

# Single-Source Unsplittable Flows in Planar Graphs\*

Vera Traub<sup>‡</sup>

Laura Vargas Koch<sup>§</sup>

Rico Zenklusen<sup>¶</sup>

## Abstract

The single-source unsplittable flow (SSUF) problem asks to send flow from a common source to different terminals with unrelated demands, each terminal being served through a single path. One of the most heavily studied SSUF objectives is to minimize the violation of some given arc capacities. A seminal result of Dinitz, Garg, and Goemans showed that, whenever a fractional flow exists respecting the capacities, then there is an unsplittable one violating the capacities by at most the maximum demand. Goemans conjectured a very natural cost version of the same result, where the unsplittable flow is required to be no more expensive than the fractional one. This intriguing conjecture remains open. More so, there are arguably no non-trivial graph classes for which it is known to hold.

We show that a slight weakening of it (with at most twice as large violations) holds for planar graphs. Our result is based on a connection to a highly structured discrepancy problem, whose repeated resolution allows us to successively reduce the number of paths used for each terminal, until we obtain an unsplittable flow. Moreover, our techniques also extend to simultaneous upper and lower bounds on the flow values. This also affirmatively answers a conjecture of Morell and Skutella for planar SSUF.

---

\*This project received funding from Swiss National Science Foundation grant 200021\_184622 and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 817750).

<sup>‡</sup>Research Institute for Discrete Mathematics and Hausdorff Center for Mathematics, University of Bonn. Email: [traub@dm.uni-bonn.de](mailto:traub@dm.uni-bonn.de).

<sup>§</sup>Department of Mathematics, ETH Zurich, Zurich, Switzerland. Email: [laura.vargas@ifor.math.ethz.ch](mailto:laura.vargas@ifor.math.ethz.ch).

<sup>¶</sup>Department of Mathematics, ETH Zurich, Zurich, Switzerland. Email: [ricoz@ethz.ch](mailto:ricoz@ethz.ch).

# 1 Introduction

Flow problems are among the most classical combinatorial optimization problems. To their prominent role in the field significantly contributed both their wide, and sometimes surprising, range of applications and also the existence of very fast algorithms (in practice as well as in theory) for basic flow optimization problems. (We refer the interested reader to the textbooks of Ahuja, Magnanti, and Orlin [AMO93], and Williamson [Wil19] for more information.) In their most canonical version, flow from a source to a sink can be split over many paths. Being able to split flows naturally allows for casting flow problems—even with multiple sources, sinks, and commodities—as linear programs. Flow problems become much harder, and are also much less understood, as soon as the flow is not allowed to be split, i.e., is unsplittable.

In this work, we consider a heavily studied variant thereof, namely the *single-source unsplittable flow problem* (SSUF), originally introduced by Kleinberg [Kle96a; Kle96b]. Here, we are given a directed graph  $G = (V, A)$ , with arc capacities  $u \in \mathbb{Q}_{\geq 0}^A$ , a single source  $s \in V$ , and terminals/sinks  $T \subseteq V$ , where each  $t \in T$  has a demand  $d_t \in \mathbb{Q}_{\geq 0}$ . (The notions of sink and terminal are used interchangeably in the literature.) For simplicity, and without loss of generality, we assume that the source  $s$  and sinks  $T$  are all distinct. Ideally, one would like to route the demanded flow from the source to each terminal on a single path such that all capacities are respected. Formally, this corresponds to determining one  $s$ - $t$  path  $P^t \subseteq A$  for each  $t \in T$  such that  $\sum_{t \in T: a \in P^t} d_t \leq u(a)$  for each  $a \in A$ . Deciding whether an *unsplittable flow*, i.e., such a family  $\mathcal{P} = \{P^t\}_{t \in T}$  of paths, exists, can easily be seen to be NP-hard. Actually, even in a 2-vertex graph, one can reduce Bin Packing or Subset Sum to it. Thus, significant attention has been devoted to the setting where one is allowed to exceed the capacities and the goal is to minimize the largest capacity violation. This is also the setting we consider here.

The research in this area has been largely influenced by a seminal result of Dinitz, Garg, and Goemans [DGG99] and a subsequent conjecture of Goemans, which we discuss next. These are statements about the existence of unsplittable flows with limited capacity violation assuming that there exists a *splittable flow*  $x \in \mathbb{Q}^A$  that respects the capacities, i.e., such a vector  $x$  satisfies  $x(a) \leq u(a)$  for all arcs  $a \in A$  and

$$x(\delta^+(v)) - x(\delta^-(v)) = \begin{cases} \sum_{t \in T} d_t & \text{if } v = s, \\ -d_v & \text{if } v \in T, \\ 0 & \text{if } v \in V \setminus (\{s\} \cup T), \end{cases}$$

where  $\delta^-(v)$  and  $\delta^+(v)$  denote the sets of all arcs entering and leaving  $v$ , respectively. More precisely, these statements assume the existence of a splittable flow  $x \in \mathbb{Q}_{\geq 0}^A$ , and the capacity violation of an unsplittable flow is measured by the difference between the flow value of the unsplittable flow and the value of the splittable one. In short, this can be interpreted as a worst-case assumption where the capacity  $u(a)$  of an arc  $a$  is equal to  $x(a)$ . We thus define an SSUF instance as follows.

**Definition 1.1** (Single-source unsplittable flow (SSUF) instance). *A single-source unsplittable flow (SSUF) instance is a tuple  $(G = (V, A), s, T, d, x)$ , where  $G$  is a directed graph with source  $s \in V$ , terminals  $T \subseteq V$  with corresponding demands  $d \in \mathbb{Q}_{\geq 0}^T$ , and a splittable flow  $x \in \mathbb{Q}_{\geq 0}^A$ . The source  $s$  and terminals  $T$  are assumed to be distinct.*

The result of Dinitz, Garg, and Goemans [DGG99], stated below, was the first to establish an additive capacity violation of  $O(d_{\max})$ , where  $d_{\max} := \max\{d_t : t \in T\}$  is the largest demand, whenever a splittable flow exists. For ease of notation, given an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  and

an arc  $a \in A$ , we denote by

$$\text{flow}_{\mathcal{P}}(a) := \sum_{t \in T: a \in P^t} d_t$$

the total flow value of the unsplittable flow  $\mathcal{P}$  that traverses  $a$ .

**Theorem 1.2** ([DGG99]). *Given an SSUF instance  $(G, s, T, u, d, x)$ , one can compute in polynomial time an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with  $\text{flow}_{\mathcal{P}}(a) \leq x(a) + d_{\max}$  for all  $a \in A$ .*

Shortly thereafter, Goemans conjectured that the following stronger, cost-enhanced version of the same result holds.

**Conjecture 1.3** (Goemans). *Given an SSUF instance  $(G, s, T, u, d, x)$  and a cost vector  $c \in \mathbb{Q}_{\geq 0}^A$ , one can compute in polynomial time an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with*

- (i)  $\text{flow}_{\mathcal{P}}(a) \leq x(a) + d_{\max}$  for all  $a \in A$ , and
- (ii) *cost at most  $c^T x$ , i.e.,  $\sum_{a \in A} c(a) \text{flow}_{\mathcal{P}}(a) \leq \sum_{a \in A} c(a) x(a)$ .*

Note that one can assume  $G$  to be acyclic in [Conjecture 1.3](#). Indeed, if  $G$  contains directed cycles, then flow along such cycles can be reduced and zero-flow arcs can be removed. This leads to an acyclic instance that is no easier than the original one.

The following even stronger version of Goemans' conjecture, which includes both upper and lower bounds on the flow values, has been stated by Morell and Skutella [\[MS22\]](#).

**Conjecture 1.4** ([\[MS22\]](#)). *Given an SSUF instance  $(G, s, T, u, d, x)$  on an acyclic graph  $G$  and a cost vector  $c \in \mathbb{Q}_{\geq 0}^A$ , one can compute in polynomial time an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with*

- (i)  $x(a) - d_{\max} \leq \text{flow}_{\mathcal{P}}(a) \leq x(a) + d_{\max}$  for all  $a \in A$ , and
- (ii)  $\sum_{a \in A} c(a) \text{flow}_{\mathcal{P}}(a) \leq \sum_{a \in A} c(a) x(a)$ .

The fact that  $G$  is explicitly assumed to be acyclic in [Conjecture 1.4](#) is due to the lower bounds. Without acyclicity, the nature of the problem changes. In particular, it would suddenly become important whether an unsplittable flow consists of paths or walks, because one may try to fulfill lower bounds by going along cycles. Also, one can show that statements as [Conjecture 1.4](#) (even with a violation of  $O(d_{\max})$  instead of just  $d_{\max}$ ) cannot be obtained in general (non-acyclic) graphs (see [Appendix A](#)).

Morell and Skutella [\[MS22\]](#) also explicitly conjectured the following weaker version of [Conjecture 1.4](#) without costs, which also remains open and is a natural intermediate step toward [Conjecture 1.4](#).

**Conjecture 1.5** ([\[MS22\]](#)). *Given an SSUF instance  $(G, s, T, u, d, x)$  on an acyclic graph  $G$  there exists an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with*

$$x(a) - d_{\max} \leq \text{flow}_{\mathcal{P}}(a) \leq x(a) + d_{\max} \text{ for all } a \in A.$$

[Conjecture 1.5](#) can be seen as the counterpart of [Theorem 1.2](#) with additional lower bounds. One motivation for [Conjecture 1.5](#) is that a positive resolution of it would have implications to both some open discrepancy and scheduling problems, on which we briefly expand in [Section 1.2](#).

All three conjectures, [Conjectures 1.3](#) to [1.5](#), remain open, and even a resolution of a weaker version of Goemans' conjecture with additive capacity violations in the order  $O(d_{\max})$  (instead of the conjectured  $d_{\max}$ , which is the best one can hope for in terms of constants) would likely be considered a breakthrough.

Procedures leading to an unsplittable flow with a multiplicative violation, i.e.,  $\text{flow}_{\mathcal{P}}(a) = O(x(a))$  for  $a \in A$ , have been obtained previously [\[Kle96a; Kle96b; KS02; Sku02; MS22\]](#). Several

of these approximation algorithms have also been implemented and empirically tested [DK05]. Moreover, Skutella [Sku02] showed that [Conjecture 1.3](#) holds when the demands are all multiples of each other. Furthermore, techniques presented in Lenstra, Shmoys, and Tardos [LST90] solve the problem for the special case when  $G$  has a source plus two layers of vertices, with arcs only going from the source to the first layer and from the first to the second layer.

Very little is known about [Conjectures 1.4](#) and [1.5](#). When only dealing with lower bounds, analogous results (for acyclic graphs) can be obtained as with upper bounds only. More precisely, Morell and Skutella [MS22] showed that for any SSUF instance  $(G, s, T, u, d, x)$ , there exists an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with  $\text{flow}_{\mathcal{P}}(a) \geq x(a) - d_{\max}$  for  $a \in A$ . Moreover, one can also adjust the algorithm of Dinitz, Garg, and Goemans [DGG99] to get the same lower bound guarantees. However, the combination of both lower and upper bounds seems to be much more challenging.

Prior to this work, there was arguably no non-trivial graph class for which any of [Conjectures 1.3](#) to [1.5](#) was known to hold, even if additive violations in the order of  $O(d_{\max})$  are allowed instead of only  $d_{\max}$ . The goal of this work is to address this gap, by showing that a slightly weaker version of [Conjecture 1.4](#) (allowing a violation of  $2d_{\max}$  instead of  $d_{\max}$ ) and [Conjecture 1.5](#) hold for planar graphs. This is a further positive sign regarding [Conjectures 1.3](#) to [1.5](#). We obtain this result by connecting SSUF in planar graphs to a very well-structured discrepancy problem, whose resolution allows us to successively transform a splittable flow into another one that uses fewer paths per terminal until, eventually, an unsplittable flow is obtained.

## 1.1 Our results

As mentioned, we focus on planar graphs in this work. A planar instance of SSUF is formally defined as follows.

**Definition 1.6** (Planar single-source unsplittable flow (PSSUF) instance). *An SSUF instance  $(G, s, T, d, x)$  is a planar single-source unsplittable flow (PSSUF) instance if  $G$  is acyclic and planar.*

Our first main result shows that, for planar graphs, unsplittable flows can be constructed with the lower and upper bound guarantees as claimed in [Conjecture 1.5](#).

**Theorem 1.7.** *Given a PSSUF instance  $(G, s, T, d, x)$ , there is an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with*

$$x(a) - d_{\max} \leq \text{flow}_{\mathcal{P}}(a) \leq x(a) + d_{\max} \quad \forall a \in A.$$

*Moreover, it can be computed in time  $O(|V|^2)$ .*

We obtain an analogous result involving costs when allowing for a slightly larger lower and upper bound violation of  $2d_{\max}$  instead of  $d_{\max}$ .

**Theorem 1.8.** *Given a PSSUF instance  $(G, s, T, d, x)$  and arc costs  $c \in \mathbb{Q}_{\geq 0}^A$ , there is an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  satisfying*

- (i)  $x(a) - 2d_{\max} \leq \text{flow}_{\mathcal{P}}(a) \leq x(a) + 2d_{\max} \quad \forall a \in A$ , and
- (ii)  $\sum_{a \in A} c(a) \text{flow}_{\mathcal{P}}(a) \leq \sum_{a \in A} c(a)x(a)$ .

*Moreover, it can be computed in time  $O(\langle \text{input} \rangle \cdot |V|)$ , where  $\langle \text{input} \rangle$  denotes the input size of the instance.*

[Theorem 1.8](#) settles [Conjecture 1.4](#) for planar graphs, when allowing for a violation of  $2d_{\max}$  instead of  $d_{\max}$ . Moreover, the above results are the first SSUF results to simultaneously respect upper and lower bounds with only additive errors of  $O(d_{\max})$  in a non-trivial graph class. We provide an overview of our techniques, which are based on a reduction to a well-structured discrepancy problem, in [Section 2](#).

