i.e., 0-flow), since c does not induce negative cycles (Theorem 4.10).

We divide the execution of CAPACITYSCALING into maximal phases in which γ is constant and we show that the algorithm takes at most n iterations in each phase. To this end, consider f, b at the start of a phase associated to any fixed value of γ . We define

$$\begin{aligned} S &:= \{v \in V \mid b(v) \geq \gamma\}, \\ S^+ &:= \{v \in V \mid b(v) \geq 2\gamma\}, \\ T &:= \{v \in V \mid b(v) \leq -\gamma\}, \\ T^+ &:= \{v \in V \mid b(v) \leq -2\gamma\}, \end{aligned}$$

and let R be the set of all vertices that are reachable from an element of S^+ in G_f . We have $S^+ \subseteq R$ by definition, and $R \cap T^+ = \emptyset$, as the previous phase ended because no paths remain from S^+ to T^+ .

Let (s_i, t_i) be the pairs (s, t) chosen during the current phase. Since $\delta_{G_f}^+(R) = \emptyset$ and since every augmentation is by the same amount γ , every augmentation with $s_i \in R$ and $t_i \notin R$ reduces the capacity of the cut $\delta_{G \cup \bar{G}}^+(R)$ by γ . This means that for every augmentation across $\delta_{G \cup \bar{G}}^+(R)$ we first need to augment at least once across $\delta_{G \cup \bar{G}}^-(R)$. We obtain

$$|\{i \mid s_i \in R \wedge t_i \notin R\}| \leq |\{i \mid s_i \notin R \wedge t_i \in R\}|.$$

The total number k of iterations in the phase is therefore bounded by

$$\begin{aligned} k &= |\{i \mid s_i \in R \wedge t_i \in R\}| + |\{i \mid s_i \in R \wedge t_i \notin R\}| + |\{i \mid s_i \notin R \wedge t_i \in R\}| + |\{i \mid s_i \notin R \wedge t_i \notin R\}| \\ &\leq |\{i \mid s_i \in R \wedge t_i \in R\}| + 2 \cdot |\{i \mid s_i \notin R \wedge t_i \in R\}| + |\{i \mid s_i \notin R \wedge t_i \notin R\}| \\ &= |\{i \mid t_i \in R\}| + |\{i \mid s_i \notin R\}|. \end{aligned}$$

But since $R \cap T^+ = \emptyset$, we have $b(t) > -2\gamma$ for every sink $t \in R$, which means that every sink can be used at most once. This implies $|\{i \mid t_i \in R\}| \leq |T \cap R|$. Analogously, every source $s \notin R$ can only be used once, since $S^+ \setminus R = \emptyset$. This implies $|\{i \mid s_i \notin R\}| \leq |S \setminus R|$. We obtain

$$k \leq |T \cap R| + |S \setminus R| \leq |T| + |S| \leq n.$$

This establishes that the algorithm needs at most n iterations per phase. Obviously, there are $\mathcal{O}(\log b_{\max})$ phases overall, and every phase can be implemented as in `SUCCESSIVESHORTESTPATH` with potential. This yields the claimed running time. \square

4.3 Orlin's Algorithm

In `CAPACITYSCALING`, we reduce γ by factor 2 after each phase. Assuming, for the moment, that we have p augmentations per phase, we can conclude that the total magnitude of all augmentations after a γ -phase is $p \sum_{i \geq 1} 2^{-i} \gamma = p\gamma$. This means that all arcs $\bar{e} \in \bar{E}$ for which $f(e) > p\gamma$ at the end of the γ -phase remain in the residual graph from now on (and so does e , because of $\mu = \infty$). This insight allows us to improve efficiency by treating connected components formed by such arcs e, \bar{e} as single vertices. Importantly, we know that there cannot be cheaper paths connecting the endpoints of e, \bar{e} , since this would induce negative cycles. We can therefore discard all other arcs between the components of the two endpoints of e, \bar{e} before merging these components. This gradual simplification of the graph significantly improves the running time of the algorithm, as we will see. We first formalize the notion of merging components (cf. Figure 4.3).

Definition 4.30. Let $\{u, v\} \in E$ be an edge of an undirected graph $G = (V, E)$, and let $c : E \rightarrow \mathbb{R}$ be edge costs. The result of *contracting* $\{u, v\}$ in G, c is $G' = (V', E')$, c' with $V' := (V \cup \{x_{uv}\}) \setminus \{u, v\}$ and $E' := (E \setminus \delta_G(\{u, v\})) \cup \{\{x_{uv}, w\} \mid w \in \Gamma_G(\{u, v\})\}$ and $c'(e) = c(e)$ for $e \in E' \cap E$ and $c'(x_{uv}, w) := \min\{c(u, w), c(v, w)\}$ for $\{x_{uv}, w\} \in E' \setminus E$. Contracting arcs of directed graphs is defined analogously.

We do not explicitly contract arcs, but rather keep track of the forest F of all arcs that have been contracted and the subgraph $H \subseteq G$ in which all other arcs within components have been removed. To further simplify the algorithm, we refine a *representative* r for each component in F and we maintain $r_v = r$ for all vertices v of this component. When merging components we ensure that $b(v) = 0$ if v is not the representative of the new component, by augmenting flow once along the unique path between the two representatives. Finally, we further improve the algorithm by reducing γ by more than factor 2 whenever possible. With these modifications we obtain Orlin's algorithm.




Figure 4.3: Illustration of the contraction of an arc (u, v) . Left: original graph, right: contracted graph.

Algorithm: ORLIN(G, b, c)

input: circ. network $(G = (V, E), \infty, b)$, costs $c: E \rightarrow \mathbb{R}$

output: minimum cost b -flow

$f \leftarrow 0$

$r_v \leftarrow v \quad \forall v \in V$

$F \leftarrow \emptyset, H \leftarrow G, \varepsilon \leftarrow 1/n$

$\gamma \leftarrow \max_{v \in V} |b(v)|$

while $b \neq 0$:

while $\exists v \in V$ **with** $|b(v)| > (1 - \varepsilon)\gamma$:

if $\exists s$ - t -path in G_f **with** $b(s) > (1 - \varepsilon)\gamma \wedge b(t) < -\varepsilon\gamma$

or $b(s) > \varepsilon\gamma \wedge b(t) < -(1 - \varepsilon)\gamma$:

 | $P \leftarrow \text{SHORTESTPATH}(H_f, c, s, t)$

 | $\text{PUSH}(P, \gamma)$

else

 | **fail**

if $f(e) = 0 \forall e \in H \setminus F$:

 | $\gamma \leftarrow \min\{\gamma/2, \max_{v \in V} |b(v)|\}$

else

 | $\gamma \leftarrow \gamma/2$

while $\exists e \in H \setminus F$ **with** $f(e) > 6n\gamma$:

 | $\text{CONTRACT}(e)$

return f

Algorithm: CONTRACT(e)

input: arc $e \equiv (u, v) \in H_f$

$F \leftarrow F \cup \{e, \bar{e}\}$
 $Q \leftarrow r_u$ - r_v -path in F
PUSH($Q, b(r_u)$)
for $e' \equiv (x, y) \in H \setminus F$ **with** $\{r_x, r_y\} = \{r_u, r_v\}$:
 $\lfloor H \leftarrow H \setminus \{e'\}$
for $y \in V$ **with** $r_y = r_u$:
 $\lfloor r_y \leftarrow r_v$

An important consequence of our contraction strategy is that all arcs that have not been contracted or removed carry a flow that is an integer multiple of γ . This means we do not have to worry that our PUSH operations may not be well-defined.

Observation 4.31. Throughout the execution of ORLIN, it holds that $f(e)$ is an integer multiple of γ for all $e \in H \setminus F$.

