

# In defense of the Simplex Algorithm’s worst-case behavior\*

Yann Disser and Martin Skutella

TU Berlin, Institut für Mathematik,  
disser@math.tu-berlin.de, martin.skutella@tu-berlin.de

**Abstract.** In the early 1970s, by work of Klee and Minty (1972) and Zadeh (1973), the Simplex Method, the Network Simplex Method, and the Successive Shortest Path Algorithm have been proved guilty of exponential worst-case behavior (for certain pivot rules). Since then, the common perception is that these algorithms can be fooled into investing senseless effort by ‘bad instances’ such as, e. g., Klee-Minty cubes. This paper promotes a more favorable stance towards the algorithms’ worst-case behavior. We argue that the exponential worst-case performance is not necessarily a senseless waste of time, but may rather be due to the algorithms performing meaningful operations and solving difficult problems on their way. Given one of the above algorithms as a black box, we show that using this black box, with polynomial overhead and a limited interface, we can solve any problem in NP. This also allows us to derive NP-hardness results for some related problems.

## 1 Introduction

Dantzig’s Simplex Method [6, 7] is probably the most important mathematical contribution to solving real-world optimization problems. Empirically, it belongs to the most efficient methods for solving linear programs. From a worst-case point of view, however, Klee and Minty gave a disillusioning answer to the question ‘How good is the simplex algorithm?’ in their eponymous article [16]. For their class of linear programming problems on deformed  $d$ -dimensional cubes, nowadays known as Klee-Minty cubes, the Simplex Method with Dantzig’s pivot rule requires  $2^d - 1$  iterations. Similar results are known for many other popular pivot rules; see, e. g., Amenta and Ziegler [2] and the recent work of Friedmann [9]. On the other hand, by the work of Khachian [14, 15] and later Karmarkar [13], it is known since the late 1970s and early 1980s that linear programs can be solved in polynomial time. We refer to standard textbooks, such as, e. g., Dantzig [7], Schrijver [20], and Matoušek and Gärtner [18], for a thorough treatment of linear programming and the simplex method.

Minimum-cost flow problems form a class of linear programs featuring a particularly rich combinatorial structure allowing for numerous specialized algorithms. The first such algorithm is Dantzig’s Network Simplex Method [5, 7]

---

\* This work was supported in part by the DFG Research Center MATHEON in Berlin and by the Alexander von Humboldt-Foundation.

which is an interpretation of the general Simplex Method applied to this class of problems. One of the simplest and most basic algorithms for minimum-cost flow problems is the Successive Shortest Path Algorithm which iteratively augments flow along paths of minimum cost in the residual network [4, 11]. According to Ford and Fulkerson, the underlying theorem stating that such an augmentation step preserves optimality “may properly be regarded as the central one concerning minimal cost flows” [8]. Zadeh [24] presented a family of instances forcing the Successive Shortest Path Algorithm and also the Network Simplex Method with Dantzig’s pivot rule into exponentially many iterations. On the other hand, Tardos [23] proved that minimum-cost flows can be computed in strongly polynomial time, and Orlin [19] gave a polynomial variant of the Network Simplex Method. We refer to standard textbooks, such as, e. g., Ford and Fulkerson [8], Ahuja, Magnanti, and Orlin [1], and Korte and Vygen [17], for a thorough treatment of minimum-cost flow algorithms.

Spielman and Teng [22] developed the concept of smoothed analysis in order to explain the practical efficiency of the Simplex Method despite its poor worst-case behavior. They showed that worst-case instances forcing the Simplex Method into exponentially many iterations are very fragile. More precisely, a slight random perturbation of a given instance is sufficient to guarantee expected polynomial running time of the Simplex Method. Brunsch et al. [3] applied smoothed analysis to explain the discrepancy between the poor worst-case behavior of the Successive Shortest Path Algorithm and its good empirical behavior.

In this paper, we consider the primal (Network) Simplex Method together with Dantzig’s pivot rule, which selects the nonbasic variable with the most negative reduced cost. We refer to this variant of the (Network) Simplex Method as the *(Network) Simplex Algorithm*.

*Our contribution.* We argue that the exponential worst-case running time of the (Network) Simplex Algorithm and the Successive Shortest Path Algorithm is not purely a waste of time. While these algorithms sometimes take longer than necessary to reach their primary objective (namely to find an optimum solution to a particular linear program), they collect meaningful information on their detours and implicitly solve difficult problems. To make this statement more precise, we introduce the following definition.

**Definition 1.** *An algorithm given by a Turing machine  $T$  implicitly solves a decision problem  $\mathcal{P}$  if, for a given instance  $I$  of  $\mathcal{P}$ , it is possible to compute an input  $I'$  for  $T$ , a cell  $C$ , and a symbol  $\sigma$  in polynomial time, such that  $I$  is a yes-instance if and only if at some point during its execution on input  $I'$  the Turing machine  $T$  writes symbol  $\sigma$  in cell  $C$  of its tape.*

This definition turns out to be sufficient for our purposes. We remark, however, that slightly more general definitions involving constantly many symbols, tape cells, and states of the Turing machine  $T$  also seem reasonable.

Notice that an algorithm that implicitly solves a particular NP-hard decision problem implicitly solves all problems in NP. We call such algorithms NP-*mighty*.

**Definition 2.** *An algorithm is NP-mighty if it implicitly solves every decision problem in NP.*

It is clear from the definitions that, unless  $P=NP$ , no polynomial-time algorithm is NP-mighty. Notice that these definitions have been formulated with some care in an attempt to distinguish ‘clever’ exponential-time algorithms from those that rather ‘waste time’ on less meaningful operations. For example, a Turing machine that receives a number  $n \in \mathbb{N}$  in binary encoding as input and then counts from  $n$  down to 0 is *not* NP-mighty.

**Theorem 1.** *The Simplex Algorithm, the Network Simplex Algorithm, and the Successive Shortest Path Algorithm are NP-mighty.*


We prove this theorem by showing that these algorithms implicitly solve the NP-complete PARTITION problem (cf. [10]). To this end, we show how to turn a given instance of PARTITION into a minimum-cost flow network with a distinguished arc  $e$  in polynomial time, such that the Network Simplex Algorithm (or the Successive Shortest Path Algorithm) augments flow along arc  $e$  in one of its iterations if and only if the PARTITION instance has a solution. Under the mild assumption that in an implementation of the Network Simplex Algorithm or the Successive Shortest Path Algorithm fixed memory cells are used to store the flow variables of arcs, this implies that these algorithms implicitly solve PARTITION in terms of Definition 1.

The core of our network construction is a family of counting gadgets on which these minimum-cost flow algorithms take exponentially many iterations. These counting gadgets are, in some sense, simpler than Zadeh’s ‘bad networks’ [24]. By slightly perturbing the costs of the arcs according to the values of a given PARTITION instance, we manage to force the considered minimum-cost flow algorithms into enumerating all possible solutions.

*Outline.* After establishing some minimal notation in Section 2, we proceed to proving Theorem 1 for the Successive Shortest Path Algorithm in Section 3. In Section 4, we adapt the construction for the Network Simplex Algorithm. Finally, Section 5 highlights some of the implications of our results for related problems.

## 2 Preliminaries

In the following sections we show that the Successive Shortest Path Algorithm and the Network Simplex Algorithm implicitly solve the classical PARTITION problem. An instance of PARTITION is given by a vector of positive numbers  $\vec{a} = (a_1, \dots, a_n) \in \mathbb{Q}^n$  and the problem is to decide whether there is a subset



**Fig. 1.** Recursive definition of the counting gadget  $N_i^{\vec{v}}$  for the Successive Shortest Path Algorithm and  $\vec{v} \in \{\vec{a}, -\vec{a}\}$ . Arcs are labeled by their cost and capacity in this order. The cost of the shortest  $s_i$ - $t_i$ -path in iteration  $j = 0, \dots, 2^i - 1$  is  $j + \vec{v}_i^{[j]}$ .

$I \subseteq \{1, \dots, n\}$  with  $\sum_{i \in I} a_i = \sum_{i \notin I} a_i$ . This problem is well-known to be NP-complete (cf. [10]). Throughout this paper we consider an arbitrary fixed instance  $\vec{a}$  of PARTITION. Without loss of generality, we assume  $A := \sum_{i=1}^n a_i < 1/6$  and that all values  $a_i, i \in \{1, \dots, n\}$ , are multiples of  $\varepsilon$  for some constant  $\varepsilon > 0$ .