## 1.2 Connections to discrepancy problems and scheduling

SSUF is closely connected to scheduling problems. Already Kleinberg [Kle96a] observed that a natural special case of makespan minimization on unrelated machines can be cast as SSUF with upper bounds only. When allowing for upper and lower bounds, an interesting more general scheduling problem can be captured as observed by Lars Rohwedder. This is a non-preemptive maximum flow time minimization problem with release times, where each job can only be processed on a subset of machines. It remains open whether this setting allows for  $O(1)$ -approximations, and a constructive resolution of Conjecture 1.5, even with a violation of  $O(d_{\max})$  instead of  $d_{\max}$ , would resolve this question. (See [MS22] for more details including a discussion of Rohwedder’s reduction.) The currently best known approximation factor is  $O(\log n)$  by Bansal and Kulkarni [BK15], while the best known lower bound for the variant with unrelated machines is  $3/2$  (unless  $P = NP$ ) by Lenstra, Shmoys, and Tardos [LST90]. Moreover, Bansal, Rohwedder, and Svensson [BRS22] showed very recently that the integrality gap of the natural LP relaxation for this scheduling problem is  $O(\sqrt{\log n})$ . (A non-constructive resolution of Conjecture 1.5 would imply an  $O(1)$  integrality gap for this LP relaxation.)

Furthermore, in [BRS22] an intriguing equivalence between the existence of an LP-based  $O(1)$ -approximation for this scheduling problem and the resolution of an interesting open discrepancy problem was revealed. This discrepancy problem is a special case of the prefix version of Beck-Fiala, which is a well-studied discrepancy problem. A positive answer to Conjecture 1.5 would imply a positive answer to the open discrepancy problem, in an interesting special case, which would be of independent interest.

Interestingly, in our algorithm for PSSUF, we also rely on a prefix version of a discrepancy problem. However, our notion of prefixes is different from the one in the prefix Beck-Fiala problem.

## 1.3 Further related work

Martens, Salazar, and Skutella [MSS07] took an approach toward Goemans’ conjecture where they allow for slight demand modifications. More precisely, they provide a polynomial-time procedure that, given any splittable single-source flow, writes it as a convex combination of unsplittable flows for slightly rounded demands (by a factor of at most 2), such that the average demands used in the convex combination correspond to the original ones. If each term in the convex combination sent the original demands, then this would imply Goemans’ conjecture. (Actually the existence of such a convex combination is equivalent to the existence of an unsplittable flow fulfilling the conditions of Goemans’ conjecture.)

Also, we would like to mention that there has been extensive work on unsplittable flows in a variety of settings, including multiple sources and sinks and with different objectives. We refer the interested reader to the survey by Kolliopoulos [Kol07], and the discussion and references in a very recent contribution of Grandoni, Mömke, and Wiese [GMW22] to unsplittable flows on a path. Moreover, intermediate notions between splittable and unsplittable flows have been considered. In particular,  $k$ -splittable flows, which have been introduced by Baier, Köhler, and Skutella [BKS05], allow for splitting the flow among up to  $k$  paths. (See [Kol05; KSS08; SS09] for further results on  $k$ -splittable flows.)

## 1.4 Organization of the paper

In Section 2, we provide an overview of our approach. A key ingredient is the construction of a nice path decomposition of the fractional flow, which we discuss in Section 3. Section 4 discusses and proves properties of nice path decompositions, which we crucially exploit later. In Section 5, we





Lemma 2.14).<sup>1</sup> These allow for reducing the selection of the paths  $P^t \in \mathcal{P}^t$  to a highly structured discrepancy problem.

In the remainder of this section we provide a more detailed description of our approach. We first describe the choice of our flow decomposition and the resulting structure of the sets  $\mathcal{P}^t$  in Section 2.1 and then describe the discrepancy problem and our solution for it in Section 2.2. Finally, we show how the structure of the sets  $\{i \in [\ell] : a \in P_i\}$  allows for reducing the selection of the paths  $P^t \in \mathcal{P}^t$  to this discrepancy problem.

## 2.1 Choosing the path decomposition


We now provide an overview of how we choose the path decomposition of the flow  $x$ . Further details and proofs will be provided in Section 3.

We start by computing a planar embedding of the graph  $G$ , i.e., an embedding of the vertices and arcs in the plane such that vertices are mapped to distinct points and no two arcs intersect in a point distinct from their endpoints. This can be done in linear time (see [HT74; CNAO85]). Our goal will be to choose the paths  $P_1, \dots, P_\ell$  such that they do not cross in a geometric sense. To this end, it is helpful to first represent the paths  $P_1, \dots, P_\ell$  in an auxiliary graph  $H = (V, F)$  obtained from  $G = (V, A)$ , by replacing each arc  $a \in A$  by a well-chosen number of parallel copies. We call such graphs  $H$  *arc-split graphs* of  $G$ , and, for every  $f \in F$ , we denote by  $\phi(f) \in A$  the arc in  $G$  corresponding to  $f$ , i.e.,  $f$  is a parallel copy of  $\phi(f)$ . The idea is to choose  $H$  such that  $P_1, \dots, P_\ell$  can be chosen to be arc-disjoint in  $H$ .

More precisely, given a PSSUF instance  $(G, s, T, d, x)$ , we start by constructing the following:

- (i) an arc-split graph  $H = (V, F)$  of  $G$  (with the associated mapping  $\phi$ );
- (ii) a partition of  $F$  into paths  $P_1, \dots, P_\ell$ , each of which is a path in  $H$  from  $s$  to some terminal  $t \in T$ , and such that these paths are (geometrically) non-crossing, and numbered counterclockwise around  $s$ ;
- (iii) coefficients  $\lambda_1, \dots, \lambda_\ell \in \mathbb{Q}_{\geq 0}$  such that  $\sum_{i=1}^\ell \lambda_i \chi^{\phi(P_i)} = x$ .

See Figure 2 for an example.



**Figure 2:** The left figure shows a PSSUF instance  $(G, s, T, d, x)$  with  $T = \{t_1, t_2, t_3, t_4\}$ , and the  $x$ -values are shown next to the arcs. The right figure shows an arc-split graph  $H$  of  $G$  together with a source-numbered nice  $s$ -path decomposition  $(P_i, \lambda_i)_{i \in [11]}$ .

<sup>1</sup>The paths  $P_i$  in  $G$  will be called  $\phi(P_i)$  in Lemma 2.14. This is due to technical reasons explained in Section 3.

We call paths  $P_1, \dots, P_\ell$  together with their coefficients  $\lambda_1, \dots, \lambda_\ell$  fulfilling items (ii) and (iii) a *source-numbered nice s-path decomposition* of  $H$ . (For a more formal definition, see Section 3.) It can be thought of, as a very structured flow decomposition of  $x$ .

First, we show that such a decomposition can be obtained in polynomial time.

**Theorem 2.1.** *Let  $(G = (V, A), s, T, d, x)$  be a PSSUF instance. We can compute in  $O(|V|^2)$  time an arc-split graph  $H$  of  $G$  together with a source-numbered nice s-path decomposition  $(P_i, \lambda_i)_{i \in [\ell]}$  of  $H$  with  $\ell = O(|V|)$ .*

Recall, that in the second step of our algorithm, we will choose our PSSUF solution from the paths  $P_1, \dots, P_\ell$  (or more formally, the corresponding paths  $\phi(P_1), \dots, \phi(P_\ell)$  in  $G$ ). More precisely, for each terminal  $t \in T$ , we will pick one path among

$$\mathcal{P}^t := \{P_i : i \in [\ell] \text{ and } P_i \text{ ends at terminal } t\}.$$

One important consequence of the paths  $(P_i)_{i \in [\ell]}$  being non-crossing, is that the sets  $S^t := \{i \in [\ell] : P_i \in \mathcal{P}^t\}$  for  $t \in T$  are *non-interleaving*, as we will prove in Section 4.2.

**Definition 2.2** (Non-interleaving). *Two disjoint sets  $S_1, S_2 \subseteq \mathbb{Z}_{>0}$  are interleaving if there exist  $a_1, b_1 \in S_1, a_2, b_2 \in S_2$  with  $a_1 < a_2 < b_1 < b_2$  or  $a_2 < a_1 < b_2 < b_1$ . A partition  $\mathcal{S}$  of  $[\ell]$  is non-interleaving if for any  $S_1, S_2 \in \mathcal{S}$  with  $S_1 \neq S_2$  the sets  $S_1$  and  $S_2$  are non-interleaving.*

**Lemma 2.3.** *For  $t \in T$  let  $S^t := \{i \in [\ell] : P_i \in \mathcal{P}^t\}$ . Then the partition  $\{S^t : t \in T\}$  of  $[\ell]$  is non-interleaving.*

This is a key structure we will exploit in our discrepancy-based approach to select one path in each  $\mathcal{P}^t$  for each  $t \in T$ , which we discuss next.

## 2.2 Interval-discrepancy on non-interleaving partitions

We now discuss the selection of paths  $P^t \in \mathcal{P}^t$  for all  $t \in T$ , which we formulate as a discrepancy problem.

**Definition 2.4** (Selections). *Let  $\mathcal{S}$  be a partition of  $[\ell]$ . Then*

- a fractional selection for  $\mathcal{S}$  is a vector  $y \in [0, 1]^\ell : q(S) = 1 \ \forall S \in \mathcal{S}$ , and
- an integral selection for  $\mathcal{S}$  is a vector  $z \in \{0, 1\}^\ell : q(S) = 1 \ \forall S \in \mathcal{S}$ .

We will apply this definition to the partition  $\mathcal{S}$  consisting of the sets  $S^t := \{i \in [\ell] : P_i \in \mathcal{P}^t\}$  with  $t \in T$ . The s-path decomposition  $(P_i, \lambda_i)_{i \in [\ell]}$  naturally gives rise to a fractional selection  $y$  for  $\mathcal{S}$  by setting  $y_i := \frac{\lambda_i}{d_t}$  where  $t \in T$  is the terminal where the path  $P_i$  ends. Then the vector  $(\lambda_i)_{i \in [\ell]}$  will be called the *load vector*.

**Definition 2.5** (Load vector). *Given a vector  $y \in [0, 1]^\ell$  and a demand vector  $d \in \mathbb{Q}_{\geq 0}^S$ , the load vector  $y^d \in \mathbb{Q}_{\geq 0}^\ell$  is defined by*

$$y_i^d := d_S \cdot y_i,$$

where  $S \in \mathcal{S}$  is the set containing  $i \in [\ell]$ .

An instance of our discrepancy problem consists of the integer  $\ell$ , the partition  $\mathcal{S}$ , the demand vector  $d$ , and a fractional selection  $y$ .



**Definition 2.6** (Weighted partition-constrained selection instance, WPCS instance). We call a tuple  $(\ell, \mathcal{S}, d, y)$  a weighted partition-constrained selection instance (for short WPCS instance) if

- $\ell$  is a nonnegative integer,
- $\mathcal{S}$  is a partition of  $[\ell]$ ,
- $d$  is a vector in  $\mathbb{Q}_{\geq 0}^{\mathcal{S}}$ , and
- $y$  is a fractional selection for  $\mathcal{S}$ .

We prove Theorem 1.7 by finding an integral selection  $z$  such that for each arc  $a \in A$ ,

$$x(a) - d_{\max} \leq \sum_{i \in [\ell]: a \in \phi(P_i)} z_i^d \leq x(a) + d_{\max}. \quad (1)$$

Because  $(P_i, \lambda_i)_{i \in [\ell]}$  is an  $s$ -path decomposition, we have

$$x(a) = \sum_{i \in [\ell]: a \in \phi(P_i)} \lambda_i = \sum_{i \in [\ell]: a \in \phi(P_i)} y_i^d,$$

and thus (1) is equivalent to the condition that the  $(y, z)$ -discrepancy of the set  $\{i \in [\ell] : a \in \phi(P_i)\}$  is bounded by  $d_{\max}$ , where the  $(y, z)$ -discrepancy is defined as follows:

**Definition 2.7** (Discrepancy). For a set  $I \subseteq [\ell]$  and fractional selections  $y, z$ , we define the  $(y, z)$ -discrepancy of  $I$  as

$$D_{y,z}(I) := |y^d(I) - z^d(I)|.$$

Clearly, in general there is no integral selection  $z$  such that the  $(y, z)$ -discrepancy for every set  $I \subseteq [\ell]$  is small. (For example, if  $y$  is a fractional selection with only small entries, the index set corresponding to the selected elements in any integral selection has very high discrepancy.) However, we will be able to find an integral selection  $z$  such that the  $(y, z)$ -discrepancy is at most  $d_{\max}$  for all *circular intervals*.

**Definition 2.8** (Circular interval). For  $i, j \in [\ell]$ , the circular interval from  $i$  to  $j$  is the set

- $\{i, i+1, \dots, j\}$  if  $i \leq j$ , and
- $\{i, i+1, \dots, \ell, 1, \dots, j\}$  if  $j < i$ .

A set  $I \subseteq [\ell]$  is called a circular interval if  $I = \emptyset$  or if  $I$  is the circular interval from  $i$  to  $j$  for some  $i, j \in [\ell]$ .

We call the maximum  $(y, z)$ -discrepancy of any circular interval the  $(y, z)$ -interval-discrepancy.

**Definition 2.9** (Interval-discrepancy). For a circular interval  $I \subseteq [\ell]$  and fractional selections  $y, z$ , we define the  $(y, z)$ -interval-discrepancy as

$$D_{y,z} := \max_{\substack{I \subseteq [\ell] \\ \text{circular interval}}} D_{y,z}(I).$$

In Section 5 we will prove the following statement.