Proof. As long as $f(e) \neq 0$, we halve γ in each iteration. Since we only augment by γ in H_f and since CONTRACT only uses arcs in F , it is ensured that $f(e)$ remains an integer multiple of γ . Note that we can reduce γ by more than half if $f(e) = 0$ for all $e \in H \setminus F$ without violating this invariant. \square

We divide the execution of ORLIN into maximal phases in which γ stays constant. For $\gamma' \in \mathbb{R}$, the γ' -phase is the phase in which $\gamma = \gamma'$. We first relate the number of augmentations outside of and within CONTRACT.

Lemma 4.32. In each phase of ORLIN, the number of augmentations outside of CONTRACT is not larger than the number of augmentations within CONTRACT plus the number of vertices v with $|b(v)| > (1 - \varepsilon)\gamma$ at the beginning of the phase.

Proof. Consider the potential function $\Phi := \sum_{v \in V} k_v := \sum_{v \in V} [|b(v)|/\gamma - (1 - \varepsilon)] \geq 0$. We show that during each augmentation along some s - t -path outside of CONTRACT, this potential Φ is reduced by at least 1. Without loss of generality, let $b(s) > (1 - \varepsilon)\gamma$ and $b(t) < -\varepsilon\gamma$. In case $|b(s)| \geq \gamma$, we reduce $|b(s)|/\gamma$ by 1 and therefore also k_s by 1. Otherwise, we reduce k_s by 1 to 0. If $b(t) \leq -\gamma/2$, we do not increase $|b(t)|$, therefore k_t does not increase. Otherwise, we have $b(t) \in (-\gamma/2, -\varepsilon\gamma)$, which implies $|b(t)|/\gamma \in (\varepsilon, 1/2)$ and $|b(t) + \gamma|/\gamma \in (1/2, 1 - \varepsilon)$. Hence, k_t stays 0.

Now consider an augmentation by δ inside of CONTRACT. We increase $|b(r_v)|$ by at most δ and hence k_{r_v} by at most $\lceil \delta/\gamma \rceil$. On the other hand, we reduce k_{r_u} by at least $\lfloor \delta/\gamma \rfloor$. The increase of Φ is therefore at most 1.

At the beginning of the phase, Φ is bounded by the number of vertices v with $|b(v)| > (1 - \varepsilon)\gamma$, since $|b(v)| \leq (1 - \varepsilon)2\gamma$ and thus $k_v \leq 1$, otherwise the previous phase would not have ended. The number of augmentations outside of CONTRACT is therefore bounded by the number of such vertices plus the number of augmentations within CONTRACT. \square

We now establish that the algorithm is well-defined in the sense that it maintains a b' -flow for some b' .

Lemma 4.33. ORLIN maintains the invariant that f is a b' -flow with $f(e) > 0$ for all $e \in G \cap F$.

Proof. To show that f is a b' -flow, we need to show that $f(e) \geq 0$ for all $e \in G$. Arcs in $G \setminus H$ do not change their flow values. Consider an arc $e \in H \setminus F$. By Observation 4.31, the flow along e , and thus the residual capacity of \bar{e} , is always an integer multiple of γ . In CONTRACT we only augment along arcs in F . Since, otherwise, we only augment flow by γ and only along arcs of positive residual capacity, we never reduce $f(e)$ below 0. For $e \in G \cap F$ we additionally need to show that we never reduce $f(e)$ to 0. We show that even the total magnitude of *all* augmentations after e was added to F would not suffice to reduce $f(e)$ to 0.

We show that the total magnitude of all augmentations after a γ -phase is at most $6n\gamma$ along each arc. We know that there are at most $n - 1$ CONTRACT operations overall, since only arcs between vertices in different components of F are contracted. Let k be the number of CONTRACT operations in a γ' -phase with $\gamma' \leq \gamma$. We know that, at the beginning of the γ' -phase, we have $|b(v)| \leq 2(1 - \varepsilon)\gamma' < 2\gamma$ for all vertices $v \in V$, because the previous phase ended. It follows that the total magnitude of all augmentations along some arc within CONTRACT during the γ' -phase is at most $2k\gamma$, i.e., at most 2γ on average per CONTRACT operation (we ignore flow changes that are undone in the same phase). In total, the magnitude of all remaining augmentations along a fixed arc inside CONTRACT is strictly less than $2n\gamma$.

By Lemma 4.32, in each γ' -phase with $\gamma' \leq \gamma$, the number of augmentations outside of CONTRACT is at most equal to the number of other augmentations plus the number of vertices v with $|b(v)| > (1 - \varepsilon)\gamma'$, i.e., at most $2n - 1$. Since each such augmentation is by $\gamma' \leq \gamma$, the magnitude of these augmentations is less than $2n\gamma$.

We only insert $e \in E$ into F once $f(e) > 6n\gamma$. Note that γ is at least halved after every phase. The total magnitude in the γ -phase and all subsequent phases is thus at most

$$2n\gamma + \sum_{i=0}^{\infty} 2n \frac{\gamma}{2^i} = 6n\gamma.$$

It follows that we do not augment enough flow along \bar{e} to ever reduce the flow along e to 0. \square

The following is the key lemma for the correctness of ORLIN. We show that our simplification strategy via contractions does not introduce errors.

Lemma 4.34. ORLIN maintains the invariant that f is a minimum-cost b' -flow for some b' , provided that c does not induce negative cycles.

Proof. We show that the invariant is maintained and, additionally, shortest paths in H_f always are shortest paths in G_f . Both holds initially for $f = 0$, since we assume that c does not induce negative cycles.

We show that, as long as f has minimum cost, every shortest u - v -path P in H_f is also a shortest path in G_f . Assume there was a shorter u - v -path P' in G_f . We can replace every arc $e \in P' \setminus H_f$ by a (unique) path P_e in F to obtain a u - v -walk W in H_f . Since f has minimum cost, c cannot induce negative cycles in G_f (Theorem 4.10). By assumption, we thus have $c(P') < c(P) \leq c(W)$. This means that there is an arc $e \in P' \setminus H_f$ for which $c(P_e) > c(e)$. By Lemma 4.33 and since capacities are unbounded, all reverse arcs of arcs $e' \in P_e$ have a positive residual capacity, thus $\bar{P}_e \subseteq G_f$. We consider the cycle $e \oplus \bar{P}_e \subseteq G_f$. The cost of this cycle is

$$c(e) + c(\bar{P}_e) = c(e) - c(P_e) < 0,$$

hence c induces a negative cycle in G_f . This is a contradiction to optimality of f (Theorem 4.10), and we conclude that P must be a shortest path in G_f .

Consequently, in every step, we augment along a shortest path in G_f . By Lemma 4.21, it follows that the resulting flow still has minimum cost, which means that our invariant is maintained. \square

Finally, we conclude that ORLIN computes the correct result.

Theorem 4.35. If c does not induce negative cycles, ORLIN is correct, i.e., if it terminates without failing, then it returns a minimum-cost b -flow.

Proof. If ORLIN terminates without failing, the statement follows from $b = 0$ and Lemma 4.34. Now assume, without loss of generality, that the algorithm fails because there is vertex s with $b(s) > (1 - \varepsilon)\gamma$ but no vertex t with $b(t) < -\varepsilon\gamma$ that is reachable from s in G_f . Let R be the set of all vertices reachable from s in G_f . Since capacities are unbounded, we have $\delta_G^+(R) = \emptyset$ (and not just $\delta_{G_f}^+(R) = \emptyset$). Moreover, we have $f(e) = 0$ for all $e \in \delta_G^-(R)$, since $\delta_{G_f}^+(R) = \emptyset$. Because f is a $(b^* - b)$ -flow, where b^* is the original value of b , it follows that $\sum_{v \in R} (b^*(v) - b(v)) = \sum_{e \in \delta_G^+(R)} f(e) - \sum_{e \in \delta_G^-(R)} f(e) = 0$. We obtain

$$\sum_{v \in R} b^*(v) = \sum_{v \in R} b(v) = b(s) + \sum_{v \in R \setminus \{s\}} b(v) > (1 - \varepsilon)\gamma - (n - 1)\varepsilon\gamma = 0,$$

but $\delta_G^+(R) = \emptyset$ and thus $\sum_{e \in \delta_G^+(R)} \mu(e) = 0 < \sum_{v \in R} b^*(v)$. By Proposition 4.4, this means that no b^* -flow exists in (G, ∞, b^*) , and ORLIN can safely fail. \square

To bound the running time of the algorithm, much like for the previous algorithms, we need to establish that we make progress over time.