Let  $\vec{v} = (v_1, \dots, v_n) \in \mathbb{Q}^n$  and  $k \in \mathbb{N}$ , with  $k_j \in \{0, 1\}, j \in \mathbb{Z}_{\geq 0}$ , being the  $j$ -th bit in the binary representation of  $k$ , i. e.,  $k_j := \lfloor k/2^j \rfloor \bmod 2$ . We define  $\vec{v}_{i_1, i_2}^{[k]} := \sum_{j=i_1+1}^{i_2} (-1)^{k_{j-1}} v_j, \vec{v}_i^{[k]} := \vec{v}_{0, i}^{[k]}$ , and  $\vec{v}_{i, i}^{[k]} = 0$ .

The following characterization will be useful later.

**Proposition 1.** *The PARTITION instance  $\vec{a}$  admits a solution if and only if there is a  $k \in \{0, \dots, 2^n - 1\}$  for which  $\vec{a}_n^{[k]} = 0$ .*

### 3 Successive Shortest Path Algorithm

Consider a network  $N$  with a source node  $s$ , a sink node  $t$ , and non-negative arc costs. The Successive Shortest Path Algorithm starts with the zero-flow and iteratively augments flow along a minimum-cost  $s$ - $t$ -path in the current residual network, until a maximum  $s$ - $t$ -flow has been found. Notice that the residual network is a sub-network of  $N$ 's bidirected network, where the cost of a backward arc is the negative of the cost of the corresponding forward arc.

#### 3.1 A Counting gadget for the Successive Shortest Path Algorithm

In this section we construct a family of networks for which the Successive Shortest Path Algorithm takes an exponential number of iterations. Assume we have a

network  $N_{i-1}$  with source  $s_{i-1}$  and sink  $t_{i-1}$  which requires  $2^{i-1}$  iterations that each augment one unit of flow. We can obtain a new network  $N_i$  with only two additional nodes  $s_i, t_i$  for which the Successive Shortest Path Algorithm takes  $2^i$  iterations. To do this we add two arcs  $(s_i, s_{i-1}), (t_{i-1}, t_i)$  with capacity  $2^{i-1}$  and cost 0, and two arcs  $(s_i, t_{i-1}), (s_{i-1}, t_i)$  with capacity  $2^{i-1}$  and very high cost. The idea is that in the first  $2^{i-1}$  iterations one unit of flow is routed along the arcs of cost 0 and through  $N_{i-1}$ . After  $2^{i-1}$  iterations both the arcs  $(s_i, s_{i-1}), (t_{i-1}, t_i)$  and the subnetwork  $N_{i-1}$  are completely saturated and the Successive Shortest Path Algorithm starts to use the expensive arcs  $(s_i, t_{i-1}), (s_{i-1}, t_i)$ . Each of the next  $2^{i-1}$  iteration adds one unit of flow along the expensive arcs and removes one unit of flow from the subnetwork  $N_{i-1}$ .

We tune the cost of the expensive arcs to  $2^{i-1} - \frac{1}{2}$  which turns out to be just expensive enough (cf. Figure 1, with  $v_i = 0$ ). This leads to a particularly nice progression of the costs of shortest paths, where the shortest path in iteration  $j = 0, 1, \dots, 2^i - 1$  simply has cost  $j$ .

Our goal is to use this counting gadget to iterate over all candidate solutions for a PARTITION instance  $\vec{v}$  (we later use the gadget for  $\vec{v} \in \{\vec{a}, -\vec{a}\}$ ). Motivated by Proposition 1, we perturb the costs of the arcs in such a way that the shortest path in iteration  $j$  has cost  $j + \vec{v}_i^{[j]}$ . We achieve this by adding  $\frac{1}{2}v_i$  to the cheap arcs  $(s_i, s_{i-1}), (t_{i-1}, t_i)$  and subtracting  $\frac{1}{2}v_i$  from the expensive arcs  $(s_i, t_{i-1}), (s_{i-1}, t_i)$ . If the value of  $v_i$  is small enough, this modification does not affect the overall behavior of the gadget. The first  $2^{i-1}$  iterations now have an additional cost of  $v_i$  while the next  $2^{i-1}$  iterations have an additional cost of  $-v_i$ , which leads to the desired cost when the modification is applied recursively.

Figure 1 shows the recursive construction of our counting gadget  $N_n^{\vec{v}}$  that encodes the PARTITION instance  $\vec{v}$ . The following lemma formally establishes the crucial properties of the construction.

**Lemma 1.** *For  $\vec{v} \in \{\vec{a}, -\vec{a}\}$  and  $i = 1, \dots, n$ , the Successive Shortest Path Algorithm applied to network  $N_i^{\vec{v}}$  with source  $s_i$  and sink  $t_i$  needs  $2^i$  iterations to find a maximum  $s_i$ - $t_i$ -flow of minimum cost. In each iteration  $j = 0, 1, \dots, 2^i - 1$ , the algorithm augments one unit of flow along a path of cost  $j + \vec{v}_i^{[j]}$  in the residual network.*

*Proof.* We prove the lemma by induction on  $i$ , together with the additional property that after  $2^i$  iterations none of the arcs in  $N_{i-1}^{\vec{v}}$  carries any flow, while the arcs in  $N_i^{\vec{v}} \setminus N_{i-1}^{\vec{v}}$  are fully saturated. First consider the network  $N_0^{\vec{v}}$ . In each iteration where  $N_0^{\vec{v}}$  does not carry flow, one unit of flow can be routed from  $s_0$  to  $t_0$ . Conversely, when  $N_0^{\vec{v}}$  is saturated, one unit of flow can be routed from  $t_0$  to  $s_0$ . In either case the associated cost is 0. With this in mind, it is clear that on  $N_1^{\vec{v}}$  the Successive Shortest Path Algorithm terminates after two iterations. In the first, one unit of flow is sent along the path  $s_1, s_0, t_0, t_1$  of cost  $v_1 = \vec{v}_1^{[0]}$ . In the second iteration, one unit of flow is sent along the path  $s_1, t_0, s_0, t_1$  of cost  $-v_1 = \vec{v}_1^{[1]}$ . Afterwards, the arc  $(s_0, t_0)$  does not carry any flow, while all other arcs are fully saturated.

Now assume the claim holds for  $N_{i-1}^{\bar{v}}$  and consider network  $N_i^{\bar{v}}$ ,  $i > 1$ . Observe that every path using either of the arcs  $(s_i, t_{i-1})$  or  $(s_{i-1}, t_i)$  has a cost of more than  $2^{i-1} - 3/4$ . To see this, note that the cost of these arcs is bounded individually by  $\frac{1}{2}(2^i - 1 - v_i) > 2^{i-1} - 3/4$ , since  $|v_i| < A < 1/4$ . On the other hand, it can be seen inductively that the shortest  $t_{i-1}$ - $s_{i-1}$ -path in the bidirected network associated with  $N_{i-1}^{\bar{v}}$  has cost at least  $-2^{i-1} + 1 - A > -2^{i-1} + 3/4$ . Hence, using both  $(s_i, t_{i-1})$  and  $(s_{i-1}, t_i)$  in addition to a path from  $t_{i-1}$  to  $s_{i-1}$  incurs cost at least  $2^{i-1} - 3/4$ . By induction, in every iteration  $j < 2^{i-1}$ , the Successive Shortest Path Algorithm thus does not use the arcs  $(s_i, t_{i-1})$  or  $(s_{i-1}, t_i)$  but instead augments one unit of flow along the arcs  $(s_i, s_{i-1})$ ,  $(t_{i-1}, t_i)$  and along an  $s_{i-1}$ - $t_{i-1}$ -path of cost  $j + \bar{v}_{i-1}^{[j]} < 2^{i-1} - 3/4$  through the subnetwork  $N_{i-1}^{\bar{v}}$ . The total cost of this  $s_i$ - $t_i$ -path is  $v_i + (j + \bar{v}_{i-1}^{[j]}) = j + \bar{v}_i^{[j]}$ , since  $j < 2^{i-1}$ .