**Theorem 2.10.** Consider a WPCS instance  $(\ell, \mathcal{S}, d, y)$ . If  $\mathcal{S}$  is non-interleaving, there exists an integral selection  $z$  with

$$D_{y,z} \leq d_{\max},$$

where  $d_{\max} := \max_{S \in \mathcal{S}} d_S$ . Moreover, it can be computed in time  $O(\ell)$ .

We will also prove that if we are given costs  $c \in \mathbb{Q}^\ell$  and require the selection  $z$  to be no more expensive than the fractional selection  $y$ , then we can still achieve an upper bound of  $2d_{\max}$  on the interval discrepancy. This cost version of the selection problem will be used to prove [Theorem 1.8](#).

**Theorem 2.11.** *Consider a WPCS instance  $(\ell, \mathcal{S}, d, y)$  and let  $c \in \mathbb{Q}^\ell$ . If  $\mathcal{S}$  is non-interleaving, there exists an integral selection  $z$  with*


$$D_{y,z} \leq 2d_{\max} \quad \text{and} \\ c^T z \leq c^T y,$$

where  $d_{\max} := \max_{S \in \mathcal{S}} d_S$ . Moreover, it can be computed in time  $O(\ell \cdot \langle \text{input} \rangle)$ , where  $\langle \text{input} \rangle$  refers to the length of the bit encoding of the input.

To obtain [Theorem 2.10](#), we exploit that  $\mathcal{S}$  is non-interleaving to prove that a simple greedy algorithm yields an integral selection with the desired properties. For showing [Theorem 2.11](#), we first use a common argument in discrepancy theory (see, e.g., [\[LSV86\]](#), and [\[BRS22\]](#) for a recent application in a scheduling context), allowing us to reduce to the special case where each set  $S \in \mathcal{S}$  has size two and the given fractional selection  $y$  fulfills  $y_i = \frac{1}{2}$  for all  $i \in [\ell]$ ; this reduction comes at the cost of losing a factor of two in our discrepancy bound. In the special case we just described, we can for each integral selection  $z$  define a complementary integral selection  $\bar{z}$  that selects for each  $S \in \mathcal{S}$  the element not selected by  $z$ . We observe that whenever  $z$  fulfills the desired discrepancy bounds, then so does  $\bar{z}$ . Moreover, one of the selections  $z$  and  $\bar{z}$  is not more expensive than  $y$ . Hence, we can apply the same algorithm used to prove [Theorem 2.10](#) to this half-integral discrepancy problem in order to find an integral selection  $z$  fulfilling the desired discrepancy bounds and then return the cheaper selection of  $z$  and  $\bar{z}$ . For details, see [Section 5](#).

### 2.3 Bounding interval-discrepancy suffices

Recall that our goal is to find an integral selection  $z$  such that the discrepancy of the set  $\{i \in [\ell] : a \in \phi(P_i)\}$  is small for every arc  $a \in A$ . By [Theorem 2.10](#) and [Theorem 2.11](#) we can achieve a small discrepancy for every circular interval. However, not every set  $\{i \in [\ell] : a \in \phi(P_i)\}$  with  $a \in A$  is a circular interval, as the example in [Figure 3](#) shows.



**Figure 3:** The figure shows an example of an arc-split graph  $H$  (left) with a family of paths  $P_1, \dots, P_6$  from a source-numbered nice  $s$ -path decomposition. The graph  $G$  is shown on the right with an arc  $a$  highlighted in blue. In this example we have  $\{i \in [\ell] : a \in \phi(P_i)\} = \{1, 4\}$ , which is not a circular interval.

Nevertheless, we will prove that [Theorem 2.10](#) and [Theorem 2.11](#) can be used to bound the discrepancy of all sets  $\{i \in [\ell] : a \in \phi(P_i)\}$  with  $a \in A$ . In order to prove this, we define an equivalence relation on the set  $2^{[\ell]}$  of all subsets of  $[\ell]$ . Recall that  $S^t := \{i \in [\ell] : P_i \text{ ends at terminal } t\}$ .

**Definition 2.12** ( $S^t$ -addition and  $S^t$ -removal). *Let  $t \in T$ . We say that a set  $Y \subseteq [\ell]$  arises from  $X \subseteq [\ell]$  by  $S^t$ -addition, if  $X \cap S^t = \emptyset$  and  $Y = X \cup S^t$ . We say that a set  $Y \subseteq [\ell]$  arises from  $X \subseteq [\ell]$  by  $S^t$ -deletion, if  $S^t \subseteq X$  and  $Y = X \setminus S^t$ .*

If a set  $Y \subseteq [\ell]$  arises from  $X \subseteq [\ell]$  by (potentially several)  $S^t$ -additions or  $S^t$ -removals (for some terminals  $t \in T$ ), the  $(y, z)$ -discrepancy of  $X$  and  $Y$  is identical for every integral selection  $z$ . For this reason we will then say that  $X$  and  $Y$  are *discrepancy-equivalent*.

**Definition 2.13** (Discrepancy-equivalent). *Two sets  $X, Y \subseteq [\ell]$  are discrepancy-equivalent if  $Y$  arises from  $X$  by a sequence of  $S^t$ -additions and  $S^t$ -removals for some terminals  $t \in T$ .*

Note that this indeed defines an equivalence relation. We will prove in [Section 4.1](#) the statement below, which says that the set of indices of all paths in  $\mathcal{P}$  containing a copy of the same given arc  $a \in A$  are discrepancy-equivalent to a circular interval. For example, the indices of the paths in [Figure 3](#) that go over (a copy of) arc  $a$ , which are 1 and 4, are discrepancy-equivalent to the circular interval  $\{4, 5, 6, 1\}$ , which is obtained from  $\{1, 4\}$  by  $S^{t_3}$ -addition.

**Lemma 2.14.** *For every arc  $a \in A$ , the set  $\{i \in [\ell] : a \in \phi(P_i)\}$  is discrepancy-equivalent to a circular interval.*

The proof of this statement (in [Section 4](#)) crucially exploits that  $P_1, \dots, P_\ell$  are paths from a source-numbered nice  $s$ -path decomposition of the arc-split graph  $H$ . Using [Lemma 2.3](#) and [Lemma 2.14](#), we can now prove that our discrepancy statements ([Theorem 2.10](#) and [Theorem 2.11](#)) imply our main results ([Theorem 1.7](#) and [Theorem 1.8](#)).

*Proof of Theorem 1.7 and Theorem 1.8.* Let  $(G = (V, A), s, T, d, x)$  be a PSSUF instance. By [Theorem 3.7](#) we can compute in  $O(|V|^2)$  time an arc-split  $(H, \pi, \phi)$  of  $G$  and a nice  $s$ -path decomposition  $(P_i, \lambda_i)_{i \in [\ell]}$  of  $H$  with  $\ell = O(|V|)$ . We define a WPCS instance  $(\ell, \mathcal{S}, d, y)$  as follows. For  $t \in T$  we let  $S^t := \{i \in [\ell] : P_i \text{ is an } s\text{-}t \text{ path}\}$  and we define  $\mathcal{S} := \{S^t : t \in T\}$ . Because  $G$  is acyclic and  $\sum_{i=1}^\ell \lambda_i \chi^{\phi(P_i)} = x$ , every path  $P_i$  with  $i \in [\ell]$  starts at  $s$  and ends at some terminal  $t \in T$ . Therefore,  $\mathcal{S}$  is a partition of  $[\ell]$ . We define the demand vector  $d \in \mathbb{Q}_{\geq 0}^{\mathcal{S}}$  by  $d_{S^t} := d_t$ . For  $i \in [\ell]$  we define  $y_i := \frac{\lambda_i}{d_t}$ , where  $t \in T$  is the terminal where the path  $P_i$  ends. Because  $\sum_{i=1}^\ell \lambda_i \chi^{\phi(P_i)} = x$  and  $x(\delta^-(t)) - x(\delta^+(t)) = d_t$  for all  $t \in T$ , we have  $\sum_{i \in S^t} \lambda_i = d_t$  and thus  $\sum_{i \in S^t} y_i = 1$  for all  $t \in T$ . This shows that  $y$  is a fractional selection for  $\mathcal{S}$ . By [Lemma 2.3](#), the partition  $\mathcal{S}$  is non-interleaving.

Now consider an integral selection  $z$  for  $\mathcal{S}$ . For each  $t \in T$ , we select the  $s$ - $t$  path  $P^t := P_i$ , where  $i$  is the unique element of  $S^t$  with  $z_i = 1$ . Consider the unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  and an arc  $a \in A$ . By [Lemma 2.14](#), the set  $\{i \in [\ell] : a \in \phi(P_i)\}$  is discrepancy equivalent to some circular interval  $I_a$ , implying  $D_{y,z}(\{i \in [\ell] : a \in \phi(P_i)\}) = D_{y,z}(I_a)$ . Using

$$x(a) = \sum_{i \in [\ell] : a \in \phi(P_i)} \lambda_i = \sum_{i \in [\ell] : a \in \phi(P_i)} y_i^d,$$

and

$$\text{flow}_{\mathcal{P}}(a) = \sum_{i \in [\ell] : a \in \phi(P_i)} z_i^d,$$

this implies

$$|x(a) - \text{flow}_{\mathcal{P}}(a)| = \left| \sum_{i \in [\ell] : a \in \phi(P_i)} y_i^d - \sum_{i \in [\ell] : a \in \phi(P_i)} z_i^d \right| = D_{y,z}(\{i \in [\ell] : a \in \phi(P_i)\}) = D_{y,z}(I_a) \leq D_{y,z}.$$

We conclude that [Theorem 1.7](#) follows from the discrepancy statement [Theorem 2.10](#) applied to the WPCS instance  $(\ell, \mathcal{S}, d, y)$ . As an overall runtime we obtain  $O(|V|^2)$ . Moreover, [Theorem 1.8](#) follows from [Theorem 2.11](#) applied to the WPCS instance  $(\ell, \mathcal{S}, d, y)$  and the cost vector  $c \in \mathbb{Q}_{\geq 0}^\ell$  defined by  $c_i := c(P_i) := \sum_{a \in P_i} c(a)$ . As an overall runtime we obtain  $O(|V|^2 + \langle \text{input} \rangle \cdot |V|)$ .  $\square$


The remainder of this paper is structured as follows. In Section 3 we show how to obtain an arc split graph  $H$  with a nice  $s$ -path decomposition and prove Theorem 3.7. In Section 4, we prove that a nice  $s$ -path decomposition has the properties claimed in Lemma 2.3 and Lemma 2.14. Section 5 contains the proof of the two discrepancy statements Theorem 2.10 and Theorem 2.11. Section 6 contains some concluding remarks. Finally, Appendix A discusses the necessity of the acyclicity assumption when dealing with lower bounds.

### 3 Computing a Nice Path Decomposition

We now provide details on how we compute a source-numbered nice path decomposition of a flow  $x$ , which proves Theorem 2.1. Moreover, we expand on the discussion of Section 2.1 and further formalize some of the concepts introduced there.

For a vertex  $v \in V$ , let  $\delta(v) := \delta^+(v) \cup \delta^-(v)$  be the set of all arcs incident to  $v$ . Each planar geometric embedding induces a cyclic ordering  $\pi_v$  of the arcs in  $\delta(v)$ , obtained by traversing the arcs in  $\delta(v)$  in counterclockwise sense. We say that  $b \in \delta(v)$  is the *successor* of  $a \in \delta(v)$  in  $\pi_v$  if, in the geometric embedding,  $b$  is the next arc after  $a$  in counterclockwise direction. In this case, we also say that  $a$  is the *predecessor* of  $b$  in  $\pi_v$ .

A collection  $\pi = (\pi_v)_{v \in V}$  of cyclic orderings of the arcs incident to  $v$  for all  $v \in V$  is called a *combinatorial embedding*. A combinatorial embedding is *planar* if it is induced by some planar geometric embedding. See Figure 4 for an example.




**Figure 4:** A vertex  $v$  with incident arcs  $a_1, \dots, a_6$ , geometrically embedded into the plane. Let  $\pi$  be the induced combinatorial embedding. In the cyclic ordering  $\pi_v$ , the arc  $a_6$  is the predecessor of  $a_1$  and the arc  $a_2$  is the successor of  $a_1$ .

The tuple  $(a_5, a_2, a_3)$  is a  $\pi_v$ -progression because  $a_5, a_2, a_3$  is a subsequence of  $a_5, a_6, a_1, a_2, a_3, a_4$ . The tuple  $(a_5, a_3, a_2)$  is not a  $\pi_v$ -progression.

In order to formally define when two paths are crossing, we introduce the notion of a  $\pi_v$ -progression. Figure 4 shows an example.

**Definition 3.1** ( $\pi_v$ -progression). *Let  $v \in V$  and let  $b_1, \dots, b_k \in \delta(v)$  be  $k$  distinct arcs. Let  $\delta(v) = \{a_1, \dots, a_m\}$ , where the numbering is chosen such that  $a_1 = b_1$  and  $a_{i+1}$  is the successor of  $a_i$  in the cyclic order  $\pi_v$  for all  $i \in [m-1]$ . Then  $(b_1, b_2, \dots, b_k)$  is a  $\pi_v$ -progression if  $b_1, b_2, \dots, b_k$  is a subsequence of  $a_1, \dots, a_m$ .*


We are now ready to formally define when two edge-disjoint paths are crossing. See Figure 5.