Lemma 4.36. If we have $|b(z)| > (1 - \varepsilon)\gamma$ for $z \in V$ at some point during the execution of ORLIN, then the connected component of z in F grows within the next $\lceil 2 \log_2 n + \log_2 m \rceil + 3$ phases.

Proof. Let $|b(z)| > (1 - \varepsilon)\gamma_1$ for some $z \in V$ at the beginning of the γ_1 -phase for $\gamma_1 \in \mathbb{R}_{>0}$. Let $\gamma_0 \geq 2\gamma_1$ be the value of γ in the previous phase (or $\gamma_0 := 2\gamma_1$ if we are in the first phase) and let γ_2 be the value of γ after $k := \lceil 2 \log_2 n + \log_2 m \rceil + 2$ additional phases. Since γ is reduced by factor 2 or more after each phase, we have

$$\gamma_1 \geq 2^k \gamma_2 \geq 2^{2 \log_2 n} 2^{\log_2 m} 2^2 \gamma_2 = 4n^2 m \gamma_2. \quad (4.8)$$

Let f_1 and b_1 be the values of f and b , respectively, at the beginning of the γ_1 -phase, and let f_2 and b_2 be the values at the end of the γ_2 -phase. Let further b^* denote the original value of b , and let C_z be the connected component of z in F at the beginning of the γ_1 -phase. Assume that C_z does not grow in the next k phases after the γ_1 -phase. In CONTRACT we ensure that $b(v) = 0$ for all v with $r_v \neq v$. Therefore, we have $r_z = z$ and $b_1(v) = 0$ for all $v \in C_z \setminus \{z\}$ at the beginning of the γ_1 -phase. Since f_1 is a $(b^* - b_1)$ -flow, we further have

$$\sum_{v \in C_z} b^*(v) - b_1(z) = \sum_{v \in C_z} (b^*(v) - b_1(v)) = \sum_{e \in \delta_G^+(C_z)} f_1(e) - \sum_{e \in \delta_G^-(C_z)} f_1(e) =: r\gamma_0,$$

with $r \in \mathbb{Z}$ by Observation 4.31, since $(\delta_G^-(C_z) \cup \delta_G^+(C_z)) \subseteq H \setminus F$. If $r \cdot b_1(z) \geq 0$, we obtain

$$\left| \sum_{v \in C_z} b^*(v) \right| = |r\gamma_0 + b_1(z)| \geq |b_1(z)| > (1 - \varepsilon)\gamma_1 \stackrel{n \geq 2}{\geq} \frac{1}{n}\gamma_1, \quad (4.9)$$

and otherwise (i.e., if $r \cdot b_1(z) < 0$) we have $r \neq 0$ and thus (note that $|b_1(z)| < (1 - \varepsilon)\gamma_0$, since the previous phase ended)

$$\left| \sum_{v \in C_z} b^*(v) \right| = |r\gamma_0 + b_1(z)| = ||r\gamma_0| - |b_1(z)|| > \gamma_0 - (1 - \varepsilon)\gamma_0 = \frac{1}{n}\gamma_0 > \frac{1}{n}\gamma_1. \quad (4.10)$$

Since f_2 is a $(b^* - b_2)$ -flow, we obtain as above

$$\sum_{v \in C_z} b^*(v) - b_2(z) = \sum_{e \in \delta_G^+(C_z)} f_2(e) - \sum_{e \in \delta_G^-(C_z)} f_2(e). \quad (4.11)$$

Together with $|b_2(z)| \leq (1 - \varepsilon)\gamma_2$ (since the γ_2 -phase ended), this yields

$$\begin{aligned} \sum_{e \in \delta_G^+(C_z) \cup \delta_G^-(C_z)} f_2(e) &\geq \left| \sum_{e \in \delta_G^+(C_z)} f_2(e) - \sum_{e \in \delta_G^-(C_z)} f_2(e) \right| \\ &\stackrel{(4.11)}{=} \left| \sum_{v \in C_z} b^*(v) - b_2(z) \right| \\ &\geq \left| \sum_{v \in C_z} b^*(v) \right| - |b_2(z)| \\ &\stackrel{(4.9), (4.10)}{\geq} \frac{1}{n}\gamma_1 - (1 - \varepsilon)\gamma_2 \\ &\stackrel{(4.8)}{\geq} 4nm\gamma_2 - (1 - \varepsilon)\gamma_2 \\ &> (4nm - 1)\gamma_2 \\ &> m(6n \frac{\gamma_2}{2}). \end{aligned}$$

But this means that at least one arc in $\delta_G^+(C_z) \cup \delta_G^-(C_z)$ carries a flow of at least $6n\gamma$ at the beginning of the next phase and is therefore contracted. This causes the component of z in F to grow. \square

Now, we have everything we need to bound the running time of ORLIN. Note that it is crucial that we reduce γ by more than half whenever possible.

Theorem 4.37. ORLIN can be implemented to run in time $\mathcal{O}(n \log n(m + n \log n))$.

Proof. We can make use of a vertex potential as in SUCCESSIVESHORTESTPATH with potential. This allows to compute shortest paths in time $\mathcal{O}(m + n \log n)$. It remains to bound the number of augmentations and phases. There are at most $n - 1$ augmentations within CONTRACT, since we merge two distinct components in each CONTRACT operation.

We count the maximum number of phases until some arc is contracted. As soon as an augmentation is performed outside of CONTRACT, by Lemma 4.36, the number of phases until the next CONTRACT is bounded by $\mathcal{O}(\log m) = \mathcal{O}(\log n)$. Before that, however, there may be additional phases without vertices v satisfying $|b(v)| > (1 - \varepsilon)\gamma$. If there is at least one arc $e \in G \setminus F$ with $f(e) > 0$, then this arc carries at least γ units of flow (Observation 4.31). After $\lceil \log_2 n \rceil + 3$ phases without change in flow, we have $f(e) \geq 2^{\lceil \log_2 n \rceil + 3} \gamma > 6n\gamma$, since we reduce γ by at least half after each phase. Then, e can be contracted. On the other hand, if $f(e) = 0$ for all arcs $e \in G \setminus F$, the algorithm sets $\gamma = \max_{v \in V} |b(v)|$, which leads to an augmentation in the next phase. It follows that after at most $\mathcal{O}(\log n)$ phases some arc is contracted, which means that we have $\mathcal{O}(n \log n)$ phases overall.

By Lemma 4.32, the number of augmentations outside of CONTRACT is at most $n - 1$ plus the number of pairs (γ, z) with $|b(z)| > (1 - \varepsilon)\gamma$ at the beginning of the γ -phase. At such a point in time, it holds that $r(z) = z$, since only the representative of a connected component may have $b(z) \neq 0$. By Lemma 4.36, the connected component grows within the next $\mathcal{O}(\log m + \log n) = \mathcal{O}(\log n)$ phases. A connected component can therefore be responsible for at most $\mathcal{O}(\log n)$ augmentations, before it changes. The total number of different connected components throughout the course of the algorithm is $2n - 1$, since we start with n components and reduce the number of components by one in each CONTRACT operation, until we are left with a single component. Hence, the total number of augmentations is $\mathcal{O}(n \log n)$. \square



5 Maximum Cardinality Matchings

We now turn to the maximum matching problem on general graphs. Elementary results and algorithms for bipartite graphs were discussed in the lecture *Algorithmic Discrete Mathematics*. We will present a better algorithm for bipartite graphs and extend results to general graphs. We begin by recalling basic definitions.