After  $2^{i-1}$  iterations, the arcs  $(s_i, s_{i-1})$  and  $(t_{i-1}, t_i)$  are both fully saturated, as well as (by induction) the arcs in  $N_{i-1}^{\bar{v}} \setminus N_{i-2}^{\bar{v}}$ , while all other arcs are without flow. Consider the residual network of  $N_{i-1}^{\bar{v}}$  at this point. If we increase the costs of the four residual arcs in  $N_{i-1}^{\bar{v}} \setminus N_{i-2}^{\bar{v}}$  by  $\frac{1}{2}(2^{i-1} - 1)$  and switch the roles of  $s_{i-1}$  and  $t_{i-1}$ , we obtain back the original subnetwork  $N_{i-1}^{\bar{v}}$ . The shift of the residual costs effectively makes every  $t_{i-1}$ - $s_{i-1}$ -path more expensive by  $2^{i-1} - 1$ , but does not otherwise affect the behavior of the network. We can thus use induction again to infer that in every iteration  $j = 2^{i-1}, \dots, 2^i - 1$  the Successive Shortest Path Algorithm augments one unit of flow along a path via  $s_i, t_{i-1}, N_{i-1}^{\bar{v}}, s_{i-1}, t_i$ . Accounting for the shift in cost by  $2^{i-1} - 1$ , we obtain that this path has a total cost of


$$(2^i - 1 - v_i) + (j - 2^{i-1} + \bar{v}_{i-1}^{[j-2^{i-1}]}) - (2^{i-1} - 1) = j + \bar{v}_i^{[j]},$$

where we used  $\bar{v}_{i-1}^{[j-2^{i-1}]} = \bar{v}_{i-1}^{[j]}$  and  $\bar{v}_{i-1}^{[j]} - v_i = \bar{v}_i^{[j]}$  for  $j \in [2^{i-1}, 2^i)$ . After  $2^i$  iterations the arcs in  $N_i^{\bar{v}} \setminus N_{i-1}^{\bar{v}}$  are fully saturated and all other arcs carry no flow.  $\square$

### 3.2 The Successive Shortest Path Algorithm implicitly solves PARTITION

We use the counting gadget of the previous section to prove Theorem 1 for the Successive Shortest Path Algorithm. Let  $G_{\text{ssp}}^{\bar{a}}$  be the network consisting of the two gadgets  $N_n^{\bar{a}}$ ,  $N_n^{-\bar{a}}$ , connected to a new source node  $s$  and a new sink  $t$  (cf. Figure 2). For both of the gadgets, we add the arcs  $(s, s_n)$  and  $(t_n, t)$  with capacity  $2^n$  and cost 0. We introduce one additional arc  $e$  (dashed in the figure) of capacity 1 and cost 0 from node  $s_0$  of gadget  $N_n^{\bar{a}}$  to node  $t_0$  of gadget  $N_n^{-\bar{a}}$ . Finally, we increase the costs of the arcs  $(s_0, t_0)$  in both gadgets from 0 to  $\frac{1}{5}\varepsilon$ . Recall that  $\varepsilon > 0$  is related to  $\bar{a}$  by the fact that all  $a_i$ 's are multiples of  $\varepsilon$ , i. e., a cost smaller than  $\varepsilon$  is insignificant compared to all other costs.

**Lemma 2.** *The Successive Shortest Path Algorithm on network  $G_{\text{ssp}}^{\bar{a}}$  augments flow along arc  $e$  if and only if the PARTITION instance  $\bar{a}$  has a solution.*



**Fig. 2.** Illustration of network  $G_{\text{ssp}}^{\vec{a}}$ . The subnetworks  $N_n^{\vec{a}}$  and  $N_n^{-\vec{a}}$  are advanced independently by the Successive Shortest Path Algorithm without using arc  $e$ , unless the PARTITION instance  $\vec{a}$  has a solution.

*Proof.* First observe that our slight modification of the cost of arc  $(s_0, t_0)$  in both gadgets  $N_n^{\vec{a}}$  and  $N_n^{-\vec{a}}$  does not affect the behavior of the Successive Shortest Path Algorithm. This is because the cost of any path in  $G$  is perturbed by at most  $\frac{2}{5}\epsilon$ , and hence the shortest path remains the same in every iteration. The only purpose of the modification is tie-breaking.

Consider the behavior of the Successive Shortest Path Algorithm on the network  $G_{\text{ssp}}^{\vec{a}}$  with arc  $e$  removed. In each iteration, the shortest  $s$ - $t$ -path goes via one of the two gadgets. By Lemma 1, each gadget can be in one of  $2^n + 1$  states and we number these states increasingly from 0 to  $2^n$  by the order of their appearance during the execution of the Successive Shortest Path Algorithm. The shortest  $s$ - $t$ -path through either gadget in state  $j = 0, \dots, 2^n - 1$  has a cost in the range  $[j - A, j + A]$ , and hence it is cheaper to use a gadget in state  $j$  than the other gadget in state  $j + 1$ . This means that after every two iterations both gadgets are in the same state.

Now consider the network  $G_{\text{ssp}}^{\vec{a}}$  with arc  $e$  put back. We show that, as before, if the two gadgets are in the same state before iteration  $2j$ ,  $j = 0, \dots, 2^n - 1$ , then they are again in the same state two iterations later. More importantly, arc  $e$  is used in iterations  $2j$  and  $2j + 1$  if and only if  $\vec{a}_n^{[j]} = 0$ . This proves the lemma since, by Proposition 1,  $\vec{a}_n^{[j]} = 0$  for some  $j < 2^n$  if and only if the PARTITION instance  $\vec{a}$  has a solution.

To prove our claim, assume that both gadgets are in the same state before iteration  $2j$ . Let  $P^+$  be the shortest  $s$ - $t$ -path that does not use any arc of  $N_n^{-\vec{a}}$ ,  $P^-$  be the shortest  $s$ - $t$ -path that does not use any arc of  $N_n^{\vec{a}}$ , and  $P$  be the shortest  $s$ - $t$ -path using arc  $e$ . Note that one of these paths is the overall short-

est  $s$ - $t$ -path. We distinguish two cases, depending on whether the arc  $(s_0, t_0)$  currently carries flow 0 or 1 in both gadgets.

If  $(s_0, t_0)$  carries flow 0, then  $P^+$ ,  $P^-$  use arc  $(s_0, t_0)$  in forward direction. Therefore, by Lemma 1, the cost of  $P^+$  is  $j + \bar{a}_n^{[j]} + \frac{1}{5}\varepsilon$ , while the cost of  $P^-$  is  $j - \bar{a}_n^{[j]} + \frac{1}{5}\varepsilon$ . On the other hand, path  $P$  follows  $P^+$  to node  $s_0$  of  $N_n^{\bar{a}}$ , then uses arc  $e$ , and finally follows  $P^-$  to  $t$ . The cost of this path is exactly  $j$ . If  $\bar{a}_n^{[j]} \neq 0$ , then one of  $P^+$ ,  $P^-$  is cheaper than  $P$ , and the next two iterations augment flow along paths  $P^+$  and  $P^-$ . Otherwise, if  $\bar{a}_n^{[j]} = 0$ , then  $P$  is the shortest path, followed in the next iteration by the path from  $s$  to node  $t_0$  of  $N_n^{-\bar{a}}$  along  $P^-$ , along arc  $e$  in backwards direction to node  $s_0$  of  $N_n^{\bar{a}}$ , and finally to  $t$  along  $P^+$ , for a total cost of  $j + \frac{2}{5}\varepsilon$ .

If  $(s_0, t_0)$  carries flow 1, then  $P^+$ ,  $P^-$  use arc  $(s_0, t_0)$  in backward direction. By Lemma 1, the cost of  $P^+$  is  $j + \bar{a}_n^{[j]} - \frac{1}{5}\varepsilon$ , while the cost of  $P^-$  is  $j - \bar{a}_n^{[j]} - \frac{1}{5}\varepsilon$ . On the other hand, path  $P$  follows  $P^+$  to node  $s_0$  of  $N_n^{\bar{a}}$ , then uses arc  $e$ , and finally follows  $P^-$  to  $t$ . The cost of this path is  $j - \frac{2}{5}\varepsilon$ . If  $\bar{a}_n^{[j]} \neq 0$ , then one of  $P^+$ ,  $P^-$  is cheaper than  $P$ , and the next two iterations augment flow along paths  $P^+$  and  $P^-$ . Otherwise, if  $\bar{a}_n^{[j]} = 0$ , then  $P$  is the shortest path, followed in the next iteration by the path from  $s$  to node  $t_0$  of  $N_n^{-\bar{a}}$  along  $P^-$ , along arc  $e$  in backwards direction to node  $s_0$  of  $N_n^{\bar{a}}$ , and finally to  $t$  along  $P^+$ , for a total cost of  $j$ .  $\square$

We assume that a single cell of the Turing machine corresponding to the Successive Shortest Path Algorithm can be used to distinguish whether arc  $e$  carries a flow of 0 or a flow of 1 during the execution of the algorithm and that the identity of this cell can be determined in polynomial time. Under this natural assumption, we get the following result, which implies Theorem 1 for the Successive Shortest Path Algorithm.