**Figure 5:**  $P_1$  and  $P_2$  are crossing because  $(a_1, a_2, b_1, b_2)$  is a  $\pi_v$ -progression.

**Definition 3.2** (Crossing paths). *Let  $H = (V, F)$  be a graph with planar combinatorial embedding  $\pi$ . Two arc-disjoint paths  $P_1, P_2$  are crossing if they have a common mutually internal vertex  $v \in V$  such that  $(a_1, a_2, b_1, b_2)$  or  $(a_2, a_1, b_2, b_1)$  is a  $\pi_v$ -progression, where  $\{a_1\} = P_1 \cap \delta^-(v)$ ,  $\{b_1\} = P_1 \cap \delta^+(v)$ ,  $\{a_2\} = P_2 \cap \delta^-(v)$ , and  $\{b_2\} = P_2 \cap \delta^+(v)$ .*

This notion works well for arc-disjoint paths but fails to capture what would naturally be considered crossing paths if the paths have arcs in common, as illustrated in Figure 6. For this reason, we work with the arc-split graph  $H$  of  $G$ , where arcs are replaced by (a well-chosen arc-dependent number of) parallel copies. This allows for assuming that the paths  $P_1, \dots, P_\ell$  we construct later are disjoint in  $H$ .



**Figure 6:** The left-hand side shows paths  $P_1$  and  $P_2$  that use a common arc  $(v_1, w_1)$  and paths  $P_3$  and  $P_4$  that use a common arc  $(v_2, w_2)$ . We would like to consider  $P_1$  and  $P_2$  as crossing, but not  $P_3$  and  $P_4$ . However, when considering the local cyclic arc-orderings at vertices, we cannot distinguish these two cases. This is why we introduce arc-split graphs, where  $(v_i, w_i)$  is replaced by two copies, each used by one path, as depicted on the right-hand side. This leads to arc-disjoint paths  $P_1$  and  $P_2$  that are crossing according to Definition 3.2. (In the example, they cross at  $w_1$ , and if we swap the red and blue copy of the arc  $(v_1, w_1)$ , they cross at  $v_1$ .) The paths  $P_3$  and  $P_4$  are not crossing.

Arc-split graphs are formally defined as follows, where for some arc  $a = (v, w)$ , we write  $\text{head}(a) = w$  and  $\text{tail}(a) = v$ .

**Definition 3.3** (Arc-split graph  $H$  of  $G$ ). *Let  $G = (V, A)$  be a planar graph. An arc-split graph of  $G$  is a tuple  $(H, \pi)$ , where  $H = (V, F)$  is a planar graph with planar combinatorial embedding  $\pi$ , such that there exists a mapping  $\phi: F \rightarrow A$  fulfilling for any  $f \in F$ :*

$$\text{head}(f) = \text{head}(\phi(f)) \quad \text{and} \quad \text{tail}(f) = \text{tail}(\phi(f)),$$

*and, for all  $a \in A$ , the set  $\phi^{-1}(a) \subseteq F$  is consecutive at both  $\text{head}(a)$  and  $\text{tail}(a)$  in  $\pi$ . We call  $(H, \pi, \phi)$  an arc-split of  $G$ .*

As mentioned above, we want the paths  $P_1, \dots, P_\ell$  to be arc-disjoint in the arc split graph  $H$ . The number of copies of an arc  $a \in A$  in the arc-split graph  $H$  that we construct will be chosen such that each copy is used by some path  $P_i$ . Thus, the paths  $P_1, \dots, P_\ell$  that we construct will form a *nice s-path partitioning* of  $H$ , which is defined as follows.

**Definition 3.4** (Nice  $s$ -path partitioning). *Let  $H = (V, F)$  be a graph with planar combinatorial embedding  $\pi$ , and let  $s \in V$ . A collection of paths  $P_1, \dots, P_\ell \subseteq F$  in  $H$  is called an  $s$ -path partitioning of  $H$  if*

- $P_i$  starts at  $s$  for all  $i \in [\ell]$ , and
- $F = P_1 \dot{\cup} \dots \dot{\cup} P_\ell$ .

*It is called nice if the paths  $P_1, \dots, P_\ell$  are pairwise non-crossing.*

Together with the nice  $s$ -path partitioning of  $H$  we will also compute non-negative weights for the paths  $P_i$  in order to obtain a path decomposition of the flow  $x$ .

**Definition 3.5** (Nice  $s$ -path decomposition). *Let  $(G, s, T, d, x)$  be a PSSUF instance and let  $(H, \pi, \phi)$  be an arc-split of  $G = (V, A)$ . A nice  $s$ -path decomposition of  $H$  is a sequence of tuples  $(P_i, \lambda_i)_{i \in [\ell]}$ , where  $P_1, \dots, P_\ell$  is a nice  $s$ -path partitioning of  $H$  and  $\lambda_i \in \mathbb{Q}_{\geq 0}$  for  $i \in [\ell]$ , such that*

$$\sum_{i=1}^{\ell} \lambda_i \chi^{\phi(P_i)} = x.^2$$

Recall that the paths  $P_1, \dots, P_\ell$  are all paths starting in  $s$ . Thus, the cyclic ordering  $\pi_s$  of the arcs incident to  $s$  naturally induces a cyclic ordering on the paths  $P_1, \dots, P_\ell$ . As mentioned in Section 2, if the numbering of these paths is consistent with this cyclic ordering, we say that  $P_1, \dots, P_\ell$  are *source-numbered*.

**Definition 3.6** (Source-numbered). *An  $s$ -path partitioning is called source-numbered if  $(a_1, \dots, a_\ell)$  is a  $\pi_s$ -progression, where  $a_i$  is the first arc of  $P_i$  for  $i \in [\ell]$ .*

We now prove Theorem 2.1, which we restate here using the just introduced notation, making explicit that we compute a combinatorial embedding of  $H$ .

**Theorem 3.7.** *Let  $(G = (V, A), s, T, d, x)$  be a PSSUF instance. We can compute in  $O(|V|^2)$  time an arc-split  $(H, \pi, \phi)$  of  $G$  and a source-numbered nice  $s$ -path decomposition  $(P_i, \lambda_i)_{i \in [\ell]}$  of  $H$  with  $\ell = O(|V|)$ .*

### 3.1 Proof of Theorem 3.7

Note that source-numbering a set of paths starting at  $s$  simply means that the paths are numbered in counterclockwise sense around  $s$  (starting at an arbitrary path). As this is trivial to do at the end, we will ignore this aspect when constructing a source-numbered nice  $s$ -path decomposition fulfilling the conditions of Theorem 3.7, and thus just focus on finding a nice  $s$ -path decomposition.

There are different ways to obtain nice  $s$ -path decompositions. One option is to start with an arbitrary flow decomposition and then iteratively reduce the number of crossings by uncrossing steps. We present another approach that avoids crossings upfront and leads to a running time that is linear in the total size  $\sum_{i=1}^{\ell} |P_i|$  of the computed paths.

To obtain the desired nice  $s$ -path decomposition in an arc-split graph  $H$  of  $G$ , as claimed by Theorem 3.7, we first decide locally on rules how paths can traverse each vertex  $v \in V \setminus (T \cup \{s\})$ . These rules are set up such that paths that we later construct from them will be non-crossing. We formalize these rules through the notion of a *good  $v$ -wiring*. A good  $v$ -wiring is a sequence of pairs, each coupling one  $v$ -incoming arc  $f \in \delta^-(v)$  with one  $v$ -outgoing arc  $g \in \delta^+(v)$ , together with a maximum load  $\mu_v((f, g))$  with which we can use this pair. The paths in our nice  $s$ -path decomposition that go over  $v$  will use one of these pairs, and the load with which each pair  $(f, g) \in \delta^-(v) \times \delta^+(v)$  is used is given by  $\mu_v((f, g))$ .

**Definition 3.8** (Good  $v$ -wiring). *Let  $(G = (V, A), s, T, d, x)$  be a PSSUF instance, and let  $v \in V \setminus (T \cup \{s\})$ . Moreover, let  $\pi^G$  be a planar combinatorial embedding of  $G$ . A good  $v$ -wiring is a tuple  $(W_v, \mu_v)$ , where  $W_v = ((f_1, g_1), \dots, (f_k, g_k))$  is a sequence of distinct arc pairs  $(f_i, g_i) \in$*

---

<sup>2</sup> $\phi(P_i) := \{\phi(f) : f \in P_i\}$  denotes the image of  $P_i$  by  $\phi$ . Moreover, for a set  $U \subseteq A$ , we denote by  $\chi^U \in \{0, 1\}^A$  the characteristic vector of  $U$ , i.e., the zero-one vector in  $\{0, 1\}^A$  with a 1 precisely in the entries corresponding to arcs in  $U$ .



$\delta^-(v) \times \delta^+(v)$  for  $i \in [k]$ , and  $\mu_v: W_v \rightarrow \mathbb{Q}_{\geq 0}$ , such that

$$\sum_{\substack{i \in [k]: \\ a \in \{f_i, g_i\}}} \mu_v((f_i, g_i)) = x(a) \quad \forall a \in \delta(v),$$

and, for  $i, j \in [k]$  with  $i < j$ , we have that neither  $(f_i, f_j, g_i)$  nor  $(f_i, g_j, g_i)$  is a  $\pi_v^G$ -progression.

Note that forbidding, for  $i, j \in [k]$  with  $i < j$ , the  $\pi_v^G$ -progressions  $(f_i, f_j, g_i)$  and  $(f_i, g_j, g_i)$ , excludes the  $\pi_v^G$ -progressions  $(f_i, f_j, g_i, g_j)$  and  $(f_j, f_i, g_j, f_i)$ .

We start by showing that a good  $v$ -wiring can be computed fast.

**Lemma 3.9.** *Let  $(G = (V, A), s, T, d, x)$  be a PSSUF instance with planar combinatorial embedding  $\pi^G$ , and let  $v \in V \setminus (T \cup \{s\})$ . We can find in  $O(\deg(v))$  time a good  $v$ -wiring  $(W_v, \mu_v)$  with  $|W_v| \leq \deg(v)$ .*

*Proof.* We first describe a simple procedure to obtain the desired result, which, when implemented in a straightforward way, only runs in polynomial time instead of the desired  $O(\deg(v))$  running time bound. In a second step, we expand on how to refine this procedure to obtain the claimed linear time bound.

We construct a good  $v$ -wiring by iteratively finding a pair of arcs  $a \in \delta^-(v)$ ,  $b \in \delta^+(v)$  with  $b$  being  $a$ 's successor in  $\pi_v^G$ , adding the pair  $(a, b)$  to  $W_v$  with a  $\mu_v$ -value equal to  $\min\{x(a), x(b)\}$ , reducing the  $x$ -values of  $a$  and  $b$  correspondingly, and deleting  $a$  and/or  $b$  if their  $x$ -value reached 0. Algorithm 1 formalizes this procedure. In its implementation, we use a circular doubly linked list  $L$  of the arcs in  $\delta(v)$ . This list contains all arcs in  $\delta(v)$  in the order given by  $\pi_v^G$ . For an arc  $a \in L$ , we denote by  $\text{succ}_L(a)$  the arc in  $L$  that comes after  $a$  with respect to the circular ordering  $\pi_v^G$ , i.e., the one after  $a$  in  $L$  in counterclockwise order.

---

**Algorithm 1:** Simple algorithm to create good  $v$ -wiring.

---

```

1 Let  $y = x$ .
2 Let  $L$  be a circular doubly linked list of  $\delta(v)$ .
3  $W_v = ()$ .
4 while  $L \neq \emptyset$  do
5   Find an arc  $a \in L \cap \delta^-(v)$  such that  $b = \text{succ}_L(a) \in \delta^+(v)$ .
6   Add  $(a, b)$  to  $W_v$ .
7   Set  $\mu_v((a, b)) = \min\{y(a), y(b)\}$ .
8    $y(a) = y(a) - \mu_v((a, b))$ .
9    $y(b) = y(b) - \mu_v((a, b))$ .
10  if  $y(a) = 0$  then delete  $a$  from  $L$ .
11  if  $y(b) = 0$  then delete  $b$  from  $L$ .
12 return  $(W_v, \mu_v)$ .
```

---

Note that we have  $y(\delta^+(v)) = y(\delta^-(v))$  during the algorithm. At the start, this holds because  $x(\delta^+(v)) = x(\delta^-(v))$ , and the property is preserved because, whenever a pair gets added to  $W_v$ , we remove the same value from one incoming and one outgoing arc of  $v$ . This implies that, indeed, in every iteration of the algorithm there exists a pair  $(a, b)$  of arcs as desired.

Also observe that this algorithm can clearly be implemented to have polynomial running time. Indeed, at every iteration of the while loop, at least one of the arcs  $a$  or  $b$  will be removed from  $L$ ,

leading to at most  $O(\deg(v))$  many iterations. Moreover, the bottleneck operation in one iteration is finding an arc  $a$  in line 5. A trivial implementation of this step, where we simply check all arcs in  $L$ , leads to a running time of  $O(\deg(v)^2)$ .

Before expanding on how to obtain a faster implementation, we show that  $(W_v, \mu_v)$  is a good  $v$ -wiring. Let  $W_v = ((f_1, g_1), \dots, (f_k, g_k))$ . Note that we clearly have

$$\sum_{\substack{i \in [k]: \\ a \in \{f_i, g_i\}}} \mu_v((f_i, g_i)) = x(a) \quad \forall a \in \delta(v),$$

because the procedure stops once all  $y$ -values, which are the leftover  $x$ -values, are set to zero.

Now consider  $i, j \in [k]$  with  $i < j$ , and we observe that between  $f_i$  and  $g_i$ , in counterclockwise sense, there is neither the arc  $f_j$  nor  $g_j$ . This holds because, when Algorithm 1 adds the pair  $(f_i, g_i)$ , then the arc  $g_i$  appears right after  $f_i$  in the circular doubly linked list  $L$ , i.e.,  $g_i = \text{succ}_L(f_i)$ . As neither  $f_j$  nor  $g_j$  have been deleted at this point because  $i < j$ , they do not lie between  $f_i$  and  $g_i$  in counterclockwise sense. This shows that Algorithm 1 returns a good  $v$ -wiring as claimed.