Definition 5.1. A *matching* in an undirected graph $G = (V, E)$ is a set $M \subseteq E$ of pairwise disjoint edges. A vertex $v \in V \setminus \bigcup_{e \in M} e$ is *(M-)exposed* and we write $v \notin M$. A vertex $v \in \bigcup_{e \in M} e$ is *(M-)covered* and we write $v \in M$. The matching M is

- *perfect* if all vertices of G are M -covered (i.e., if $|M| = \frac{1}{2} |V|$),
- *maximal* if there is no matching $M' \supsetneq M$ in G ,
- *maximum (cardinality)* if it maximizes the number of edges in the matching, i.e., if $|M| = \max\{|M'| \mid M' \text{ is a matching in } G\}$.

With this, the maximum matching problem is defined as follows. We sometimes explicitly speak of maximum *cardinality* matchings to distinguish them from maximum *weight* matchings.

Maximum (Cardinality) Matching Problem

input: graph $G = (V, E)$

problem: find maximum cardinality matching $M \subseteq E$ of G

As before, from now on we fix an undirected graph $G = (V, E)$. We recall the following characterization of maximum matchings.

Definition 5.2. Let M be a matching in G . A u - v -path P in G is *(M-)alternating* if $P \setminus M$ is a matching. If, additionally, $u, v \notin M$, then P is *(M-)augmenting*.

Theorem 5.3. Let M be a matching in G . Then, M is a maximum matching in G if and only if there is no M -augmenting path in G .

5.1 Bipartite Graphs Revisited

We begin by revisiting the maximum matching problem in bipartite graphs. This problem can be reduced to the maximum flow problem by introducing the following auxiliary network (see Figure 5.1).

Definition 5.4. Let $G = (U \cup V, E)$ be a bipartite graph with $s, t \notin (U \cup V)$. The *flow network* $N_G = (H, \mu, s, t)$ induced by G is given by $H = (V_H, E_H)$ with vertices $V_H = U \cup V \cup \{s, t\}$ and arcs $E_H = \{(u, v) \in U \times V \mid \{u, v\} \in E\} \cup (\{s\} \times U) \cup (V \times \{t\})$, and $\mu(e) = 1$ for all $e \in E_H$.

With this, the maximum matching problem can easily be solved by invoking any maximum flow algorithm. For example, we may use the following algorithms. Note that FORDFULKERSON and DINIC are guaranteed to compute a maximum flow f with $f(e) \in \{0, 1\}$ for all edges e , since $\mu(e) \in \{0, 1\}$.




Figure 5.1: Left: bipartite graph G . Right: flow network N_G induced by G ($\mu = 1$).

Algorithm: FORDFULKERSONMATCHING(G)

input: bipartite graph $G = (U \cup V, E)$

output: maximum card. matching $M \subseteq E$ in G

$(H, \mu, s, t) \leftarrow N_G$

$f \leftarrow \text{FORDFULKERSON}(H, \mu, s, t)$

return $\{\{u, v\} \in E \mid (u, v) \in U \times V \wedge f((u, v)) = 1\}$

Algorithm: DINICMATCHING(G)

input: bipartite graph $G = (U \cup V, E)$

output: maximum card. matching $M \subseteq E$ in G

$(H, \mu, s, t) \leftarrow N_G$

$f \leftarrow \text{DINIC}(H, \mu, s, t)$

return $\{\{u, v\} \in E \mid (u, v) \in U \times V \wedge f((u, v)) = 1\}$

Importantly, in this application, FORDFULKERSONMATCHING beats DINICMATCHING in terms of running time.

Theorem 5.5. FORDFULKERSONMATCHING solves the maximum matching problem in time $\mathcal{O}(nm)$, and DINICMATCHING solves it in time $\mathcal{O}(n^3)$.

Proof. The first part of the statement follows from the fact that FORDFULKERSON has running time $\mathcal{O}(m \cdot |f^*|)$ (see lecture *Algorithmic Discrete Mathematics*) and $|f^*| < n$ in N_G . The running time of DINICMATCHING follows from Theorem 3.10. \square

The same way that the analysis of FORDFULKERSON can (trivially) be strengthened for flow networks induced by bipartite graphs, we can modify DINICMATCHING to exploit the special structure of these flow networks. In particular, we can compute blocking flows more efficiently in these flow networks. The resulting improvement yields the Hopcroft-Karp algorithm:

Algorithm: HOPCROFTKARP(G)

input: bipartite graph $G = (U \cup V, E)$

output: maximum cardinality matching $M \subseteq E$ in G

$(H, \mu, s, t) \leftarrow N_G$

$f \leftarrow 0$

while \exists s - t -path in H_f :

$H_{f,L} \leftarrow (V, \{(v, w) \in H_f \mid d_{H_f}(s, w) = d_{H_f}(s, v) + 1\})$

$f' \leftarrow \text{BIPARTITEBLOCKINGFLOW}(H_{f,L}, \mu_f, s, t)$

$f \leftarrow f + f'$

return $\{\{u, v\} \in E \mid (u, v) \in U \times V \wedge f((u, v)) = 1\}$

The key insight to make BIPARTITEBLOCKINGFLOW more efficient than BLOCKINGFLOW is that we can greedily push flow and do not need to rebalance at all. To obtain a linear running time, we also need to avoid having to sort the vertices of the graph topologically. This can be achieved as follows.

Algorithm: BIPARTITEBLOCKINGFLOW(G, μ, s, t)**input:** network $(G = (V, E), \mu = 1, s, t)$
induced by a bipartite graph**output:** blocking flow $f \leftarrow 0$ **for** $v \in V$:└ **unmark** v **while** $P \leftarrow \text{DFS}(G, s, t) \neq \emptyset$:└ AUGMENT($P, 1$)**return** f

Algorithm: DFS(G, u, t)**input:** graph $G = (V, E)$, vertices $u, t \in V$ **output:** u - t -path in G along unmarked vertices**if** $u = t$:└ **return** (t)**mark** u *(t is never marked)***for** $v \in \Gamma(u)$ with v **unmarked**:└ **if** $P \leftarrow \text{DFS}(G, v, t) \neq \emptyset$:└└ **return** (u) \oplus P **return** \emptyset

We show that this method of computing blocking flows is correct for flow networks.

Proposition 5.6. BIPARTITEBLOCKINGFLOW computes a blocking flow in time $\mathcal{O}(m)$.

Proof. We show that BIPARTITEBLOCKINGFLOW computes a blocking flow. The running time then follows, because every vertex is visited at most once and only if it is reachable from s .

For the sake of contradiction, assume that there is still an s - t -path P in $G \cap G_f$ upon termination of BIPARTITEBLOCKINGFLOW. Then, at least one vertex of $P \setminus \{s, t\}$ must be marked, otherwise DFS would have still found some s - t -path. Let $v \in V \setminus \{s, t\}$ the last such vertex along P . Then, a v - t -path without marked vertices exists. This means that DFS either did not visit all neighbors of v , or $t \in \Gamma(v)$. In either case, DFS must have found a v - t -path at some point, which means that BIPARTITEBLOCKINGFLOW must have augmented flow over some path P' that contains v . From $\mu = 1$ it follows that P and P' do not share arcs. Since both P and P' contain v , this implies that v has both in- and out-degree greater one. This is a contradiction to the fact that the network is induced by a bipartite graph. \square

With Proposition 3.4, this immediately gives a running time of $\mathcal{O}(nm)$. This is not yet an improvement over FORDFULKERSONMATCHING, but we can further strengthen the analysis.