**Corollary 1.** *The Successive Shortest Path Algorithm solves PARTITION implicitly.*

By slightly modifying network  $G_{\text{ssp}}^{\bar{a}}$ , we obtain the following result.


**Corollary 2.** *Determining the number of iterations needed by the Successive Shortest Path Algorithm for a given minimum-cost flow instance is NP-hard.*

*Proof.* We replace the arc  $e$  in  $G_{\text{ssp}}^{\bar{a}}$  by two parallel arcs, each with a capacity of  $1/2$  and slightly perturbed costs. This way, every execution of the Successive Shortest Path Algorithm that previously did not use arc  $e$  is unaffected, while executions using  $e$  require additional iterations. Thus, by Lemma 2, the Successive Shortest Path Algorithm on network  $G_{\text{ssp}}^{\bar{a}}$  takes more than  $2^{n+1}$  iterations if and only if the PARTITION instance  $\bar{a}$  has a solution.  $\square$

## 4 Simplex Algorithm and Network Simplex Algorithm

In this section we adapt our construction for the Simplex Algorithm and, in particular, for its interpretation for the minimum-cost flow problem, the Network






**Fig. 3.** Recursive definition of the counting gadget  $S_i^{\vec{v}, r}$  for the Network Simplex Algorithm,  $\vec{v} \in \{\vec{a}, -\vec{a}\}$ , and a parameter  $r \in (2A, 1 - 2A)$ ,  $r \neq 1/2$ . The capacities of the arcs of  $S_i^{\vec{v}, r} \setminus S_{i-1}^{\vec{v}, r}$  are  $x_i + 1 = 3 \cdot 2^{i-1}$ . If we guarantee that there always exists a tree-path from  $t_i$  to  $s_i$  with sufficiently negative cost outside of the gadget, the cost of iteration  $3k$ ,  $k = 0, \dots, 2^i - 1$ , within the gadget is  $k + \vec{v}_i^{[k]}$ . Bold arcs are in the initial basis and carry a flow of at least 1 throughout the execution of the algorithm.

Simplex Algorithm. In this specialized version of the Simplex Algorithm, a basic feasible solution is specified by a spanning tree  $T$  such that the flow value on each arc of the network not contained in  $T$  is either zero or equal to its capacity. We refer to this tree simply as the basis or the spanning tree. The reduced cost of a residual non-tree arc  $e$  equals the cost of sending one unit of flow in the direction of  $e$  around the unique cycle obtained by adding  $e$  to  $T$ . For a pair of nodes, the unique path connecting these nodes in the spanning tree  $T$  is referred to as the *tree-path* between the two nodes. Note that while we setup the initial basis and flow manually in the constructions of the following sections, determining the initial feasible flow algorithmically via the algorithm of Edmonds and Karp, ignoring arc costs, yields the same result.

Our construction ensures that all intermediate solutions of the Network Simplex Algorithm are non-degenerate. Moreover, in every iteration there is a unique non-tree arc of minimum reduced cost which is used as a pivot element.

#### 4.1 A Counting gadget for the Network Simplex Algorithm

We design a counting gadget for the Network Simplex Algorithm(cf. Figure 3), similar to the gadget  $N_i^{\vec{v}}$  of Section 3.1 for the Successive Shortest Path Algorithm. Since the Network Simplex Algorithm augments flow along cycles obtained by adding one arc to the current spanning tree, we assume that the tree always contains an external tree-path from the sink of the gadget to its source with a very low (negative) cost. This assumption will be justified below in Section 4.2 when we embed the counting gadget into a larger network.



**Fig. 4.** Illustration of the iterations performed by the Network Simplex Algorithm on the counting gadget  $S_2^{\vec{a},r}$  for  $r < 1/2$ . The external tree-path from  $t_2$  to  $s_2$  is not shown. Bold arcs are in the basis before each iteration, the red arc enters the basis and the dashed arc exits the basis. Arcs are oriented in the direction in which they are used next. Note that after  $2x_2 = 3 \cdot 2^2 - 2 = 10$  iterations the configuration is the same as in the beginning if we switch the roles of  $s_2$  and  $t_2$ .

The main challenge when adapting the gadget  $N_i^{\vec{v}}$  is that the spanning trees in consecutive iterations of the Network Simplex Algorithm differ in one arc only, since in each iteration a single arc may enter the basis. However, successive shortest paths in  $N_i^{\vec{v}}$  differ by exactly two tree-arcs between consecutive iterations. We obtain a new gadget  $S_i^{\vec{v}}$  from  $N_i^{\vec{v}}$  by modifying arc capacities in such a way that we get two intermediate iterations for every two successive shortest paths in  $N_i^{\vec{v}}$  that serve as a transition between the two paths and their corresponding spanning trees. Recall that in  $N_i^{\vec{v}}$  the capacities of the arcs of  $N_i^{\vec{v}} \setminus N_{i-1}^{\vec{v}}$  are exactly the same as the capacity of the subnetwork  $N_{i-1}^{\vec{v}}$ . In  $S_i^{\vec{v}}$ , we increase the capacity by one unit relative to the capacity of  $S_{i-1}^{\vec{v}}$ . The resulting capacities of the arcs in  $S_i^{\vec{v}} \setminus S_{i-1}^{\vec{v}}$  are  $x_i$  (for the moment), where  $x_i = 2x_{i-1} + 1$  and  $x_1 = 2$ , i.e.,  $x_i = 3 \cdot 2^{i-1} - 1$ .

Similar to before, after  $2x_{i-1}$  iterations the subnetwork  $S_{i-1}^{\vec{v}}$  is saturated. In contrast however, at this point the arcs  $(s_i, s_{i-1})$ ,  $(t_{i-1}, t_i)$  are not saturated yet. Instead, in the next two iterations, the arcs  $(s_i, t_{i-1})$ ,  $(s_{i-1}, t_i)$  enter the basis and one unit of flow gets sent via the paths  $s_i, s_{i-1}, t_i$  and  $s_i, t_{i-1}, t_i$ , which saturates the arcs  $(s_i, s_{i-1})$ ,  $(t_{i-1}, t_i)$  and eliminates them from the basis.

Afterwards, in the next  $2x_{i-1}$  iterations, flow is sent via  $(s_i, t_{i-1}), (s_{i-1}, t_i)$  and through  $S_{i-1}^{\vec{v}}$  as before (cf. Figure 4 for an example execution of the Network Simplex Algorithm on  $S_2^{\vec{v}}$ ).

For the construction to work, we need that, in every non-intermediate iteration, arc  $(s_0, t_0)$  not only enters the basis but, more importantly, is also the unique arc to leave the basis. In other words, we want to ensure that no other arc becomes tight in these iterations. For this purpose, we add an initial flow of 1 along the paths  $s_i, s_{i-1}, \dots, s_0$  and  $t_0, t_1, \dots, t_i$  by adding supply 1 to  $s_i, t_0$  and demand 1 to  $s_0, t_i$  and increasing the capacities of the affected arcs by 1. The arcs of these two paths are the only arcs from the gadget that are contained in the initial spanning tree. We also increase the capacities of the arcs  $(s_i, t_{i-1}), (s_{i-1}, t_i)$  by one to ensure that these arcs are never saturated.

Finally, we also make sure that in every iteration the arc entering the basis is unique. To achieve this, we introduce a parameter  $r \in (2A, 1 - 2A)$ ,  $r \neq 1/2$  and replace the costs of  $2^{i-1} - \frac{1}{2} - \frac{1}{2}v_i$  of the arcs  $(s_i, t_{i-1}), (s_{i-1}, t_i)$  by new costs  $2^{i-1} - r - \frac{1}{2}v_i$  and  $2^{i-1} - (1-r) - \frac{1}{2}v_i$ , respectively.

We later use the final gadget  $S_n^{\vec{v}, r}$  as part of a larger network  $G$  by connecting the nodes  $s_n, t_n$  to nodes in  $G \setminus S_n^{\vec{v}, r}$ . The following lemma establishes the crucial properties of the gadget used in such a way as a part of a larger network  $G$ .