It remains to show that it can be sped up to run in  $O(\deg(v))$  time. To this end, we show how, after an initial  $O(\deg(v))$ -time preprocessing step, one can find in each iteration a pair  $(a, b) \in \delta^-(v) \times \delta^+(v)$  of arcs in  $L$  with  $b = \text{succ}_L(a)$  to add to the  $v$ -wiring in constant time. To this end, we maintain throughout the iterations of the while loop of Algorithm 1 the set  $Z \subseteq \delta^-(v) \cap L$  of all arcs  $a \in L \cap \delta^-(v)$  with  $\text{succ}_L(a) \in \delta^+(v)$ . The set  $Z$  thus contains the arcs  $a$  of all pairs  $(a, b)$  that could be chosen in line 5. Before the start of the while loop in Algorithm 1, we initialize  $Z$  in  $O(\deg(v))$  time by exhaustively going through all arcs in  $\delta(v)$ . Then, at the end of each iteration of the while loop, we update  $Z$  in constant time as follows. If the arc  $a$  is deleted in line 10 of Algorithm 1, we remove it from  $Z$  and check whether its predecessor in  $L$  should be added to  $Z$ . Otherwise,  $\text{succ}_L(a)$  will get deleted in line 11, and we check whether, after deleting  $\text{succ}_L(a)$ , the arc  $a$  still fulfills that its successor in  $L$  is a  $v$ -outgoing arc, i.e., is in  $\delta^+(v)$ . If so, we keep  $a$  in  $Z$ ; otherwise, we remove it from  $Z$ . One can easily verify that this procedure correctly maintains  $Z$ . Hence, in each iteration of the while loop, we can find a pair  $(a, b) \in Z$  in constant time, and the total time spent on initializing and updating  $Z$  is  $O(\deg(v))$ , as desired. (These time bounds are for example achieved when implementing  $Z$  as a stack.)  $\square$

Starting with good  $v$ -wirings, we now show how to compute an arc-split of  $G$  together with a nice  $s$ -path decomposition as claimed in Theorem 3.7.

*Proof of Theorem 3.7.* We start by obtaining a planar combinatorial embedding  $\pi^G$  of  $G$ , which can be done in linear time, and compute a structured  $s$ -path decomposition of  $x$  as described in Algorithm 2. Note that Algorithm 2 requires a planar combinatorial embedding  $\pi^G$  of  $G$  as it first computes a good  $v$ -wiring  $(W_v, \mu_v)$  for each vertex  $v \in V \setminus (T \cup \{s\})$ . Moreover, also line 6 relies on  $\pi_s^G$ . To perform line 6, we interpret  $\pi_s^G$  as a linear order on  $\delta^+(s)$  instead of a cyclic one by simply choosing an arbitrary arc  $a \in \delta^+(s)$ , declaring  $a$  to be the first arc, and then numbering the remaining arcs with respect to  $\pi_s^G$ , i.e., in counterclockwise order.

Algorithm 2 greedily performs a path-decomposition of the flow  $x$ . The algorithm builds paths step-by-step according to these  $v$ -wirings; whenever we enter a vertex  $v$  through some arc  $f$ , we look for the first pair  $(f, g)$  in the  $v$ -wiring that still has a strictly positive  $\mu_v$ -value and use the arc  $g$  to leave  $v$ . To simplify notation, we drop the index  $v$  from  $\mu_v$  in Algorithm 2. This is without risk of ambiguity, because each pair  $(f, g)$  of consecutive links that we consider is contained in precisely one  $W_v$ , namely the one for  $v = \text{head}(f) = \text{tail}(g)$ ; hence, in this case  $\mu((f, g)) := \mu_v((f, g))$ .

Given the family of pairs  $(Q_i, \lambda_i)_{i \in [\ell]}$  obtained from Algorithm 2, we then construct the claimed arc-split  $(H, \pi, \phi)$  of  $G$  together with a nice  $s$ -path decomposition as follows. Each tuple  $(Q_i, \lambda_i)$ ,

---

**Algorithm 2:** Compute structured flow decomposition of  $x$ .

---

```

1  Compute a good  $v$ -wiring  $(W_v, \mu_v)$  for each  $v \in V \setminus (T \cup \{s\})$ .
2   $y = x$ .
3   $\ell = 0$ .
4  while  $y(\delta^+(s)) > 0$  do
5       $\ell = \ell + 1$ .
6      Let  $a \in \delta^+(s)$  be the first arc with respect to  $\pi_s^G$  in  $\delta^+(s)$  with  $y(a) > 0$ .
7       $Q_\ell = \{a\}$ .
8      while  $\text{head}(a) \notin T$  do
9          Let  $(f, g)$  be the first pair in  $W_v$  with  $f = a$  and  $\mu((f, g)) > 0$ , where  $v = \text{head}(a)$ .;
           $Q_\ell = Q_\ell \cup \{g\}$ .
10          $a = g$ .
11     Let  $\lambda_\ell$  be the smallest value  $\mu((f, g))$  for any pair  $(f, g)$  of consecutive arcs in  $Q_\ell$ .
12     For every pair  $(f, g)$  of consecutive arcs in  $Q_\ell$ , set  $\mu((f, g)) = \mu((f, g)) - \lambda_\ell$ .
13      $y(a) = y(a) - \lambda_\ell \quad \forall a \in Q_\ell$ .
14 return  $(Q_i, \lambda_i)_{i \in [\ell]}$ .
```

---

for  $i \in [\ell]$ , will correspond to one path  $P_i$  in our  $s$ -path decomposition with coefficient being  $\lambda_i$ , i.e., the nice  $s$ -path decomposition we construct will be  $(P_i, \lambda_i)_{i \in [\ell]}$ . We start with an empty graph  $H$  on the vertices  $V$  and then add, one-by-one for  $i$  starting from 1 to  $\ell$ , the paths  $P_i$ . When considering an index  $i \in [\ell]$ , we add, for each arc  $a = (u, v) \in Q_i$ , a fresh arc  $h$  from  $u$  to  $v$  to the graph  $H$  and set  $\phi(h) = a$ . The path  $P_i$  consists of all the arcs added to  $H$  when considering the index  $i$ . Hence,  $P_i$  is a path traversing the same vertices as  $Q_i$  (in the same order), simply on a fresh set of arcs. To conclude the construction, we have to specify how to define the planar combinatorial embedding  $\pi$  of  $H$ . The idea is that we largely inherit the planar combinatorial embedding  $\pi^G$  of  $G$ . More precisely, the copy of the arc  $(u, v) \in Q_i$  that we added to  $H$  will appear in  $\pi_u$  and  $\pi_v$  at the same spot as it appears in  $\pi_u^G$  and  $\pi_v^G$ , respectively. In other words, we define the planar combinatorial embedding such that, for any vertex  $u \in V$ , any subset of non-parallel arcs in  $H$  incident to  $u$  appear in  $\pi_u$  in the precise same order as the order they have in  $\pi_u^G$ . It remains to decide how to order different copies in  $H$  of the same arc  $(u, v)$  of  $G$ . Hence, assume that  $h_1, h_2$  are two arcs in  $H$ , both going from  $u \in V$  to  $v \in V$ , and assume that we added  $h_1$  in an earlier iteration than  $h_2$ . We apply the following rule:

- Within all arcs from  $u$  to  $v$  in  $H$ , the arc  $h_1$  comes before  $h_2$  in  $\pi_u$ , i.e., the arc  $h_2$  is on the counterclockwise side of  $h_1$  within all arcs from  $u$  to  $v$  in  $H$ ,
- Within all arcs from  $u$  to  $v$  in  $H$ , the arc  $h_1$  comes after  $h_2$  in  $\pi_v$ , i.e., the arc  $h_2$  is on the clockwise side of  $h_1$  within all arcs from  $u$  to  $v$  in  $H$ .

This finishes the procedure to compute the arc-split  $(H, \pi, \phi)$  and the  $s$ -path decomposition  $(P_i, \lambda_i)_{i \in [\ell]}$ .

We start by showing the desired bound of  $\ell = O(|V|)$  on the number of paths in our  $s$ -path decomposition. Note that this number of paths is the same as the number of paths  $Q_i$ , computed in Algorithm 2. Each time we compute a path  $Q_i$  in Algorithm 2, one of the values  $\mu_v((f, g))$  is

set to zero for some pair  $(f, g)$  in one of the sets  $W_v$  for some  $v \in V \setminus (T \cup \{s\})$ . We thus have

$$\ell \leq \sum_{v \in V \setminus (T \cup \{s\})} |W_v| = O \left( \sum_{v \in V \setminus (T \cup \{s\})} \deg(v) \right) = O(|A|) = O(|V|),$$

where the first equality follows from  $|W_v| \leq \deg(v)$ , which holds by [Lemma 3.9](#), the second equality is due to the fact that the sum of all degrees of a graph is twice the number of its arcs, and the last inequality holds because  $G$  is planar.

We now discuss the running time of the suggested algorithm. First, note that [Algorithm 2](#) can clearly be implemented to run in  $O(|V|^2)$  time. More precisely, due to [Lemma 3.9](#), computing the good  $v$ -wirings at the beginning of the algorithm takes  $O(\sum_{v \in V \setminus (T \cup \{s\})} \deg(v)) = O(|A|) = O(|V|)$  time. Each iteration of the outer while loop constructs one path  $Q_i$ , and as we construct  $\ell = O(|V|)$  many paths, this while loop has  $O(|V|)$  iterations. Moreover, the inner while loop successively adds arcs to the path  $Q_i$ , each such step taking  $O(1)$  time. Because  $|Q_i| = O(|V|)$  for each path  $Q_i$ , this leads to a total running time of  $O(|V|^2)$  for [Algorithm 2](#). Finally, the construction of the arc-split  $(H, \pi, \phi)$  and nice  $s$ -path decomposition  $(P_i, \lambda_i)_{i \in [\ell]}$  takes time linear in the size of  $H$ , which is bounded by  $O(\sum_{i=1}^{\ell} |Q_i|) = O(|V|^2)$  as desired, where the equality follows from  $\ell = O(|V|)$  and  $|Q_i| \leq |V|$  for  $i \in [\ell]$ .

It remains to show that the arc-split  $(H, \pi, \phi)$  and the  $s$ -path decomposition  $(P_i, \lambda_i)_{i \in [\ell]}$  have the desired properties. First, note that  $\pi$  is indeed a planar combinatorial embedding of  $H$ . This readily follows from the fact that a (geometric) planar embedding of  $H$  corresponding to  $\pi$  can be obtained by first obtaining a (geometric) planar embedding of  $G$  that corresponds to  $\pi^G$  and then replacing arcs of  $G$  by parallel arcs. Also, the properties of  $\phi$  for  $(H, \pi, \phi)$  to be an arc-split of  $G$  (see [Definition 3.3](#)) are clearly fulfilled. Hence,  $(H, \pi, \phi)$  is an arc-split of  $G$ .

Because the arcs of  $H$  are by construction the disjoint union of the paths  $P_i$  for  $i \in [\ell]$ , and all these paths start at  $s$ , we have that  $P_1, \dots, P_{\ell}$  is an  $s$ -path partitioning of  $H$ . To show that it is a nice  $s$ -path partitioning, we have to show that they are pairwise non-crossing. Hence, let  $i, j \in [\ell]$  with  $i < j$  and let  $v \in V$  be a common internal vertex of  $P_i$  and  $P_j$ . We name the arcs of  $P_i$  and  $P_j$  that are incident to  $v$  as follows:

$$\begin{aligned} \{a_i^H\} &= P_i \cap \delta^-(v) & a_i^G &= \phi(a_i^H) \\ \{b_i^H\} &= P_i \cap \delta^+(v) & b_i^G &= \phi(b_i^H) \\ \{a_j^H\} &= P_j \cap \delta^-(v) & a_j^G &= \phi(a_j^H) \\ \{b_j^H\} &= P_j \cap \delta^+(v) & b_j^G &= \phi(b_j^H). \end{aligned}$$

Note that we may have  $a_i^G = a_j^G$ , or  $b_i^G = b_j^G$ , or both. We need to show that neither  $(a_i^H, a_j^H, b_i^H, b_j^H)$  nor  $(a_j^H, a_i^H, b_j^H, b_i^H)$  is a  $\pi_v$ -progression. To this end, we exploit that, by construction of  $\pi_v$ , we have the following property:

$$\begin{aligned} &\text{For any subset } U \subseteq \delta_H(v) \text{ of non-parallel arcs incident to } v, \text{ the ordering of } U \\ &\text{with respect to } \pi_v \text{ is the same as the ordering of } \phi(U) \text{ with respect to } \pi_v^G. \end{aligned} \quad (2)$$

We first observe that  $(a_i^H, a_j^H, b_i^H, b_j^H)$  is not a  $\pi_v$ -progression. If  $a_i^G \neq a_j^G$ , then  $(a_i^G, a_j^G, b_i^G, b_j^G)$  is not a  $\pi_v^G$ -progression, which follows from  $(W_v, \mu_v)$  being a good  $v$ -wiring. By (2), this implies that  $(a_i^H, a_j^H, b_i^H, b_j^H)$  is not a  $\pi_v$ -progression, and therefore neither is  $(a_i^H, a_j^H, b_j^H, b_i^H)$ . Otherwise, if  $a_i^G = a_j^G$ , then, in the ordering  $\pi_v$  of the arcs  $\phi^{-1}(a_i^G)$ , the arc  $a_i^H$  appears after  $a_j^H$ . Again, this implies that  $(a_i^H, a_j^H, b_i^H, b_j^H)$  is not a  $\pi_v$ -progression as  $b_i^H, b_j^H \notin \phi^{-1}(a_i^G)$ .