Theorem 5.7. HOPCROFTKARP computes a maximum matching in bipartite graphs in time $\mathcal{O}(m\sqrt{n})$.

Proof. We have already seen that the length of a shortest s - t -path in G_f increases in every iteration of DINIC (Proposition 3.4). Consider the situation after iteration $\sqrt{n} + 2$. Let M be the bipartite matching in G corresponding to the current flow in H , and let M^* be a maximum matching in G . The symmetric difference $M \Delta M^*$ has at most one edge of M and at most one of M^* at every vertex. This allows to partition $M \Delta M^*$ into M -alternating cycles and paths. The number of M -augmenting paths in this partition is exactly $|M^*| - |M|$, since there cannot be any M^* -augmenting paths (Theorem 5.3). We also know that every M -augmenting path corresponds to an augmenting path in H_f , which, in turn, has length at least $\sqrt{n} + 2$. It follows that every M -augmenting path in $M \Delta M^*$ has length at least \sqrt{n} . Since these paths are disjoint, there can be at most \sqrt{n} such paths, and, hence, $|M^*| - |M| \leq \sqrt{n}$. Since the size of the matching increases in every iteration (Proposition 3.4), HOPCROFTKARP needs at most \sqrt{n} more iterations to reach a maximum cardinality matching. Overall, the number of iterations is bounded by $2\sqrt{n} + 2 = \mathcal{O}(\sqrt{n})$, and each iteration takes time $\mathcal{O}(m)$ by Proposition 5.6. \square

5.2 Characterization of Matchings

We turn to matchings in general (not necessarily bipartite) undirected graphs. We first give an algebraic characterization of graphs that admit perfect matchings. This characterization relies on the following

matrix similar to an adjacency matrix, but with variables as its entries. Recall that an *orientation* of an undirected graph is a directed graph that is obtained by assigning a direction to each edge.

Definition 5.8. For $\mathbf{x} = (x_e)_{e \in E}$ and an arbitrary orientation G' of G , the *Tutte matrix* $T(\mathbf{x}) = (t_{u,v})_{u,v \in V}$ is given by

$$t_{u,v} := \begin{cases} x_e, & \text{if } e = (u, v) \in G', \\ -x_e, & \text{if } e = (v, u) \in G', \\ 0, & \text{otherwise.} \end{cases}$$

It is a beautiful mathematical connection that the determinant of the Tutte matrix vanishes if and only if the underlying undirected graph has no perfect matching.

Theorem 5.9 (Tutte). An undirected graph G does not have a perfect matching if and only if all terms in the polynomial $\det T(\mathbf{x})$ cancel out (symbolically).

Proof. Let $V = \{1, \dots, n\}$ and let S_n denote the set of all permutations of V . By the Leibniz formula for determinants, we have

$$\det T(\mathbf{x}) = \sum_{\pi \in S_n} \text{sgn}(\pi) \prod_{v \in V} t_{v, \pi(v)}.$$

First assume that a perfect matching M exists in G . Consider the permutation $\pi(u) = v$ and $\pi(v) = u$ for all $\{u, v\} \in M$. For this and no other permutation we have

$$\prod_{v \in V} t_{v, \pi(v)} = \prod_{e \in M} (-x_e^2).$$

Therefore, this term does not cancel with any other term in $\det T(\mathbf{x})$.

Now, conversely, assume that $\det T(\mathbf{x}) \neq 0$. Let $S'_n := \{\pi \in S_n \mid \prod_{v \in V} t_{v, \pi(v)} \neq 0\}$ denote the set of all permutations that contribute to $\det T(\mathbf{x})$. We interpret every permutation $\pi \in S_n$ as a directed graph H_π with arcs (u, v) between every $u, v \in V$ with $v = \pi(u)$. Every vertex of H_π has in- and out-degree 1, and, hence, H_π consists of cycles – the permutation cycles of π . For all $\pi \in S'_n$ we have that H_π is a subgraph of $G' \cup \bar{G}'$, where G' is an arbitrary orientation of G .

We show that the contributions of all permutations $\pi \in S'_n$ that contain at least one odd cycle in H_π cancel out in $\det T(\mathbf{x})$. Let $r \in S_n$ be the permutation for which the first odd cycle in H_π (i.e., the cycle containing the smallest vertex) and in $H_{r \circ \pi}$ have opposite orientation and all other cycles are unchanged.

We first show that $\text{sgn}(\pi) = \text{sgn}(r \circ \pi)$, i.e., π can be translated to $r \circ \pi$ with an even number of swaps. Let v_1, \dots, v_{2k+1} be the first odd cycle in H_π . For $i = 1, \dots, k$, we swap (in this order) v_{2i-1} with v_{2k} and, analogously, v_{2i} with v_{2k+1} . Let $v_{-1} := v_{2k}$ and $v_0 := v_{2k+1}$. Now, for all $i \in \{1, \dots, k\}$, we ensured that $r(v_i) = v_{i-2}$, since v_{i-2} was first swapped with v_{2k} or v_{2k-1} and then (from there) with v_i . Hence, we have

$$r(\pi(v_i)) = r(v_{i+1}) = v_{i-1},$$

i.e., the first odd cycle is flipped. As we used an even number of swaps, we have $\text{sgn}(\pi) = \text{sgn}(r(\pi))$.

Since we changed the orientation of an odd cycle, the sign of an odd number of entries of $\prod_{v \in V} t_{v, \pi(v)}$ changed. In particular, $\pi \in S'_n$ implies $r \circ \pi \in S'_n$. Consequently, $\prod_{v \in V} t_{v, \pi(v)} = -\prod_{v \in V} t_{v, r(\pi(v))}$ and the contributions of π and $r \circ \pi$ cancel out. Furthermore, $r \circ r \circ \pi = \pi$, therefore we can pair the elements of S'_n with odd cycles, such that the contributions of each pair to $\det T(\mathbf{x})$ cancels out.

Now, since $\det T(\mathbf{x}) \neq 0$, a permutation with only even cycles must exist. But then we can construct a perfect matching by taking every second edge of each cycle (ignoring orientations). \square

While the above characterization is beautiful, it is often difficult to apply. We now derive a much more useful characterization. To this end, we need some definitions.

Definition 5.10. Let $G = (V, E)$ be an undirected graph and $X \subseteq V$. Then $q(X) := q_G(X)$ denotes the number of connected components with an odd number of vertices in the induced subgraph $G[V \setminus X]$.

In order to inductively characterize graphs that admit perfect matchings, we need a way of splitting a graph into smaller parts that allows us to relate the existence of a perfect matching in the original graph to the existence of a perfect matching in every part. Sets of vertices $X \subseteq V$ with $q(X) = |X|$ allow us to do this nicely, because they split the graph into parts in such a way that every part of odd cardinality in $V \setminus X$ needs to have exactly one edge to X in every perfect matching. All other edges need to be within parts. This justifies the following definitions.

Definition 5.11. A *barrier* is a set of vertices $X \subseteq V$ with $q(X) = |X|$.

Definition 5.12. An undirected graph $G = (V, E)$ is called *factor critical* if $G - v$ has a perfect matching for all $v \in V$.

The following lemma will be useful.

Lemma 5.13. For all undirected graphs $G = (V, E)$ and all $X \subseteq V$ it holds that $|V| \pm q(X) \pm |X|$ is even.

Proof. Let \mathcal{Z} be the partition of V into the vertex sets of the connected components of $G - X$ and the set X . Then, $q(X) \pm |X|$ is even if and only if \mathcal{Z} contains an even number of odd (cardinality) sets. Since \mathcal{Z} is a partition of V , this is the case if and only if $|V|$ is even. It follows that $|V| \pm q(X) \pm |X|$ is always even. \square

We now present a characterization, similar to the marriage theorem for bipartite graphs.