**Lemma 3.** *Let  $S_i^{\vec{v}, r}$ ,  $\vec{v} \in \{\vec{a}, -\vec{a}\}$ , be part of a larger network  $G$  and assume that before every iteration of the Network Simplex Algorithm on  $G$  where flow is routed through  $S_i^{\vec{v}, r}$  there is a tree-path from  $t_i$  to  $s_i$  in the residual network of  $G$  that has cost smaller than  $-2^{i+1}$  and capacity greater than 1. Then, there are exactly  $2x_i = 3 \cdot 2^i - 2$  iterations in which one unit of flow is routed from  $s_i$  to  $t_i$  along arcs of  $S_i^{\vec{v}, r}$ . Moreover:*

1. *In iteration  $j = 3k$ ,  $k = 0, \dots, 2^i - 1$ , arc  $(s_0, t_0)$  enters the basis carrying flow  $k \bmod 2$  and immediately exits the basis again carrying flow  $(k + 1) \bmod 2$ . The cost incurred by arcs of  $S_i^{\vec{v}, r}$  is  $k + \vec{v}_i^{[k]}$ .*
2. *In iterations  $j = 3k + 1, 3k + 2$ ,  $k = 0, \dots, 2^i - 2$ , for some  $0 \leq i' \leq i$ , the cost incurred by arcs of  $S_i^{\vec{v}, r}$  is  $k + r + \vec{v}_{i', i}^{[k]}$  and  $k + (1 - r) + \vec{v}_{i', i}^{[k]}$  in order of increasing cost. One of the arcs  $(s_{i'}, s_{i'-1}), (s_{i'-1}, t_{i'})$  and one of the arcs  $(s_{i'}, t_{i'-1}), (t_{i'-1}, t_{i'})$  each enter and leave the basis in these iterations.*

*Proof.* First observe that throughout the execution of the Network Simplex Algorithm on  $G$ , one unit of flow must always be routed along both of the paths  $s_i, s_{i-1}, \dots, s_0$  and  $t_0, t_1, \dots, t_i$ . This is because there is an initial flow of one along these paths, all of  $s_0, \dots, s_{n-1}$  have in-degree 1, and all of  $t_0, \dots, t_{n-1}$  have out-degree 1, which means that the flow cannot be rerouted.

We prove the lemma by induction on  $i > 0$ , together with the additional property, that after  $2x_i$  iterations the arcs in  $S_{i-1}^{\vec{v}, r}$  carry their initial flow values, while the arcs in  $S_i^{\vec{v}, r} \setminus S_{i-1}^{\vec{v}, r}$  all carry  $x_i$  additional units of flow (which implies that  $(s_i, s_{i-1})$  and  $(t_{i-1}, t_i)$  are saturated). Also, the configuration of the basis is identical to the initial configuration, except that the membership in the basis of arcs in  $S_i^{\vec{v}, r} \setminus S_{i-1}^{\vec{v}, r}$  is inverted. In the following, we assume that  $r \in (2A, 1/2)$ , the

case where  $r \in (1/2, 1 - 2A)$  is analogous. In each iteration  $j$ , let  $P_j$  denote the tree-path outside of  $S_i^{\vec{v}, r}$  from  $t_i$  to  $s_i$  of cost  $c_j < -2^{i+1}$  and capacity greater than 1.

For  $i = 1$ , the Network Simplex Algorithm performs the following four iterations involving  $S_1^{\vec{v}, r}$  (cf. Figure 4 for an illustration embedded in  $S_2^{\vec{v}, r}$ ). In the first iteration,  $(s_0, t_0)$  enters the basis and one unit of flow is routed along the cycle  $s_1, s_0, t_0, t_1, P_0$  of cost  $v_1 + c_0 = \vec{v}_1^{[0]} + c_0$ . This saturates arc  $(s_0, t_0)$  which is the unique arc to become tight (since  $P_0$  has capacity greater than 1) and thus exits the basis again. In the second iteration,  $(s_0, t_1)$  enters the basis and one unit of flow is routed along the cycle  $s_1, s_0, t_1, P_1$  of cost  $r + c_1 = r + \vec{v}_{1,1}^{[0]} + c_1$ , thus saturating (together with the initial flow of 1) arc  $(s_0, s_1)$  of capacity  $x_1 + 1 = 3$ . Since  $P_1$  has capacity greater than 1, this is the only arc to become tight and it thus exits the basis. In the third iteration,  $(s_1, t_0)$  enters the basis and one unit of flow is routed along the cycle  $s_1, t_0, t_1, P_2$  of cost  $(1 - r) + c_2 = (1 - r) + \vec{v}_{1,1}^{[0]} + c_2$ . Similar to before,  $(t_0, t_1)$  is the only arc to become tight and thus exits the basis. In the fourth and final iteration,  $(s_0, t_0)$  enters the basis and one unit of flow is routed along the cycle  $s_1, t_0, s_0, t_1, P_3$  of cost  $1 - v_1 + c_3 = \vec{v}_1^{[1]} + c_3$ , which causes  $(s_0, t_0)$  to become empty and leave the basis. Thus, after four iterations, arc  $(s_0, t_0)$  in  $S_0^{\vec{v}, r}$  carries its initial flow of value 0, while the arcs in  $S_1^{\vec{v}, r} \setminus S_0^{\vec{v}, r}$  all carry  $2 = x_1$  additional units of flow. Also, the arcs  $(s_0, t_1), (s_1, t_0)$  replaced the arcs  $(s_1, s_0), (t_0, t_1)$  in the basis.

To see, for  $i > 0$ , that  $S_i^{\vec{v}, r}$  is saturated after  $2x_i$  units of flow have been routed from  $s_i$  to  $t_i$ , consider the directed  $s_i$ - $t_i$ -cut in  $S_i^{\vec{v}, r}$  induced by  $\{s_i, t_{i-1}\}$  containing the arcs  $(s_i, s_{i-1}), (t_{i-1}, t_i)$ . The capacity of this cut is exactly  $2x_i + 2$  and the initial flow over the cut is 2.

Now assume our claim holds for  $S_{i-1}^{\vec{v}, r}$  and  $S_{i-1}^{\vec{v}, 1-r}$  and consider  $S_i^{\vec{v}, r}$ . Consider the first  $2x_{i-1}$  iterations  $j = 0, \dots, 2x_{i-1} - 1$  and set  $k := \lfloor j/3 \rfloor < 2^{i-1}$ . It can be seen inductively that the shortest path from  $t_{i-1}$  to  $s_{i-1}$  in the bidirected network associated with  $S_{i-1}^{\vec{v}, r}$  has cost at least  $-2^{i-1} + 1 - A > -2^{i-1} + 1 - r$ . Hence, every path from  $s_i$  to  $t_i$  using either or both of the arcs  $(s_i, t_{i-1})$  or  $(s_{i-1}, t_i)$  has cost greater than  $2^{i-1} - (1 - r) - A > 2^{i-1} - 1 + A$ . By induction, we can thus infer that none of these arcs enters the basis in iterations  $j < 2x_{i-1}$ , and instead an arc of  $S_{i-1}^{\vec{v}, r}$  enters (and exits) the basis and one unit of flow gets routed from  $s_i$  to  $t_i$  via the arcs  $(s_i, s_{i-1}), (t_{i-1}, t_i)$ . We may use induction here since, before iteration  $j$ , the path  $t_{i-1}, t_i, P_j, s_i, s_{i-1}$  has cost  $v_i + c_j < v_i - 2^{i+1} < -2^i$  and its capacity is greater than 1, since both  $(s_i, s_{i-1}), (t_{i-1}, t_i)$  have capacity  $x_i + 1 = 2x_{i-1} + 2$ , leaving one unit of spare capacity even after a flow of  $2x_{i-1}$  has been routed along them in addition to the initial unit of flow. The additional cost contributed by arcs  $(s_i, s_{i-1}), (t_{i-1}, t_i)$  is  $v_i$ , which is in accordance with our claim since  $\vec{v}_{\ell, i-1}^{[k]} + v_i = \vec{v}_{\ell, i}^{[k]}$  for all  $\ell \in \{0, \dots, i-1\}$  and  $k \in \{0, \dots, 2^{i-1} - 1\}$ .