Analogously, we now show that  $(a_j^H, a_i^H, b_j^H, b_i^H)$  is not a  $\pi_v$ -progression. If  $b_i^G \neq b_j^G$ , then  $(a_i^G, b_j^G, b_i^G)$  is not a  $\pi_v^G$ -progression, which follows from  $(W_v, \mu_v)$  being a good  $v$ -wiring. By (2), this implies that  $(a_i^H, b_j^H, b_i^H)$  is not a  $\pi_v$ -progression, and therefore neither is  $(a_j^H, a_i^H, b_j^H, b_i^H)$ . Otherwise, if  $b_i^G = b_j^G$ , then, in the ordering  $\pi_v$  of the arcs  $\phi^{-1}(b_i^G)$ , the arc  $b_i^H$  appears before  $b_j^H$ . Again, this implies that  $(a_j^H, a_i^H, b_j^H, b_i^H)$  is not a  $\pi_v$ -progression.

Thus,  $P_1, \dots, P_\ell$  are pairwise non-crossing, which shows that  $P_1, \dots, P_\ell$  is a nice  $s$ -path partitioning. To finish the proof, it remains to show that

$$\sum_{i=1}^{\ell} \lambda_i \chi^{\phi(P_i)} = x. \quad (3)$$

This holds because Algorithm 2 performs greedily an  $s$ -path decomposition that uses up all the  $\mu_v$ -values of the good  $v$ -wirings  $(W_v, \mu_v)$  for each  $v \in V \setminus (T \cup \{s\})$ . Indeed, this path decomposition is just one way to perform a greedy path decomposition of the original flow  $x$ . (Recall that a good  $v$ -wiring  $(W_v, \mu_v)$  with  $W_v = ((f_1, g_1), \dots, (f_k, g_k))$  satisfies  $\sum_{i \in [k]: a \in \{f_i, g_i\}} \mu_v((f_i, g_i)) = x(a)$  for all  $a \in \delta(v)$ , and thus the  $\mu_v$ -values indeed represent the original flow.) Moreover, a greedy path decomposition of a flow in an acyclic graph always decomposes the full flow value on all arcs. This implies Equation (3).  $\square$

## 4 Nice path partitionings have nice properties

In this section we prove Lemma 2.3 and Lemma 2.14, which are both statements about source-numbered nice  $s$ -path partitionings stemming from a PSSUF instance  $(G, s, T, d, x)$ . As these statements do not depend on the fractional flow  $x$  or the demands  $d$ , we will only fix, throughout this section, an acyclic graph  $G = (V, A)$  and an arc-split  $(H = (V, F), \pi, \phi)$  of  $G$  together with a source-numbered nice  $s$ -path partitioning  $P_1, \dots, P_\ell$  of  $H$ . Moreover, we assume throughout that no end vertex of any path of the  $s$ -path partitioning has outgoing arcs, because these end vertices correspond to terminals, which do not have outgoing arcs by assumption. We thus also call here such end vertices *terminals* and let  $T$  be the set of all such terminals.

Our proofs of the statements Lemma 2.3 and Lemma 2.14 both rely on a crucial property how the paths  $P_1, \dots, P_\ell$  can interact. This property is described in terms of how two paths that share a common internal vertex split the planar embedding in two parts. We represent these two parts as a tri-coloring, which we introduce first. Using this tri-coloring, we first prove Lemma 2.14 in Section 4.1, and then Lemma 2.3 in Section 4.2.


**Tri-coloring** Let us fix a planar geometric embedding of  $H$  realizing the combinatorial embedding  $\pi$ . Let  $\tilde{P}$  and  $\bar{P}$  be two distinct paths of the nice  $s$ -path partitioning that share a vertex  $v \in V \setminus \{s\}$ . Moreover, let  $C$  be the subgraph of  $H$  obtained from taking the union of the  $s$ - $v$  subpaths of  $\tilde{P}$  and  $\bar{P}$ . Observe that  $C$  is a connected graph in which every vertex  $v$  has even degree  $|\delta(v)|$ . Thus, the planar dual  $C^*$  of the undirected graph obtained from  $C$  by ignoring the orientation of the arcs is bipartite.<sup>3</sup> Therefore, we can color the faces of  $C$  with two colors, red and blue. Because the geometric embedding is planar, this yields a red-blue coloring of all arcs of  $H$  except for  $C$ . Finally, we color the arcs in  $C$  in black. We call this coloring of all arcs of  $H$  the *tri-coloring of  $H$  induced by  $\tilde{P}$ ,  $\bar{P}$ , and  $v$* . (It is unique up to exchanging red and blue.) See Figure 7. When we say that a

<sup>3</sup>This is well known and can be seen as follows. In the undirected graph obtained from  $C$  by ignoring the orientation of the arcs, all vertex degrees are even and hence every cut contains an even number of edges. As cycles in the dual graph  $C^*$  correspond to cuts in this graph,  $C^*$  has no odd cycle and is thus bipartite.





Thus, a path entering  $u$  with an arc of color red or blue and leaving  $u$  with an arc of the other of these colors implies a crossing among the paths  $P_1, \dots, P_\ell$ , which violates that they are a nice  $s$ -path partitioning.



**Figure 8:** Illustration of the neighborhood of a vertex  $u$  in  $C$ . The three different possible situations (up to symmetry and exchange of the colors red and blue) are illustrated in the picture. Here the solid and dashed arcs belong to the  $s$ - $v$  subpaths of  $\bar{P}$  and  $\tilde{P}$ , respectively. A path changing its color from a red face to a blue face (or the other way around) at  $u$  crosses one of the paths  $\bar{P}$  and  $\tilde{P}$ .

So assume  $|C \cap \delta(u)| = 4$ . Hence, the four primal faces adjacent to  $u$  have alternating colors between red and blue. (More precisely, as we color the arcs, the color of a face corresponds to the color of all arcs within that face.) Let us denote by  $\bar{a}$  and  $\tilde{a}$  the arc of  $\bar{P}$  and  $\tilde{P}$ , respectively, that enters  $u$ , and by  $\bar{b}$  and  $\tilde{b}$  the arc of  $\bar{P}$  and  $\tilde{P}$ , respectively, that leaves  $u$ , i.e.,  $\{\bar{a}\} = \bar{P} \cap \delta^-(u)$ ,  $\{\tilde{a}\} = \tilde{P} \cap \delta^-(u)$ ,  $\{\bar{b}\} = \bar{P} \cap \delta^+(u)$ , and  $\{\tilde{b}\} = \tilde{P} \cap \delta^+(u)$ . Because  $\bar{P}$  and  $\tilde{P}$  do not cross, we can describe the structure of  $\pi_u$  up to symmetries as follows:

$$\begin{aligned} \pi_u &= (\bar{a}, \underbrace{\dots}_{\text{red}}, \bar{b}, \underbrace{\dots}_{\text{blue}}, \tilde{a}, \underbrace{\dots}_{\text{red}}, \tilde{b}, \underbrace{\dots}_{\text{blue}}) \quad \text{or} \\ \pi_u &= (\bar{a}, \underbrace{\dots}_{\text{red}}, \tilde{b}, \underbrace{\dots}_{\text{blue}}, \bar{b}, \underbrace{\dots}_{\text{red}}, \tilde{a}, \underbrace{\dots}_{\text{blue}}). \end{aligned}$$

Again, a path  $P$  would have to cross one of the paths  $\bar{P}$  and  $\tilde{P}$  when changing color in  $u$ . See Figure 8 for a visualization. This is not possible in a nice  $s$ -path partitioning.

We conclude that  $P$  can change its color only at the vertex  $v$ . Because  $P$  is a path, it can thus change its color at most once.

Now suppose  $P \in \{\bar{P}, \tilde{P}\}$ , and let  $Q \in \{\bar{P}, \tilde{P}\}$  be the other path. By definition,  $P$  is black until  $v$  and becomes red or blue after that. A second change of color is thus only possible at a vertex that  $P$  visits after  $v$ . Because  $P$  cannot visit a vertex twice, such a vertex is not visited by the  $s$ - $v$  subpath of  $P$ . Thus, if  $P$  changes its color at a vertex  $u \neq v$ , the vertex  $u$  must be an internal vertex of the  $s$ - $v$  subpath of  $Q$ , and we have  $|C \cap \delta(u)| = 2$ . However, at such a vertex  $u$ , the path  $P$  cannot change color as this would imply a crossing of  $P$  and  $Q$  at  $u$ . Thus, after changing color from black to red or blue at vertex  $v$ , the path  $P$  does not change color again, as desired.  $\square$

#### 4.1 Characterization of paths sharing an arc

We now derive Lemma 2.14 from Lemma 4.1. For convenience, we restate the lemma below and recall that, for  $t \in T$ , the set

$$\mathcal{P}^t := \{P_i : i \in [\ell] \text{ and } P_i \text{ ends at terminal } t\},$$

are all paths in the nice  $s$ -path partitioning that end at terminal  $t$ .

**Lemma 2.14.** *For every arc  $a \in A$ , the set  $\{i \in [\ell] : a \in \phi(P_i)\}$  is discrepancy-equivalent to a circular interval.*

*Proof.* Fix some arc  $a = (v, w) \in A$  and a geometric planar embedding of  $H$  realizing the combinatorial embedding  $\pi$ . Let

$$I := \{i \in [\ell] : a \in \phi(P_i)\} \text{ and} \\ \mathcal{P} := \{P_i : i \in I\}$$


be the indices that we want to show to be discrepancy-equivalent to a circular interval, and the corresponding paths, respectively.

If  $v = s$ , then  $I$  is a circular interval because the paths  $P_1, \dots, P_\ell$  are source-numbered and the arcs in  $\phi^{-1}(a)$  appear consecutively in  $\pi_s$  (see [Definition 3.3](#), i.e., the definition of an arc-split graph). Moreover, if  $|\mathcal{P}| \leq 1$  the statement of the lemma is trivial. Hence, we will now assume  $s \neq v$  and  $|\mathcal{P}| \geq 2$ .

By definition of  $\mathcal{P}$ , every path in  $\mathcal{P}$  uses an arc of  $\phi^{-1}(a) \subseteq \delta^+(v)$ . Since  $v \neq s$ , also every path of  $\mathcal{P}$  uses an arc of  $\delta^-(v)$ . We denote these arcs as  $F_{\text{in}}$ , i.e.,  $F_{\text{in}} \subseteq F$  contains all arcs that belong to a path of  $\mathcal{P}$  and enter  $v$ .

Recall that by the definition of an arc-split graph, the arcs in  $\phi^{-1}(a)$  appear consecutively in  $\pi_v$ . Consider the first arc of  $F_{\text{in}}$  appearing after the arcs of  $\phi^{-1}(a)$  in clockwise direction in the geometric embedding and let  $P_{\text{clo}}$  be the unique path in  $\mathcal{P}$  using this arc. (Uniqueness follows from the fact that the different paths in  $\mathcal{P}$  are arc-disjoint by definition of an  $s$ -path partitioning.) Similarly, consider the first arc of  $F_{\text{in}}$  after the arcs of  $\phi^{-1}(a)$  in counterclockwise direction and let  $P_{\text{cc}}$  be the unique path using this arc. Because  $|\mathcal{P}| \geq 2$ , also  $|F_{\text{in}}| \geq 2$  and, thus, the paths  $P_{\text{cc}}$  and  $P_{\text{clo}}$  exist and are distinct.

Both  $P_{\text{clo}}$  and  $P_{\text{cc}}$  start in  $s$  and contain the vertex  $v$ . We consider the tri-coloring of  $H$  induced by  $P_{\text{clo}}$ ,  $P_{\text{cc}}$ , and  $v$ . Note that, for each terminal  $t \in T$ , either all arcs in  $\delta^-(t)$  are red or all arcs in  $\delta^-(t)$  are blue. This follows from the fact that terminals cannot lie in the interior of the paths  $P_{\text{clo}}$  or  $P_{\text{cc}}$ , because they do not have outgoing arcs. In particular,  $v$  cannot be a terminal, because it has an outgoing arc  $a = (v, w)$ . We call a terminal  $t \in T$  *red* if all arcs in  $\delta^-(t)$  are red, and we call it *blue* if all arcs in  $\delta^-(t)$  are blue. Analogously, the vertex  $w := \text{head}(a)$  does not lie on the  $s$ - $v$  subpaths of either  $P_{\text{clo}}$  or  $P_{\text{cc}}$  because both paths  $P_{\text{clo}}$  and  $P_{\text{cc}}$  contain some arc from  $\phi^{-1}((v, w))$ , and  $w$  cannot appear twice on any path. Hence, also the arcs in  $\delta(w)$  are either all red or all blue. We assume without loss of generality that the arcs in  $\delta(w)$  are all colored red. See [Figure 9](#) for an illustration.



**Figure 9:** Illustration of the neighborhood of  $v$ . The solid arcs are arcs of  $P_{\text{clo}}$  and the dashed arcs belong to  $P_{\text{cc}}$ .

Let  $\mathcal{Q}$  be the set of paths containing  $P_{\text{clo}}$ ,  $P_{\text{cc}}$ , and all paths  $P \in \{P_1, \dots, P_\ell\}$  whose first arc is colored blue. By construction of the tri-coloring and since  $s$  is adjacent to only two black arcs, the index set  $J := \{i \in [\ell] : P_i \in \mathcal{Q}\}$  is a circular interval. We claim that the set  $I$  of indices of paths in  $\mathcal{P}$  is discrepancy-equivalent to the set  $J$  of indices of paths in  $\mathcal{Q}$ . To this end, we will show that:

- (i) For every red terminal  $t \in T$ , we have  $\mathcal{P}^t \cap \mathcal{P} = \mathcal{P}^t \cap \mathcal{Q}$ .
- (ii) For every blue terminal  $t \in T$ , we have  $\mathcal{P}^t \cap \mathcal{P} = \emptyset$  and  $\mathcal{P}^t \subseteq \mathcal{Q}$ .


Hence,  $J$  can be obtained from  $I$  by  $S^t$ -addition for all terminals  $t \in T$  that are colored blue, where  $S^t := \{i \in [\ell] : P_i \in \mathcal{P}^t\}$ .