Theorem 5.14 (Tutte). An undirected graph $G = (V, E)$ has a perfect matching if and only if

$$q(X) \leq |X| \quad \forall X \subseteq V. \quad (5.1)$$

Proof. If a perfect matching M exists in G , then, for every $X \subseteq V$, every odd component of $G - X$ must have a matching edge to a distinct element of X . Hence, $q(X) \leq |X|$.

Conversely, assume that the Tutte condition (5.1) holds. We show, by induction over $n = |V|$, that a perfect matching exists. For $n = 2$, by the Tutte condition, we have $q(\emptyset) = 0$ so there must be an edge between the two vertices. Thus, a perfect matching exists.

Now consider the case $n > 2$. The Tutte condition implies $q(\emptyset) = 0$, hence \emptyset is a barrier and n must be even. We can therefore choose a barrier X of maximum size. By definition, $G - X$ has $|X|$ odd connected components. Let $x \in V$ be a vertex of an even component. By the Tutte-condition, we have $q(X \cup \{x\}) \leq |X| + 1$, but it also holds that $q(X \cup \{x\}) \geq q(X) + 1 = |X| + 1$. Hence, $X \cup \{x\}$ is a barrier – a contradiction to the maximality of X . Hence, $G - X$ cannot have even components.

We claim that every connected component of $G - X$ is factor critical. Let C be such a component and $v \in C$ be one of its vertices. Assume that $C - v$ did not have a perfect matching. By induction, it then follows that the Tutte condition is violated for the graph $C - v$. In other words, there must be a set of vertices $Y \subseteq C - v$ with $q_{C-v}(Y) > |Y|$. Since $G - X$ has no even components, C is odd and $C - v$ is even. By Lemma 5.13, $q_{C-v}(Y) - |Y|$ must be even as well. It follows that

$$q_{C-v}(Y) \geq |Y| + 2.$$

Since X , Y , and $\{v\}$ are pairwise disjoint, it follows that

$$\begin{aligned} q(X \cup Y \cup \{v\}) &= q(X) - 1 + q_C(Y \cup \{v\}) \\ &= |X| - 1 + q_{C-v}(Y) \\ &\geq |X| - 1 + |Y| + 2 \\ &= |X \cup Y \cup \{v\}|. \end{aligned}$$

By the Tutte condition, we also have $q(X \cup Y \cup \{v\}) \leq |X \cup Y \cup \{v\}|$. Therefore, $X \cup Y \cup \{v\}$ is a barrier – a contradiction to maximality of X . We conclude that every component of $G - X$ is factor critical.

Now consider the bipartite graph $G' = (X \cup \mathcal{Z}, E')$, where \mathcal{Z} is the set of components of $G - X$ and E' contains an edge between $x \in X$ and $C \in \mathcal{Z}$ if and only if $x \in \Gamma_G(C)$. If there is no perfect matching in G' , then, by the marriage theorem (see *Algorithmic Discrete Mathematics*), there exists $\mathcal{A} \subseteq \mathcal{Z}$ with $|\Gamma_{G'}(\mathcal{A})| < |\mathcal{A}|$. Since the components in \mathcal{Z} are not connected in $G - X$, the set $\Gamma_{G'}(\mathcal{A})$ isolates every component in \mathcal{A} from the rest of G . Since every $C \in \mathcal{A} \subseteq \mathcal{Z}$ is an odd component of $G - X$, it follows that $q_G(\Gamma_{G'}(\mathcal{A})) \geq |\mathcal{A}|$. We conclude that $q_G(\Gamma_{G'}(\mathcal{A})) > |\Gamma_{G'}(\mathcal{A})|$, which is a contradiction to the Tutte condition.

Thus, there must be a perfect matching in G' . Since every $C \in \mathcal{Z}$ is factor critical, we can transform this perfect matching into a perfect matching M^* of G : For every edge (x, C) in the matching, let $v \in C$ with $\{x, v\} \in E$ be a vertex of C that shares an edge with x in G . We include this edge in M^* . Because C is factor critical, there is a perfect matching in $C - v$, which we can also include in M^* . \square

While Tutte's theorem gives a characterization for the existence of perfect matchings, it does not apply for graphs that do not admit perfect matchings. We now extend this characterization to capture maximum matchings in general. Note that the following Theorem subsumes Tutte's theorem for perfect matchings, since then $|M^*| = n/2$ and $q(X) \leq |X|$, $q(\emptyset) = 0$.

Theorem 5.15 (Berge-Tutte). For all undirected graphs $G = (V, E)$ it holds that

$$2|M^*| + \max_{X \subseteq V} \{q(X) - |X|\} = |V|,$$

where M^* denotes a maximum cardinality matching in G .

Proof. For every $X \subseteq V$, every matching M must leave at least $q(X) - |X|$ vertices exposed, since every odd component of $G - X$ either has an exposed vertex or a matching edge to a vertex of X . It follows that $2|M| + q(X) - |X| \leq |V|$.

Now let

$$k := \max_{X \subseteq V} \{q(X) - |X|\}$$

and $H = (V', E')$ be the graph G with k additional vertices that each have edges to all other vertices. Assume that H has a perfect matching M_H^* . Then, for a maximum matching M^* in G , it holds that $|M^*| \geq |M_H^*| - k$, since we can construct a matching in G from M_H^* by removing all matching edges that contain a vertex of $V' \setminus V$. Since M_H^* is perfect, we have

$$2|M^*| + k \geq 2|M_H^*| - k = |V'| - k = |V|.$$

It remains to show that H always has a perfect matching. Assume this was not the case. By Tutte's theorem (Theorem 5.14), it follows that there is a set $Y \subseteq V'$ with $q_H(Y) > |Y|$. By Lemma 5.13, $|V'| = |V| + k$ is even. Hence, $q_H(\emptyset) = 0$ and thus $Y \neq \emptyset$. It follows that $q_H(Y) > |Y| \geq 1$, hence $H - Y$ consists of multiple components. But then Y must contain all vertices in $V' \setminus V$, since these are connected to all vertices. We obtain a contradiction to the definition of k (with $X = Y \cap V$):

$$q_G(Y \cap V) = q_H(Y) > |Y| = |Y \cap V| + k. \quad \square$$





Figure 5.3: Illustration of the operation $\text{AUGMENT}(u, v)$. Left: alternating tree before the augmentation; right: situation after the augmentation. Note that the alternating tree is *not* maintained.





Figure 5.4: Illustration of the operation $\text{EXTENDTREE}(u, v)$. Left: alternating tree before the extension; right: situation after the extension. Note that the alternating tree is maintained.

Algorithm: $\text{EDMONDSPERFECTBIPARTITE}(G)$

input: bipartite graph $G = (V, E)$

output: perfect matching in G

$M \leftarrow \emptyset$

$\text{MAKETREE}(\emptyset)$

while $\exists \{u, v\} \in E$ with $u \in T_{\text{even}}$ **and** $v \notin T$:

- if** $v \notin M$:
 - $\text{AUGMENT}(u, v)$
 - $\text{MAKETREE}(\emptyset)$
- else**
 - $\text{EXTENDTREE}(u, v)$

if $T = \emptyset$:

- $\text{return } M$

else

- fail

Algorithm: $\text{MAKETREE}(R)$

if $\exists r \notin R$ with $r \notin M$:

- $T \leftarrow (\{r\}, \emptyset)$
- $T_{\text{even}} \leftarrow \{r\}$
- $T_{\text{odd}} \leftarrow \emptyset$

else

- $T \leftarrow \emptyset, T_{\text{even}} \leftarrow \emptyset, T_{\text{odd}} \leftarrow \emptyset$

Algorithm: $\text{AUGMENT}(u, v)$

$P \leftarrow r\text{-}u\text{-path in } T + \{u, v\}$

$M \leftarrow M \Delta P$

Algorithm: $\text{EXTENDTREE}(u, v)$

$\{z\} \leftarrow \Gamma_M(v)$

$T \leftarrow T + \{u, v\} + \{v, z\}$

$T_{\text{odd}} \leftarrow T_{\text{odd}} \cup \{v\}$

$T_{\text{even}} \leftarrow T_{\text{even}} \cup \{z\}$

We show formally that this algorithm indeed always finds a perfect matching in bipartite graphs or concludes that none exists.