Because  $S_{i-1}^{\vec{v}, r}$  is fully saturated after  $2x_{i-1}$  iterations, in the next iteration  $j = 2x_{i-1} = 3 \cdot 2^{i-1} - 2$ ,  $k := \lfloor j/3 \rfloor = 2^{i-1} - 1$ , arc  $(s_{i-1}, t_i)$  is added to the basis and one unit of flow is sent along the path  $s_i, s_{i-1}, t_i$ , thus saturating the capacity  $x_i + 1 = 2x_{i-1} + 2$  of arc  $(s_i, s_{i-1})$  and incurring a cost of  $2^{i-1} - (1 - r) = k + r + \vec{v}_{i, i}^{[k]}$ .

Note that this cost is higher than the cost of each of the previous iterations. The saturated arc has to exit the basis since, by assumption,  $P_j$  has capacity greater than 1. Similarly, in the following iteration  $j = 2x_{i-1} + 1 = 3 \cdot 2^{i-1} - 1$ ,  $k := \lfloor j/3 \rfloor = 2^{i-1} - 1$ , the cost is  $2^{i-1} - r = k + (1 - r) + \bar{v}_{\ell,i}^{[k]}$  and arc  $(t_{i-1}, t_i)$  is replaced by  $(s_i, t_{i-1})$  in the basis.

By induction, at this point  $(s_{i-2}, t_{i-1})$  and  $(s_{i-1}, t_{i-2})$  are in the basis, the arcs of  $S_{i-1}^{\bar{v},r} \setminus S_{i-2}^{\bar{v},r}$  carry a flow of  $x_{i-1}$  in addition to their initial flow, and  $S_{i-2}^{\bar{v},r}$  is back to its initial configuration. To be able to apply induction on the residual network of  $S_{i-1}^{\bar{v},r}$ , we shift the costs of the arcs at  $s_{i-1}$  by  $-(2^{i-2} - r)$  and the costs of the arcs at  $t_{i-1}$  by  $-(2^{i-2} - (1 - r))$  in the residual network of  $S_{i-1}^{\bar{v},r}$ . Since we shift costs uniformly across cuts, this only affects the costs of paths but not the structural behavior of the gadget. Specifically, the costs of all paths from  $t_{i-1}$  to  $s_{i-1}$  in the residual network are increased by exactly  $2^{i-1} - 1$ . If we switch roles of  $s_{i-1}$  and  $t_{i-1}$ , say  $\tilde{s}_{i-1} := t_{i-1}$  and  $\tilde{t}_{i-1} := s_{i-1}$ , we obtain the residual network of  $S_{i-1}^{\bar{v},1-r}$  with its initial flow. This allows us to use induction again for the next  $2x_{i-1}$  iterations.


To apply the induction hypothesis, we need the tree-path from  $\tilde{t}_{i-1} = s_{i-1}$  to  $\tilde{s}_{i-1} = t_{i-1}$  to maintain cost smaller than  $-2^i$  and capacity greater than 1. This is fulfilled since  $P_j$  has cost smaller than  $-2^{i+1}$ , which is sufficient even with the additional cost of  $2^i - 1 - v_i$  incurred by arcs  $(s_i, \tilde{s}_{i-1})$ ,  $(\tilde{t}_{i-1}, t_i)$ . The residual capacity of  $(t_i, \tilde{t}_{i-1})$  and  $(\tilde{s}_{i-1}, s_i)$  is  $x_i > 2x_{i-1}$  and thus sufficient as well. By induction for  $S_{i-1}^{\bar{v},1-r}$ , we may thus conclude that in iterations  $j = 2x_{i-1} + 2, \dots, 2x_i - 1$ ,  $k := \lfloor j/3 \rfloor \geq 2^{i-1}$ , one unit of flow is routed via  $(s_i, t_{i-1})$ ,  $S_{i-1}^{\bar{v},r}$ ,  $(s_{i-1}, t_i)$ . The cost of  $(s_i, \tilde{s}_{i-1})$  and  $(\tilde{t}_{i-1}, t_i)$  together is  $2^i - 1 - v_i$ . The cost of iteration  $j' = j - 2x_{i-1} - 2$ ,  $k' := \lfloor j'/3 \rfloor = k - 2^{i-1}$ , in  $S_{i-1}^{\bar{v},1-r}$  is  $k' + y + \bar{v}_{\ell,i-1}^{[k']}$ , for  $y \in \{0, r, (1 - r)\}$  and  $\ell \in \{0, \dots, i - 1\}$  chosen according to the different cases of the lemma. Accounting for the shift by  $2^{i-1} - 1$  of the cost compared with the residual network of  $S_{i-1}^{\bar{v},r}$ , the incurred total cost in  $S_{i-1}^{\bar{v},r}$  is

$$\begin{aligned} & (2^i - 1 - v_i) + (k' + y + \bar{v}_{\ell,i-1}^{[k']}) - (2^{i-1} - 1) \\ & = 2^{i-1} + k' + y - v_i + \bar{v}_{\ell,i-1}^{[k']} = k + y + \bar{v}_{\ell,i}^{[k]}, \end{aligned}$$

where we used  $-v_i + \bar{v}_{\ell,i-1}^{[k']} = \bar{v}_{\ell,i}^{[k'+2^{i-1}]}$  since  $k' < 2^{i-1}$ . This concludes the proof.  $\square$

## 4.2 The Network Simplex Algorithm implicitly solves PARTITION

We construct a network  $G_{\text{ns}}^{\bar{a}}$  similar to the network  $G_{\text{ssp}}^{\bar{a}}$  of Section 3.2. Without loss of generality, we assume that  $a_1 = 0$ . The network  $G_{\text{ns}}^{\bar{a}}$  consists of the two gadgets  $S_n^+$  and  $S_n^-$  (cf. Figure 5), where  $S_i^\pm := S_i^{\pm \bar{a}, 1/3}$ ,  $i = 0, 1, \dots, n$ . Let  $s_i^\pm, t_i^\pm$  denote the nodes of  $S_n^\pm$ . The two gadgets are connected to a new source node  $s$  and a new sink  $t$  by introducing arcs  $(s, s_n^+)$ ,  $(s, s_n^-)$ ,  $(t_n^+, t)$ ,  $(t_n^-, t)$ , each with capacity  $\infty$  and cost 0. The supply 1 of  $s_n^+$  and  $s_n^-$  is moved to  $s$  and the



**Fig. 5.** Illustration of network  $G_{\text{ns}}^{\vec{a}}$ . The subnetworks  $S_n^+$  and  $S_n^-$  are advanced independently by the Network Simplex Algorithm without using the dashed arc  $e$ , unless the PARTITION instance  $\vec{a}$  has a solution. Bold arcs are in the initial basis and carry a flow of at least 1 throughout the execution of the algorithm.

initial flow on arcs  $(s, s_n^+)$  and  $(s, s_n^-)$  is set to 1. Similarly, the demand 1 of  $t_n^+$  and  $t_n^-$  is moved to  $t$  and the initial flow on arcs  $(t_n^+, t)$  and  $(t_n^-, t)$  is set to 1. Finally, we add an infinite capacity arc  $(s, t)$  of cost  $2^{n+1}$ , increase the supply of  $s$  and the demand of  $t$  by  $4x_n + 1$ , and set the initial flow on  $(s, t)$  to  $4x_n + 1$ .

In addition, we add two new nodes  $c^+$ ,  $c^-$  and replace the arc  $(s_0^+, t_0^+)$  by two arcs  $(s_0^+, c^+)$ ,  $(c^+, t_0^+)$  of capacity 2 and cost 0 (for the moment), and analogously for the arc  $(s_0^-, t_0^-)$  and  $c^-$ . Finally, we move the demand of 1 from  $s_0^+$  to  $c^+$  and the supply of 1 from  $t_0^-$  to  $c^-$ . The arcs  $(s_0^+, c^+)$  and  $(c^-, t_0^-)$  carry an initial flow of 1 and are part of the initial basis. Observe that these modifications do not change the behavior of the gadgets. In addition to the properties of Lemma 3 we have that whenever the arc  $(s_0, t_0)$  previously carried a flow of 1, now the arc  $(c^+, t_0^+)$  or  $(s_0^-, c^-)$  is in the basis, and whenever  $(s_0, t_0)$  previously did not carry flow, now the arc  $(s_0^+, c^+)$  or  $(c^-, t_0^-)$  is in the basis.

We slightly increase the costs of the arcs  $(c^+, t_0^+)$  and  $(s_0^-, c^-)$  from 0 to  $\frac{1}{5}\epsilon$ , again without affecting the behavior of the gadgets. To make every pivot step unique, we need to ensure that equivalent paths in  $S_n^+$  and  $S_n^-$  have different costs. We do this by further slightly perturbing the costs in  $S_n^-$ , such that every path using arcs of  $S_n^-$  has its cost increased by a positive amount less than  $\epsilon/5$ . Finally, we add one more arc  $e = (c^+, c^-)$  with cost 0 and capacity  $\frac{1}{2}$ .