It remains to prove (i) and (ii). By Lemma 4.1, a path  $P \in \{P_1, \dots, P_\ell\}$  can only change its color at  $v$ . We show that this implies that  $P$  can only change its color from either blue or black to red, but not the other way around. Hence, once a path is red, it stays red. Because all arcs incident with  $w$  are red, including all arcs in  $\phi^{-1}(a)$ , this implies (ii). More concretely, given a blue terminal  $t$ , a path  $P \in \mathcal{P}^t$  can never use a red arc, because it would then stay red and can therefore not end in a blue terminal. Thus,  $P \in \mathcal{Q}$ . Moreover,  $P \notin \mathcal{P}$  since otherwise it contains  $w$ , which only has red arcs incident with it. Note that this also implies that both  $P_{\text{clo}}$  and  $P_{\text{cc}}$  do not have a blue terminal since they use an arc of  $\phi^{-1}(a)$ , which is red.

To show that a path can only change the color from blue or black to red but not the other way around, remember that paths can only change their color in  $v$  by Lemma 4.1. The vertex  $v$  has only two adjacent black arcs (one incoming arc each from  $P_{\text{clo}}$  and  $P_{\text{cc}}$ ). Both  $P_{\text{clo}}$  and  $P_{\text{cc}}$  use an arc of  $\phi^{-1}(a)$  by definition and all the arcs of  $\phi^{-1}(a)$  appear consecutively in the embedding. Thus,  $\pi_v$  has the following form, where  $a_{\text{clo}}$  and  $a_{\text{cc}}$  are the arcs of  $P_{\text{clo}}$  and  $P_{\text{cc}}$ , respectively, that enter  $v$ , and  $b_{\text{clo}}$  and  $b_{\text{cc}}$  are the arcs of  $P_{\text{clo}}$  and  $P_{\text{cc}}$ , respectively, that leave  $v$ , i.e.,  $\{a_{\text{clo}}\} = P_{\text{clo}} \cap \delta^-(v)$ ,  $\{b_{\text{clo}}\} = P_{\text{clo}} \cap \delta^+(v)$ ,  $\{a_{\text{cc}}\} = P_{\text{cc}} \cap \delta^-(v)$ , and  $\{b_{\text{cc}}\} = P_{\text{cc}} \cap \delta^+(v)$ :

$$\pi_v = (a_{\text{clo}}, \underbrace{\dots}_{\text{red}}, \underbrace{b_{\text{clo}}, \dots, b_{\text{cc}}}_{\subseteq \phi^{-1}(a)}, \underbrace{\dots}_{\text{red}}, a_{\text{cc}}, \underbrace{\dots}_{\text{blue}}). \quad (4)$$

One can even show that all arcs of  $\phi^{-1}(a)$  lie counterclockwise between  $b_{\text{clo}}$  and  $b_{\text{cc}}$ , but since it is not needed for the proof, we just use the simpler fact that  $b_{\text{clo}}, b_{\text{cc}} \in \phi^{-1}(a)$  and thus also all arcs which are counterclockwise between them are contained in  $\phi^{-1}(a)$ . See Figure 10 for an illustration of the neighborhood of  $v$ .



**Figure 10:** Illustration of the arcs  $a_{\text{clo}}, b_{\text{clo}}, a_{\text{cc}}, b_{\text{cc}}$  incident to  $v$ . The solid arcs are arcs of  $P_{\text{clo}}$  and the dashed arcs belong to  $P_{\text{cc}}$ .

To change the color from red to blue, a path cannot use an arc of  $\phi^{-1}(a)$ , since they are all red and leave  $v$ . This means that any path changing color from red to blue needs to cross either  $P_{\text{clo}}$  or  $P_{\text{cc}}$ , which is forbidden by the definition of a nice  $s$ -path partitioning.

It remains to prove (i). To this end we show that every  $P \in \{P_1, \dots, P_\ell\}$  satisfies the following.

- (a) If  $P$  enters  $v$  via a red arc, then it does not use an arc from  $\phi^{-1}(a)$ , i.e.,  $P \notin \mathcal{P}$ .
- (b) If  $P$  changes its color from blue to red, then it must use an arc from  $\phi^{-1}(a)$ , i.e.,  $P \in \mathcal{P}$ .

Then (i) follows, because (a) implies that every path in  $\mathcal{P}$ , i.e., every path using an arc in  $\phi^{-1}(a)$  from  $v$  to  $w$ , must change its color at  $v$  from blue to red and thus must be contained in  $\mathcal{Q}$ . Hence,  $\mathcal{P} \subseteq \mathcal{Q}$ , and thus  $\mathcal{P}^t \cap \mathcal{P} \subseteq \mathcal{P}^t \cap \mathcal{Q}$  for every terminal  $t \in T$ . (Here we used that  $P_{\text{cc}}$  and  $P_{\text{clo}}$  are contained in  $\mathcal{Q}$  by definition.) Now consider a path  $P \in \mathcal{P}^t \cap \mathcal{Q}$  for some red terminal  $t \in T$ . Then (b) implies that  $P$  uses an arc from  $\phi^{-1}(a)$ , i.e., we have  $P \in \mathcal{P}$ . (Note that  $P_{\text{clo}}$  and  $P_{\text{cc}}$  are contained in  $\mathcal{P}$  by definition.) Thus, it remains to prove (a) and (b).

First, (a) follows by definition of  $P_{\text{clo}}$  and  $P_{\text{cc}}$ . If a path enters  $v$  using a red arc, this arc lies either clockwise between  $b_{\text{clo}}$  and  $a_{\text{clo}}$  or counterclockwise between  $b_{\text{cc}}$  and  $a_{\text{cc}}$ , see Figure 9 and Equation (4). This is because the red arcs of  $\delta(v)$  are precisely those that are counterclockwise between  $a_{\text{clo}}$  and  $a_{\text{cc}}$ ; moreover, all arcs counterclockwise between  $b_{\text{clo}}$  and  $b_{\text{cc}}$  are arcs of  $\phi^{-1}(a)$ , which leave  $v$ . By definition of  $P_{\text{clo}}$ , there is no path entering  $v$  clockwise between  $b_{\text{clo}}$  and  $a_{\text{clo}}$  and using an arc of  $\phi^{-1}(a)$  afterward. Analogously, by definition of  $P_{\text{cc}}$ , there is no path entering counterclockwise between  $b_{\text{cc}}$  and  $a_{\text{cc}}$  and then using an arc of  $\phi^{-1}(a)$ . This shows (a).

Finally, (b) follows because paths are non-crossing. By Lemma 4.1 a path can only change its color in  $v$ . Now, assume for the sake of deriving a contradiction that there is a path  $P$  that enters  $v$  via a blue arc and leaves  $v$  via a red one not in  $\phi^{-1}(a)$ . Such a path either crosses  $P_{\text{clo}}$  (when it leaves clockwise between  $b_{\text{clo}}$  and  $a_{\text{clo}}$ ) or  $P_{\text{cc}}$  (when it leaves counterclockwise between  $b_{\text{cc}}$  and  $a_{\text{cc}}$ )—see also Figure 9 and Equation (4)—which is a contradiction.  $\square$

## 4.2 Nice $s$ -path partitionings are non-interleaving

We now show Lemma 2.3, which we restate here for convenience.

**Lemma 2.3.** *For  $t \in T$  let  $S^t := \{i \in [\ell] : P_i \in \mathcal{P}^t\}$ . Then the partition  $\{S^t : t \in T\}$  of  $[\ell]$  is non-interleaving.*

*Proof.* Assume for the sake of deriving a contradiction that there are two interleaving sets  $S^{t_1}$  and  $S^{t_2}$ , for some  $t_1, t_2 \in T$ , with  $a_i, b_i \in S^{t_i}$  for  $i \in [2]$  such that  $a_1 < a_2 < b_1 < b_2$ .

Consider the paths  $P_{a_1}$  and  $P_{b_1}$ , which both go to the same terminal  $t_1$ . We consider the tri-coloring of  $H$  induced by  $P_{a_1}$ ,  $P_{b_1}$ , and  $t_1$ , see Figure 7 with  $v = t_1$ . By Lemma 4.1, a path can change the color only once and only at  $t_1$ . Because  $t_1$  does not have any outgoing arcs, we get that no path can change its color. As paths are source-numbered,  $P_{a_2}$  and  $P_{b_2}$  start with arcs of different color. Thus, their joint terminal  $t_2$  needs to lie on either  $P_{a_1}$  or  $P_{b_1}$ , because for all vertices  $w$  not lying on one of these two paths, the arcs in  $\delta(w)$  have all the same color. However, as  $t_2$  does not have any outgoing arcs and is disjoint from  $t_1$ , it cannot lie on either  $P_{a_1}$  or  $P_{b_1}$ , leading to a contradiction.  $\square$

## 5 Interval-discrepancy on non-interleaving partitions

In this section we prove our two main discrepancy statements, Theorem 2.10 and Theorem 2.11.

### 5.1 The case without costs (Proof of Theorem 2.10)

For convenience, we first recall Theorem 2.10.

**Theorem 2.10.** *Consider a WPCS instance  $(\ell, \mathcal{S}, d, y)$ . If  $\mathcal{S}$  is non-interleaving, there exists an integral selection  $z$  with*

$$D_{y,z} \leq d_{\max},$$

where  $d_{\max} := \max_{S \in \mathcal{S}} d_S$ . Moreover, it can be computed in time  $O(\ell)$ .

In order to prove Theorem 2.10, we show the following bound on the discrepancy of the prefix intervals  $[i]$  with  $i \in [\ell]$ .

**Theorem 5.1.** *Consider a WPCS instance  $(\ell, \mathcal{S}, d, y)$ . If  $\mathcal{S}$  is non-interleaving, then there exists*

an integral selection  $z$  with

$$D_{y,z}([i]) \leq \frac{d_{\max}}{2}$$

for all  $i \in [\ell]$ , where  $d_{\max} := \max_{S \in \mathcal{S}} d_S$ . Moreover, it can be computed in time  $O(\ell)$ .

Before proving [Theorem 5.1](#), we first observe that it implies [Theorem 2.10](#). For any interval  $I = \{i, i+1, \dots, j\} \subseteq [\ell]$ , we have

$$D_{y,z}(I) = D_{y,z}([j] \setminus [i-1]) \leq D_{y,z}([i-1]) + D_{y,z}([j]) \leq d_{\max}.$$

If  $I$  is a circular interval, then either  $I$  is an interval or the complement  $[\ell] \setminus I$  of  $I$  is an interval. In the latter case we have  $D_{y,z}(I) = D_{y,z}([\ell] \setminus I) \leq d_{\max}$ .

Hence, it remains to prove [Theorem 5.1](#), which we do in the remainder of this section. We fix a WPCS instance  $(\ell, \mathcal{S}, d, y)$  such that  $\mathcal{S}$  is non-interleaving.

For a given integral selection  $z \in \{0, 1\}^{\mathcal{S}}$ , we define the values

$$D^k := y^d([k]) - z^d([k]) \quad \text{for } k \in \{0, \dots, \ell\}.$$

Then  $D_{y,z}([k]) = |D^k|$ . We now show that the difference of  $D^i$  and  $D^j$  for  $i, j \in S \in \mathcal{S}$  depends only on the restriction of  $z$  to  $S$  and not on any other entry of the vector  $z$ . This is a key consequence of the fact that  $\mathcal{S}$  is non-interleaving.

**Lemma 5.2.** *Let  $z \in \{0, 1\}^{\mathcal{S}}$  be an integral selection. Let  $S \in \mathcal{S}$  and  $i, j \in S$  with  $i \leq j$ . Then*

$$D^j - D^{i-1} = \sum_{\substack{m \in S: \\ i \leq m \leq j}} (y_m^d - z_m^d).$$

*Proof.* Consider a set  $R \in \mathcal{S} \setminus \{S\}$ . Because  $\mathcal{S}$  is non-interleaving, we have either  $R \cap \{i, i+1, \dots, j\} = \emptyset$  or  $R \subseteq \{i, i+1, \dots, j\}$ . In both cases we have  $\sum_{m \in R: i \leq m \leq j} (y_m^d - z_m^d) = 0$ , and thus

$$D^j - D^{i-1} = \sum_{m=i}^j (y_m^d - z_m^d) = \sum_{U \in \mathcal{S}} \sum_{\substack{m \in U: \\ i \leq m \leq j}} (y_m^d - z_m^d) = \sum_{\substack{m \in S: \\ i \leq m \leq j}} (y_m^d - z_m^d).$$

□

We now describe the algorithm which we use to compute an integral selection  $z$  with the desired properties. The algorithm fixes the entries  $z_j$  from  $j = 1$  to  $j = \ell$  one by one in this order, while maintaining the desired prefix discrepancy bound  $|D^j| \leq \frac{d_{\max}}{2}$ .