Proposition 5.18. $\text{EDMONDSPERFECTBIPARTITE}$ is correct.

Proof. We prove that $\text{EDMONDSPERFECTBIPARTITE}$ maintains the invariant that M is a matching and T is an alternating tree with root r (or $T = \emptyset$). This is the case initially, since every vertex is M -exposed, $M = \emptyset$, and $T = (\{r\}, \emptyset)$. Now consider an iteration for edge $\{u, v\} \in E$ with $u \in T_{\text{even}}$ and $v \in V \setminus T$. First assume that v is M -exposed and let P_{ru} be the unique r - u -path in T . Then, the path $P = P_{ru} \oplus (v)$ is an M -augmenting path, since r and v are M -exposed and P_{ru} has even length. Hence, we may augment and M remains a matching. Afterwards, we either found a perfect matching and set $T = \emptyset$, or we set $T = (\{r'\}, \emptyset)$ for some new vertex $r' \in V$, hence the invariant is maintained. If v is M -covered, we

extend the tree T with the edges $\{u, v\}$ and $\{v, z\}$ with $\{u, v\} \notin M$ and $\{v, z\} \in M$. Hence, T remains an alternating tree and the invariant is maintained.

If the algorithm terminates successfully, then $T = \emptyset$ and, by the invariant, M is a matching. The only way that T can be set to \emptyset is if no M -exposed vertex exists anymore in `MAKETREE`. But then M is a perfect matching.

Finally, if the algorithm fails, then no edge $\{u, v\} \in E$ exists with $u \in T_{\text{even}}$ and $v \in V \setminus T$. Also, there cannot be any edge in $G[T_{\text{even}}]$, since such an edge would close an odd cycle, which would be a contradiction with G being bipartite. Hence, every edge in $G[T] \cup \delta_G(T)$ must contain a vertex of T_{odd} . But this means that T_{odd} is a vertex cover of $G[T] \cup \delta_G(T)$. By Lemma 5.17, there is no perfect matching, hence the algorithm legitimately fails. \square

To generalize `EDMONDSPERFECTBIPARTITE` to general graphs, we need to address edges between vertices in T_{even} . As observed above, such edges close a unique odd cycle in T . Restricted to this odd cycle, the current matching covers all vertices except for the unique vertex of the cycle closest to the root (see Figure 5.5). We can formalize this as follows.

Definition 5.19. A *blossom* w.r.t. a matching M in G is a factor critical subgraph $C = (V_C, E_C) \subseteq G$ with $|M \cap E_C| = (|V_C| - 1)/2$. The (*blossom*) *basis* of C is the unique $(M \cap E_C)$ -exposed vertex of C .

The key idea now is to *contract* the corresponding blossom C whenever we find an edge between two vertices of T_{even} , i.e., to shrink it to a single vertex v_C and to replace the endpoints in C of all edges in $\delta(C)$ by v_C (see Figure 5.5). In the resulting graph, we eliminated the critical edge, and we can continue as before. Once we find a maximum matching in the contracted graph, we can easily recover a maximum matching in the original graph. This key insight is called the *blossom lemma*.

Lemma 5.20 (blossom lemma). Let C be a blossom with respect to matching M in G , and let G' , M' be the graph and the matching obtained from G , M , respectively, by contraction of C . Let further Q be an M -alternating path of even length in G from an M -exposed vertex to the basis b of C , such that Q does not use edges in C . Then, M is a maximum matching in G if and only if M' is a maximum matching in G' .

Proof. Let M' not be a maximum matching in G' . We can extend every matching in G' to a matching in G with $\frac{|C|-1}{2}$ additional edges, since, by definition, C is factor critical. For a maximum matchings M'^* in G' and M^* in G it therefore holds that

$$|M^*| \geq |M'^*| + \frac{|C|-1}{2} > |M'| + \frac{|C|-1}{2} = |M|.$$

This means that M is not a maximum matching in G .

Conversely, let M not be a maximum matching in G . We construct the matching $N = M \Delta Q$. Since Q has even length, this matching has the same cardinality as M and is therefore not maximum. If b is M -covered, then it is N -exposed. Otherwise, we may instead use $Q = \emptyset$, which again implies that b is N -exposed. Because N is not maximum, there is an N -augmenting path P in G (Theorem 5.3). Since C contains only a single N -exposed vertex, P has (at least) one of its ends outside of C . Let \bar{P} be that maximal subpath of P that contains this end v and contains no vertex of C in its interior. If we contract C , as b is N -exposed, the resulting vertex is N' -exposed, where N' denotes N after contracting C . Let Q' be the path in G' obtained from Q by contracting C . In G' , the path \bar{P} corresponds to an N' -augmenting path. Hence, N' is not maximum in G' . But $M' = N' \Delta Q'$ has the same cardinality as N' , since Q has even length. Thus, M' is not maximum either. \square




Figure 5.5: Illustration of the operation $\text{CONTRACT}(u, v)$. Left: alternating tree before the extension; right: situation after the extension. Note that the alternating tree is maintained.

With this, we can extend $\text{EDMONDSPERFECTBIPARTITE}$ to general graphs, as follows. The operation CONTRACT is illustrated in Figure 5.5.

Algorithm: $\text{EDMONDSPERFECT}(G)$

input: undirected graph $G = (V, E)$

output: perfect matching in G

$M \leftarrow \emptyset$

$\text{MAKETREE}(\emptyset)$

$G' = (V', E') \leftarrow G$

while $\exists \{u, v\} \in E'$ with $u \in T_{\text{even}}$ and $v \notin T_{\text{odd}}$:

if $v \notin T$ and $v \notin M$:

$\text{AUGMENT}(u, v)$

$M \leftarrow \text{EXPAND}(M, \emptyset)$

$G' \leftarrow G$

$\text{MAKETREE}(\emptyset)$

else if $v \notin T$ and $v \in M$:

$\text{EXTENDTREE}(u, v)$

else

$\text{CONTRACT}(u, v)$

if $T = \emptyset$:

return M

else

fail

Algorithm: $\text{CONTRACT}(C, G = (V, E))$

$E_c \leftarrow \{e \in E \mid e \cap C = \emptyset\}$

for $v \in \Gamma(C)$:

$E_c \leftarrow E_c \cup \{\{v_C, v\}\}$

$V_c \leftarrow (V \setminus C) \cup \{v_C\}$

return (V_c, E_c)

Algorithm: $\text{CONTRACT}(u, v)$

$C \leftarrow (v\text{-}u\text{-path in } T) \oplus (v)$

$G' \leftarrow \text{CONTRACT}(C, G')$

$M \leftarrow \text{CONTRACT}(C, M)$

$T \leftarrow \text{CONTRACT}(C, T)$

Algorithm: $\text{EXPAND}(M, R)$

while $\exists v_C \in V' \setminus R$ with v_C contracted:

$M' \leftarrow \text{max. matching in } C \cup \Gamma_M(v_C)$

$M \leftarrow (M \setminus \delta_M(v_C)) \cup M'$

$V' \leftarrow V' - v_C + C$

return M

We show that this algorithm works in general to find perfect matchings.

Proposition 5.21. EDMONDSPERFECT is correct.

Proof. Again, the algorithm maintains the invariant that M is a matching in G' and that T is an alternating tree with root r . It is maintained during augmentations and extensions of T , as in the proof of Proposition 5.18. It is also maintained when a blossom C is contracted, since the basis of C is the unique vertex of C that may be part of an edge in $\delta_{G'}(C) \cap M$ (unless it coincides with r).