**Lemma 4.** *Arc  $e$  enters the basis in some iteration of the Network Simplex Algorithm on network  $G_{\text{ns}}^{\vec{a}}$  if and only if the PARTITION instance  $\vec{a}$  has a solution.*

*Proof.* First observe that  $\vec{a}_n^{[2k]} = \vec{a}_n^{[2k+1]}$  for  $k \in 0, \dots, 2^{n-1}$  since, by assumption,  $a_1 = 0$ .

Similar to the proof of Lemma 2, in isolation each of the two gadgets can be in one of  $2x_n + 1$  states (Lemma 3), which we label by the number of iterations  $0, 1, \dots, 2x_n$  needed to reach each state. Assuming that both gadgets are in state  $12k$ ,  $k \in \mathbb{Z}_{\geq 0}$ , after some number of iterations, we show that both gadgets will reach state  $12k + 12$  together as well. In addition, we show that, in the iterations in-between, arc  $e$  enters the basis if and only if  $\vec{a}_n^{[4k]} = 0$  (and thus  $\vec{a}_n^{[4k+1]} = 0$ ) or  $\vec{a}_n^{[4k+2]} = 0$  (and thus  $\vec{a}_n^{[4k+3]} = 0$ ). Consider the situation where both gadgets are in state  $12k$ . Note that in this state the arcs in  $S_1^+$  and  $S_1^-$  are back in their original configuration.

Let  $P^\pm$  denote the tree-path from  $t_1^\pm$  to  $s_1^\pm$ , and let  $P^{\pm\mp}$  denote the tree-path from  $t_1^\mp$  to  $s_1^\pm$ . We refer to these paths as the *outer* paths. Observe that, since the gadgets are in the same state, the costs of the outer paths differ by at most  $A < 1/4$ . In the next iterations, flow is sent along a cycle containing one of the outer paths, and we analyze only the part of each cycle without the outer path. Let  $P_0^\pm, P_1^\pm, P_2^\pm, P_3^\pm$  be the four successive shortest paths within the gadget  $S_1^\pm$ . The costs of these paths are  $\frac{1}{5}\varepsilon, 1/3, 2/3, 1 - \frac{1}{5}\varepsilon$ , respectively. Note that, since  $A < 1/6$ , the costs of the paths stay in the same relative order within each gadget throughout the algorithm.

If  $\vec{a}_n^{[4k]} < 0$ , then  $P^+$  is the cheapest of the outer paths by a margin of more than  $\varepsilon/2$ . Thus, in the first iteration,  $(c^+, t_0^+)$  replaces  $(s_0^+, c^+)$  in the basis closing the path  $P_0^+$ . In the next five iterations, the paths  $P_0^-, P_1^+, P_1^-, P_2^+, P_2^-$  are closed in this order. The final two iterations are  $P_3^+, P_3^-$ , similar to the first two iterations, as  $\vec{a}_n^{[4k+1]} = \vec{a}_n^{[4k]} < 0$ . At this point, 8 iterations have passed and both gadgets are in state  $12k + 4$ .

If  $\vec{a}_n^{[4k]} > 0$ , then  $P^-$  is the cheapest of the outer paths by a margin of more than  $\varepsilon/2$ . Thus, the first iteration closes the path  $P_0^-$ . The next five iterations are via  $P_0^+, P_1^-, P_1^+, P_2^-, P_2^+$ , in this order. The final two iterations are  $P_3^-, P_3^+$ , similar to the first two iterations, as  $\vec{a}_n^{[4k+1]} = \vec{a}_n^{[4k]} > 0$ . At this point, 8 iterations have passed and both gadgets are in state  $12k + 4$ .

If  $\vec{a}_n^{[4k]} = 0$ , then all four outer paths have the same cost (ignoring the slight perturbation of the costs in  $S_n^-$ ). The first iteration is via the path  $s_1^+, s_0^+, c^+, c^-, t_0^-, t_1^-$ , i. e., arc  $e$  enters and leaves the basis, for a cost of 0 and an additional flow of  $1/2$ . The next two iterations are via  $P_1^\pm$ , each for a cost of  $\frac{1}{5}\varepsilon$  and an additional flow of  $1/2$ . The fourth iteration is via the path  $s_1^-, s_0^-, c^-, c^+, t_0^+, t_1^+$ , i. e., arc  $e$  enters and leaves the basis again, for a cost of  $\frac{2}{5}\varepsilon$  and an additional flow of  $1/2$ . The next iterations are as before: via  $P_1^+, P_1^-, P_2^-, P_2^+$ , in this order. The final four iterations are similar to the first four iterations, again twice using  $e$ , as  $\vec{a}_n^{[4k+1]} = \vec{a}_n^{[4k]} = 0$ . At this point, 12 iterations have passed and both gadgets are in state  $12k + 4$ .

After the next four iterations (two for each gadget), which do not involve the subnetworks  $S_1^+, S_1^-$  and do thus not use  $e$ , both gadgets are in state  $12k + 6$ .

The iterations going from state  $12k + 6$  to state  $12k + 12$  are analogous to the above if we exchange the roles of  $s_1^\pm$  and  $t_1^\pm$ .  $\square$

The proof of Lemma 4 implies the following observation.

**Corollary 3.** *Determining the number of iterations taken by the Network Simplex Algorithm for a given instance is NP-hard.*

*Proof.* From the proof of Lemma 4 it follows that the Network Simplex Algorithm takes more than  $4x_n$  iterations for network  $G_{\text{ns}}^{\vec{a}}$  if and only if the PARTITION instance  $\vec{a}$  has a solution.  $\square$

Again, we assume that a single cell of the Turing machine corresponding to the Simplex Algorithm can be used to detect whether a variable is in the basis and that the identity of this cell can be determined in polynomial time. Under this natural assumption, we get the following result, which implies Theorem 1 for the Network Simplex Algorithm and thus the Simplex Algorithm.

**Corollary 4.** *The Network Simplex Algorithm implicitly solves PARTITION.*

Finally, we state the complexity result of Lemma 4 in terms of the Simplex Algorithm.

**Corollary 5.** *It is NP-hard to decide for a given linear program whether a given variable ever enters the basis during the execution of the Simplex Algorithm.*

## 5 Further results

We discuss interesting consequences of the results presented above. We first state complexity results for parametric flows and, more generally, parametric linear programming.

**Corollary 6.** *Determining whether a parametric minimum-cost flow uses a given arc (i. e., assigns positive flow value for any parameter value) is NP-hard. In particular, determining whether the solution to a parametric linear program uses a given variable is NP-hard. Also, determining the number of different basic solutions over all parameter values is NP-hard.*

*Proof.* This follows from the fact that the Successive Shortest Path Algorithm solves a parametric minimum-cost flow problem, together with Lemma 2 and Corollary 2.  $\square$

We also obtain a complexity result on 2-dimensional projections of polyhedra.

**Corollary 7.** *Given a  $d$ -dimensional polytope  $P$ , determining the number of vertices of  $P$ 's projection onto a given 2-dimensional subspace is NP-hard.*



*Proof.* Let  $P$  be the polytope of all feasible  $s$ - $t$ -flows in network  $G_{\text{ssp}}^{\bar{a}}$  of Section 3.2. Consider the 2-dimensional subspace  $S$  defined by flow value and cost of a flow. Let  $P'$  be the projection of  $P$  onto  $S$ . The lower envelope of  $P'$  is the parametric minimum-cost flow curve for  $G_{\text{ssp}}^{\bar{a}}$ , while the upper envelope is the parametric maximum-cost flow curve for  $G_{\text{ssp}}^{\bar{a}}$ .

The  $s$ - $t$ -paths of maximum cost in  $G_{\text{ssp}}^{\bar{a}}$  are the four paths via  $s_n, s_{n-1}, t_n$  or via  $s_n, t_{n-1}, t_n$  in both of the gadgets. Each of these paths has cost  $2^{n-1} - \frac{1}{2}$  and the total capacity of all paths together is  $2^{n+1}$  which is equal to the maximum flow value from  $s$  to  $t$ . Therefore, the upper envelope of  $P'$  consists of a single edge.

The number of edges on the lower envelope of  $P'$  is equal to the number of different costs among all successive shortest paths in  $G_{\text{ssp}}^{\bar{a}}$ . If we slightly perturb the costs of the two arcs in  $G_{\text{ssp}}^{\bar{a}}$  with cost  $\frac{1}{5}\varepsilon$ , we can ensure that each successive shortest path has a unique cost. The claim then follows by Corollary 2.  $\square$

We finally mention a result for a long-standing open problem in the area of network flows over time (see, e. g., [21] for an introduction). The goal in earliest arrival flows is to find an  $s$ - $t$ -flow over time that simultaneously maximizes the amount of flow that has reached the sink  $t$  at any point in time. The Successive Shortest Path Algorithm can be used to obtain such an earliest arrival flow. All known encodings of earliest arrival flows suffer from exponential worst-case size, and it has been an open problem whether there is a polynomial encoding which can be found in polynomial time. The following corollary implies that, in a certain sense, earliest arrival flows are NP-hard to obtain. In this context, it is interesting to mention that an  $s$ - $t$ -flow over time is an earliest arrival flow if and only if it has minimum average arrival time [12].

**Corollary 8.** *Determining the minimum average arrival time of a maximum  $s$ - $t$ -flow over time is NP-hard.*

*Proof.* Consider network  $G_{\text{ssp}}^{\bar{a}}$  introduced in Section 3.2 and scale all arc costs by a sufficiently large integer to make them integral. Moreover, let  $\xi := 1/2^{n+2}$  and change the cost of arc  $e$  in  $G_{\text{ssp}}^{\bar{a}}$  from 0 to  $\xi$ . Notice that this modification does not change the sequence of paths chosen by the Successive Shortest Path Algorithm. Denote the resulting network by  $G_{\xi}^{\bar{a}}$ . Jarvis and Ratliff [12] proved that an earliest arrival flow has minimum average arrival time [12] (and vice versa). We therefore consider in the following the earliest arrival flow on  $G_{\xi}^{\bar{a}}$  that can be obtained from the paths found by the Successive Shortest Path Algorithm; see, e. g., [21] for details.

As argued in Section 3, the Successive Shortest Path Algorithm takes  $2^{n+1}$  iterations on network  $G_{\xi}^{\bar{a}}$ . In each iteration  $i = 0, \dots, 2^{n+1} - 1$  it augments one unit of flow along some path  $P_i$  of cost  $c(P_i)$  in the residual network. Notice that  $c(P_i)$  is integral unless it contains arc  $e$ . In the latter case,  $c(P_i) = z \pm \xi$  for some  $z \in \mathbb{Z}_{>0}$ . Let  $k := |\{i \in \{0, \dots, 2^{n+1} - 1\} : P_i \text{ contains } e\}|$  such that  $0 \leq k \leq 2^{n+1}$ . In particular,  $k = 0$  if and only if the PARTITION instance  $\bar{a}$  has a solution.

An earliest arrival flow with integral time horizon  $T > c(P_{2^{n+1}}) \geq c(P_i)$  sends flow at rate 1 into path  $P_i$  from time 0 up to time  $T - c(P_i)$ , for  $i = 0, \dots, 2^{n+1} - 1$ . In particular, the total flow sent along  $P_i$  over time is  $T - c(P_i)$ . This flow arrives at the sink at rate 1 between time  $c(P_i)$  and time  $T$ ; its average arrival time is  $\frac{1}{2}(T + c(P_i))$ . Thus, the overall average arrival time of flow at the sink is

$$\frac{1}{F} \sum_{i=0}^{2^{n+1}-1} (T - c(P_i)) \frac{1}{2} (T + c(P_i)) = \frac{1}{2F} \sum_{i=0}^{2^{n+1}-1} (T^2 - c(P_i)^2) , \quad (1)$$

where  $F$  is the total amount of flow sent into the sink, i.e., the value of a maximum  $s$ - $t$ -flow over time. Since  $T$  is integral, it follows from (1) that  $2F$  times the average arrival time is of the form  $\alpha \pm \beta\xi - k\xi^2$  with  $\alpha, \beta \in \mathbb{Z}_{\geq 0}$ . Since  $0 \leq k \leq 2^{n+1}$  and  $\xi = 1/2^{n+2}$  divides  $\alpha$ , this value is a multiple of  $\xi$  if and only if  $k = 0$ , that is, if and only if the PARTITION instance  $\vec{a}$  has a solution.

Since the maximum value  $F$  of an  $s$ - $t$ -flow over time can be computed in polynomial time [8], we can decide PARTITION by observing the minimum average arrival time of a maximum  $s$ - $t$ -flow over time in  $G_{\xi}^{\vec{a}}$ .  $\square$

## 6 Conclusion

We have introduced the concept of NP-*mightiness* in order to provide the Successive Shortest Path Algorithm and the (Network) Simplex Method with Dantzig's pivot rule with an excuse and a more satisfactory explanation regarding their exponential worst-case running time: These algorithms implicitly solve NP-hard problems!

We believe that these results can be carried over to the Simplex Method with other pivot rules and hope that our approach will turn out to be useful in developing a better understanding of other algorithms that suffer from poor worst-case behavior. Preliminary work, building on the construction of Friedmann [9], suggests that the Simplex Method with Zadeh's pivot rule [25] is NP-mighty as well.

## References

1. R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows. Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ, 1993.
2. N. Amenta and G. M. Ziegler. Deformed products and maximal shadows of polytopes. In *Advances in Discrete and Computational Geometry*, pages 57–90. Amer. Math. Soc, 1996.
3. T. Brunsch, K. Cornelissen, B. Manthey, and H. Röglin. Smoothed analysis of the successive shortest path algorithm. In *Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms*, pages 1180–1189, 2013.
4. R. G. Busacker and P. J. Gowen. A procedure for determining a family of minimum-cost network flow patterns. Technical Paper ORO-TP-15, Operations Research Office, The Johns Hopkins University, Bethesda, Maryland, 1960.

5. G. B. Dantzig. Application of the simplex method to a transportation problem. In Tj. C. Koopmans, editor, *Activity Analysis of Production and Allocation – Proceedings of a Conference*, pages 359–373. Wiley, 1951.
6. G. B. Dantzig. Maximization of a linear function of variables subject to linear inequalities. In Tj. C. Koopmans, editor, *Activity Analysis of Production and Allocation – Proceedings of a Conference*, pages 339–347. Wiley, 1951.
7. G. B. Dantzig. *Linear programming and extensions*. Princeton University Press, 1962.
8. L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
9. O. Friedmann. A subexponential lower bound for Zadeh’s pivoting rule for solving linear programs and games. In O. Günlük and G. J. Woeginger, editors, *Proceedings of the 15th Conference on Integer Programming and Combinatorial Optimization*, Lecture Notes in Computer Science, pages 192–206. Springer, 2011.
10. M. R. Garey and D. S. Johnson. *Computers and Intractability, A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
11. M. Iri. A new method of solving transportation-network problems. *Journal of the Operations Research Society of Japan*, 3:27–87, 1960.
12. J. J. Jarvis and H. D. Ratliff. Some equivalent objectives for dynamic network flow problems. *Management Science*, 28:106–108, 1982.
13. N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4:373–395, 1984.
14. L. G. Khachiyan. A polynomial algorithm in linear programming. *Soviet Mathematics Doklady*, 20:191–194, 1979.
15. L. G. Khachiyan. Polynomial algorithms in linear programming. *U.S.S.R. Computational Mathematics and Mathematical Physics*, 20:53–72, 1980.
16. V. Klee and G. J. Minty. How good is the simplex algorithm? In O. Shisha, editor, *Inequalities III*, pages 159–175. Academic Press, New York, 1972.
17. B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*. Springer, 5th edition, 2012.
18. J. Matoušek and B. Gärtner. *Understanding and using linear programming*. Springer, 2007.
19. J. B. Orlin. A polynomial time primal network simplex algorithm for minimum cost flows. *Mathematical Programming*, 78:109–129, 1997.
20. A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, Chichester, 1986.
21. M. Skutella. An introduction to network flows over time. In W. Cook, L. Lovász, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*, pages 451–482. Springer, 2009.
22. D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51:385–463, 2004.
23. É. Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5:247–255, 1985.
24. N. Zadeh. A bad network problem for the simplex method and other minimum cost flow algorithms. *Mathematical Programming*, 5:255–266, 1973.
25. N. Zadeh. What is the worst case behavior of the simplex algorithm? Technical Report 27, Department of Operations Research, Stanford, 1980.