We show a second invariant: Throughout the algorithm, the expansion of M to G is maximum in G if and only if M is maximum in G' . Before every contraction of a blossom C with basis b , the r - b -path Q in T is an M -alternating path in G' of even length, r is M -exposed, and b is M -covered (or $r = b$). By the blossom lemma (Lemma 5.20), the invariant is maintained. During EXTENDTREE the invariant is trivially maintained, since M and G' are unchanged. As soon as we augment, we set $G' = G$, hence the invariant is maintained.

If the algorithm terminates successfully, we have $T = \emptyset$, which means that the last operation was `MAKETREE`, there is no M -exposed vertex, and $G' = G$. By the first invariant, M is a perfect matching in G .

If the algorithm fails, all edges in $G'[T] \cup \delta_{G'}(T)$ have one end in T_{odd} . But then T_{odd} is a vertex cover of $G'[T] \cup \delta_{G'}(T)$. By Lemma 5.17, it follows that G' does not have a perfect matching and, by the second invariant, neither does G . Hence, the algorithm legitimately fails. \square

Finally, we extend the algorithm further to settings where a perfect matching need not exist. The key idea is to store completed trees, and to start new trees, until all exposed vertices are part of some tree. This yields Edmonds' famous blossom algorithm, which is considered to be one of the most beautiful algorithms in combinatorial optimization, due to the various key ideas that it combines in elegant fashion.

Algorithm: EDMONDS(G)

input: undirected graph $G = (V, E)$
output: maximum cardinality matching in G

$M \leftarrow \emptyset, R \leftarrow \emptyset, \text{MAKETREE}(\emptyset), G' = (V', E') \leftarrow G$

while $\exists r \notin R$ with $r \notin M$:

- `MAKETREE`(R)
- while** $\exists \{u, v\} \in E'$ with $u \in T_{\text{even}}$ and $v \notin T_{\text{odd}} \cup R$:
 - if** $v \notin T$ and $v \notin M$:
 - `AUGMENT`(u, v)
 - $M \leftarrow \text{EXPAND}(M, R)$
 - $G' \leftarrow G, T \leftarrow \emptyset$
 - break** inner loop
 - else if** $v \notin T$ and $v \in M$:
 - `EXTENDTREE`(u, v)
 - else**
 - `CONTRACT`(u, v)
- $R \leftarrow R \cup T$

return `EXPAND`(M, \emptyset)

We show that Edmonds' algorithm computes a maximum matching in general graphs.

Theorem 5.22. EDMONDS is correct.

Proof. We show that M is a maximum matching in G' upon termination of the algorithm. As for `EDMONDSPERFECT`, it follows that M expanded to G is also maximum.

We first observe that, during each update of R , the vertices of an alternating tree are added to R . These vertices remain unchanged afterwards, therefore the graph $G' = (V', E')$ contains all these trees T_1, T_2, \dots at the end. Every tree contains exactly one exposed vertex, namely its root. Let V_{even} denote the set of vertices in T_{even} of all these trees $T \in \{T_1, T_2, \dots\}$, and let V_{odd} be defined analogously. There are no edges between vertices in T_{even} of a single tree, otherwise a blossom could have been contracted. Also there are no edges between vertices in V_{even} of different trees, otherwise the first of these trees could have been extended further. Now, if we remove all vertices of V_{odd} from G' , we obtain $|V_{\text{even}}|$ components consisting of a single vertex each. If k is the number of trees, then $|V_{\text{even}}| = |V_{\text{odd}}| + k$. Hence,

$q_{G'}(V_{\text{odd}}) = |V_{\text{even}}| = |V_{\text{odd}}| + k$. Since every M -exposed vertex of G' is contained in some tree, and since every tree contains exactly one M -exposed vertex, we have

$$|M| = \frac{|V'| - k}{2}.$$

By the Theorem of Berge-Tutte (Theorem 5.15) for G' , it holds that

$$|M'^*| = \frac{|V'| - \max_{X \subseteq V} \{q_{G'}(X) - |X|\}}{2} \leq \frac{|V'| - q_{G'}(V_{\text{odd}}) + |V_{\text{odd}}|}{2} = \frac{|V'| - k}{2} = |M|,$$

hence M' is maximum in G' . □

To implement EDMONDS efficiently, we need to maintain contracted components and to maintain $\delta(T_{\text{even}})$. The former can be achieved using a union-find data structure (see lecture *Algorithmic Discrete Mathematics*), and the latter can be achieved by explicitly maintaining a list L that contains all edges in $\delta(T_{\text{even}})$ that have not yet been considered.

Algorithm: EDMONDS(G) (implementation)

input: undirected graph $G = (V, E)$

output: maximum cardinality matching in G

$M \leftarrow \emptyset, R \leftarrow \emptyset$

MAKETREE(\emptyset)

while $\exists r \notin R$ with $r \notin M$:

MAKETREE(R)

$L \leftarrow \{(r, r_v) \mid v \in \Gamma_G(r)\}$

while $\exists (u, v) \in L$:

$L \leftarrow L \setminus \{(u, v)\}$ (extract first)

if FIND(u) \neq FIND(v) **and** $v \notin T_{\text{odd}}$:

if $v \notin T$ **and** $v \notin M$:

AUGMENT(u, v)

$M \leftarrow \text{EXPAND}(M, R)$

else if $v \notin T$ **and** $v \in M$:

EXTENDTREE(u, v)

else

CONTRACT(u, v)

$R \leftarrow R \cup T$

return M

Algorithm: EXTENDTREE(u, v) (impl.)

$\{z\} \leftarrow \Gamma_M(v)$

$T \leftarrow T + \{u, v\} + \{v, z\}$

$T_{\text{odd}} \leftarrow T_{\text{odd}} \cup \{v\}$

$T_{\text{even}} \leftarrow T_{\text{even}} \cup \{z\}$

for $w \in \{r_{w'} \mid w' \in \Gamma_G(z)\} \setminus \{v\}$:

CONTRACT(z, w)

Algorithm: CONTRACT(u, v) (impl.)

$C \leftarrow (v\text{-}u\text{-path in } T) \oplus (v)$

for $\{a, b\} \in C$:

UNION(a, b)

Algorithm: EXPAND(M, R)

for $v \in M$ with $v \notin R$:

$r_v \leftarrow v$

We remark (without proof) that a union-find data structure can be implemented efficiently. Recall that $\alpha(m)$ is the inverse Ackermann function that grows extremely slowly (e.g., $\alpha(9876!) = 5$).

Proposition 5.23. A union-find data structure can be implemented such that every sequence of $\mathcal{O}(m)$ operations takes time $\mathcal{O}(m \cdot \alpha(m))$.

With this, we obtain a very efficient implementation of Edmonds' algorithm.

Theorem 5.24. EDMONDS can be implemented with a running time of $\mathcal{O}(nm \cdot \alpha(m))$.

Proof. Since every augmentation increases the size of the matching, there can be at most $n/2 = \mathcal{O}(n)$ augmentations. We consider the operations between two augmentations. Obviously, every edge is added to L at most once. It follows that we have $\mathcal{O}(m)$ FIND and $\mathcal{O}(n)$ UNION operations. The total running time of these operations is $\mathcal{O}(m \cdot \alpha(m))$ if we use a suitable union-find data structure, by Proposition 5.23.

We have $\mathcal{O}(n)$ EXTENDTREE operations that consider each edge at most twice overall. The total running time for this is $\mathcal{O}(m)$.

In every contraction, the number of vertices is reduced by the size of the contracted blossom minus 1 and the running time required is linear in the size of the blossom. The total running time for contractions therefore is $\mathcal{O}(n)$ plus the running time of the involved UNION operations.

Overall, the running time is $\mathcal{O}(nm \cdot \alpha(m))$, assuming that we implement trees by storing predecessors for each node. \square