Suppose the currently considered entry is  $z_j$  with  $j \in S \in \mathcal{S}$ . If, for some  $k \in S$  with  $k < j$ , the entry  $z_k$  has already been fixed to 1 in an earlier iteration, the algorithm has to choose  $z_j = 0$  in order to obtain an integral selection. Similarly, if  $j$  is the largest element of  $S$ , and we have chosen  $z_k = 0$  for all other elements  $k$  of  $S$ , the algorithm must set  $z_j = 1$  to obtain an integral selection. We will use [Lemma 5.2](#) to prove that in both of these cases we have  $|D^j| \leq \frac{d_{\max}}{2}$  (using that prior prefix discrepancies are fine, i.e.,  $|D^i| \leq \frac{d_{\max}}{2}$  for  $i < j$ ). If neither of the above two cases applies, our algorithm greedily fixes the currently considered entry  $z_j$  to ensure  $|D^j| \leq \frac{d_{\max}}{2}$ . This is possible because  $D^{j-1} \in [-\frac{d_{\max}}{2}, \frac{d_{\max}}{2}]$  and the two different values that  $D^j$  can attain, depending on the choice of  $z^j \in \{0, 1\}$ , differ by at most  $d_S \leq d_{\max}$  from each other, where one of the two values lies below  $D^{j-1}$  and the other one lies above  $D^{j-1}$ . A formal description of the algorithm is given by [Algorithm 3](#). Note that the runtime of the algorithm is  $O(\ell)$ , since every iteration of the

---

**Algorithm 3:** Computing an integral selection  $z$  with small interval-discrepancy.

---

```

for  $j = 1, \dots, \ell$  do
    Let  $S \in \mathcal{S}$  such that  $j \in S$ .
    If  $z_k = 1$  for some  $k \in S$  with  $k < j$ , set  $z_j := 0$ .
    If  $j$  is the largest element of  $S$  and  $z_k = 0$  for all  $k \in S$  with  $k < j$ , set  $z_j := 1$ .
    If none of the above two cases apply, set
        
$$z_j := \begin{cases} 1 & \text{if } D^{j-1} + y_j^d - d_S \geq -\frac{d_{\max}}{2} \\ 0 & \text{otherwise,} \end{cases}$$

    where  $D^{j-1} := y^d([j-1]) - z^d([j-1])$ .
return  $z$ .

```

---

for-loop can be executed in constant time ( $D^{j-1}$  does not need to be recomputed from scratch in iteration  $j$  but can be obtained from the previously computed  $D^{j-2}$  in constant time).

The vector  $z$  returned by [Algorithm 3](#) is an integral selection by construction. It remains to prove that it fulfills the discrepancy bound claimed by [Theorem 5.1](#).

**Lemma 5.3.** *Let  $z$  be the integral selection returned by [Algorithm 3](#). For all  $j \in \{0, \dots, \ell\}$  we have  $D_{y,z}([j]) = |D^j| \leq \frac{d_{\max}}{2}$ .*

*Proof.* We prove the lemma by induction on  $j$ , starting with  $j = 0$ . We have  $D^0 = 0$  and hence assume  $j > 0$ . As in [Algorithm 3](#), we let  $S \in \mathcal{S}$  such that  $j \in S$ . We distinguish the same three cases as the algorithm.

First, suppose  $z_k = 1$  for some  $k \in S$  with  $k < j$ . Let  $i$  be the minimal element of  $S$ . By [Lemma 5.2](#), we have

$$D^k = D^{i-1} + y^d(\{m \in S : i \leq m \leq k\}) - d_S, \text{ and} \\ D^j = D^{i-1} + y^d(\{m \in S : i \leq m \leq j\}) - d_S.$$

Since  $y$  is a selection, we have  $y^d(\{m \in S : i < m \leq j\}) \leq d_S$ . This implies  $D^j \leq D^{i-1}$ . Moreover, because  $y^d \geq 0$  and  $k < j$ , we get  $D^k \leq D^j$ . But since  $|D^{i-1}| \leq \frac{d_{\max}}{2}$  and  $|D^k| \leq \frac{d_{\max}}{2}$  by induction, we get  $|D^j| \leq \frac{d_{\max}}{2}$ .

Now suppose  $j$  is the largest element of  $S$  and  $z_k = 0$  for all  $k \in S$  with  $k < j$ . Let  $i$  be the minimal element of  $S$ . Then [Lemma 5.2](#) implies

$$D^j - D^{i-1} = \sum_{m \in S: i \leq m \leq j} (y_m^d - z_m^d) = \sum_{m \in S} (y_m^d - z_m^d) = 0,$$

where the second equation follows from the fact that  $i$  is the smallest element of  $S$  and  $j$  the largest one, and the last equality holds due to  $y(S) = 1$  and  $z(S) = 1$  because  $y$  and  $z$  is a fractional and integral selection, respectively. By the induction hypothesis, we conclude  $|D^j| = |D^{i-1}| \leq \frac{d_{\max}}{2}$ .

Finally, we assume that none of the above cases applies. If we set  $z_j = 1$ , then we have  $D^j = D^{j-1} + y_j^d - z_j^d = D^{j-1} + y_j^d - d_S \geq -\frac{d_{\max}}{2}$ , where we used the choice of  $z_j$  in [Algorithm 3](#). Moreover, because  $y_j^d = y_j d_S \leq d_S$ , we have  $D^{j-1} \geq D^j$ , which implies  $D^j \leq \frac{d_{\max}}{2}$  by the induction



hypothesis. If we set  $z_j = 0$ , we have  $D^j \geq D^{j-1}$ , which implies  $D^j \geq -\frac{d_{\max}}{2}$  by the induction hypothesis. Moreover, by the choice of  $z_j$  in [Algorithm 3](#) we have  $D^{j-1} + y_j^d - d_S \leq -\frac{d_{\max}}{2}$ , implying  $D^j = D^{j-1} + y_j^d \leq d_S - \frac{d_{\max}}{2} \leq \frac{d_{\max}}{2}$ , where we used  $d_S \leq d_{\max}$ .  $\square$

## 5.2 The case with costs (Proof of [Theorem 2.11](#))

We first restate [Theorem 2.11](#) for convenience.

**Theorem 2.11.** *Consider a WPCS instance  $(\ell, \mathcal{S}, d, y)$  and let  $c \in \mathbb{Q}^\ell$ . If  $\mathcal{S}$  is non-interleaving, there exists an integral selection  $z$  with*

$$\begin{aligned} D_{y,z} &\leq 2d_{\max} \quad \text{and} \\ c^T z &\leq c^T y, \end{aligned}$$

where  $d_{\max} := \max_{S \in \mathcal{S}} d_S$ . Moreover, it can be computed in time  $O(\ell \cdot \langle \text{input} \rangle)$ , where  $\langle \text{input} \rangle$  refers to the length of the bit encoding of the input.

As in the proof of [Theorem 2.10](#), we start by bounding the discrepancy of the prefix intervals  $[i]$  with  $i \in [\ell]$ .

**Theorem 5.4.** *Consider a WPCS instance  $(\ell, \mathcal{S}, d, y)$  and let  $c \in \mathbb{Q}^\ell$ . If  $\mathcal{S}$  is non-interleaving, there exists an integral selection  $z$  with*

$$\begin{aligned} D_{y,z}([i]) &\leq d_{\max} \quad \text{for all } i \in [\ell] \text{ and} \\ c^T z &\leq c^T y, \end{aligned}$$

where  $d_{\max} := \max_{S \in \mathcal{S}} d_S$ . Moreover, such a selection  $z$  can be computed in time  $O(\langle \text{input} \rangle \cdot \ell)$ .

Before proving [Theorem 5.4](#), we first observe that it implies [Theorem 2.11](#). For any interval  $I = \{i, i+1, \dots, j\} \subseteq [\ell]$ , we have

$$D_{y,z}(I) = D_{y,z}([j] \setminus [i-1]) \leq D_{y,z}([i-1]) + D_{y,z}([j]) \leq 2d_{\max}.$$

If  $I$  is a circular interval, then either  $I$  is an interval or the complement  $[\ell] \setminus I$  of  $I$  is an interval. In the latter case we have  $D_{y,z}(I) = D_{y,z}([\ell] \setminus I) \leq 2d_{\max}$ .

To prove [Theorem 5.4](#), we first consider a special case of it where the given fractional selection  $y$  is half-integral. In this special case we will be able to achieve a stronger discrepancy bound. Then we will show that the general case can be reduced to the special case at the cost of increasing the interval discrepancy by a factor of two. The approach we employ to reduce to this special case is a common argument in discrepancy theory (see, e.g., [\[LSV86\]](#), and [\[BRS22\]](#) for a recent application in a scheduling context).

**Lemma 5.5.** *Consider a WPCS instance  $(\ell, \mathcal{S}, d, y)$  and let  $c \in \mathbb{Q}^\ell$ . Suppose  $|S| = 2$  for all  $S \in \mathcal{S}$  and  $y_i = \frac{1}{2}$  for all  $i \in [\ell]$ . If  $\mathcal{S}$  is non-interleaving, there exists an integral selection  $z$  with*

$$\begin{aligned} D_{y,z}([i]) &\leq \frac{d_{\max}}{2} \quad \text{for all } i \in [\ell] \text{ and} \\ c^T z &\leq c^T y, \end{aligned}$$

where  $d_{\max} := \max_{S \in \mathcal{S}} d_S$ . Moreover, such a selection  $z$  can be computed in time  $O(\ell)$ .

*Proof.* We apply [Theorem 5.1](#) to obtain an integral selection  $z \in \{0, 1\}^\ell$  with  $D_{y,z}([i]) \leq \frac{d_{\max}}{2}$  for all  $i \in [\ell]$ . We show that either  $z$  or the vector  $\bar{z} \in \{0, 1\}^\ell$ , defined by

$$\bar{z}_i := \begin{cases} 1 & \text{if } z_i = 0 \\ 0 & \text{if } z_i = 1, \end{cases}$$

has the desired properties. First observe that for any integral selection  $q \in \{0, 1\}^\ell$  and any  $j \in S \in \mathcal{S}$ , we have

$$y_j^d - q_j^d = \begin{cases} \frac{d_S}{2} & \text{if } q_j = 0 \\ -\frac{d_S}{2} & \text{if } q_j = 1, \end{cases}$$

where we used  $y_i = \frac{1}{2}$  for all  $i \in [\ell]$ . This implies  $y_j^d - \bar{z}_j^d = -(y_j^d - z_j^d)$  for all  $j \in [\ell]$ . Thus, we have for all  $i \in [\ell]$ :

$$D_{y,\bar{z}}([i]) = \left| \sum_{j \in [i]} (y_j^d - \bar{z}_j^d) \right| = \left| - \sum_{j \in [i]} (y_j^d - z_j^d) \right| = D_{y,z}([i]) \leq \frac{d_{\max}}{2}.$$

Moreover, we have  $z + \bar{z} = 2y$  and thus  $\min\{c^T z, c^T \bar{z}\} \leq c^T y$ . We conclude that one of the integral selections  $z$  or  $\bar{z}$  has the desired properties. The runtime is  $O(\ell)$ , since [Theorem 5.1](#) gives an integral selection  $z$  in that time and computing  $\bar{z}$  and evaluating the cost is both also linear in  $\ell$ .  $\square$

For  $\alpha \in \mathbb{Q}$ , we call a vector  $y \in \mathbb{Q}^\ell$  an  $\alpha$ -integral vector if every entry of  $y$  is an integer multiple of  $\alpha$ . As a consequence of [Lemma 5.5](#) we obtain the following.

**Lemma 5.6.** *Consider a WPCS instance  $(\ell, \mathcal{S}, d, y)$  with  $\mathcal{S}$  being non-interleaving, let  $c \in \mathbb{Q}^\ell$ , and let  $k \in \mathbb{Z}_{>0}$ . Given a  $2^{-k}$ -integral fractional selection  $y$ , we can compute in polynomial time a  $2^{-k+1}$ -integral fractional selection  $\bar{y}$  such that*

$$\begin{aligned} |y^d([i]) - \bar{y}^d([i])| &\leq 2^{-k} \cdot d_{\max} \quad \text{for all } i \in [\ell] \text{ and} \\ c^T \bar{y} &\leq c^T y, \end{aligned}$$

where  $d_{\max} := \max_{S \in \mathcal{S}} d_S$ . Moreover, such a selection  $\bar{y}$  can be computed in time  $O(\ell)$ .

*Proof.* We construct a WPCS instance  $(\hat{\ell}, \hat{\mathcal{S}}, \hat{d}, \hat{y})$  to which we will then apply [Lemma 5.5](#). Let

$$F := \{i \in [\ell] : y_i \text{ is not } 2^{-k+1}\text{-integral}\},$$

and we denote by  $\hat{\ell} := |F|$  the number of elements in  $F$ . Moreover, we define a renumbering  $f: F \rightarrow [\hat{\ell}]$  such that  $f$  is the (unique) bijection preserving the order of elements, i.e.,  $f(i) < f(j)$  if and only if  $i < j$ . Because  $y$  is  $2^{-k}$ -integral and  $y(S) = 1$  for  $S \in \mathcal{S}$ , we have that  $|F \cap S|$  is even for every set  $S \in \mathcal{S}$ . Hence, there is a partition  $\mathcal{R}_S$  of  $F \cap S$  into sets of size two, which we choose such that no two sets in  $\mathcal{R}_S$  are interleaving. This can be achieved, e.g., by considering the elements of  $F \cap S$  in increasing order and repeatedly putting two consecutive elements in the same set in  $\mathcal{R}_S$ . We define

$$\hat{\mathcal{S}} := \{\{f(i), f(j)\} : \{i, j\} \in \mathcal{R}_S, S \in \mathcal{S}\}.$$

Then  $\hat{\mathcal{S}}$  is a partition of  $[\hat{\ell}]$ . Moreover,  $\hat{\mathcal{S}}$  is non-interleaving, because any two distinct sets from the same set  $\mathcal{R}_S$  are non-interleaving by construction, and for  $R_1 \in \mathcal{R}_{S_1}$  and  $R_2 \in \mathcal{R}_{S_2}$  for distinct

sets  $S_1, S_2 \in \mathcal{S}$ , the sets  $R_1 \subseteq S_1$  and  $R_2 \subseteq S_2$  are non-interleaving because  $S_1$  and  $S_2$  are non-interleaving. For a set  $\hat{S} = \{f(i), f(j)\}$  with  $\{i, j\} \subseteq S \in \mathcal{S}$ , we define  $\hat{d}_{\hat{S}} := d_S$ . Moreover, we set  $\hat{y}_i := \frac{1}{2}$  for all  $i \in [\hat{\ell}]$  and  $\hat{c}_{f(i)} := c_i$  for all  $i \in [\hat{\ell}]$ .

We apply [Lemma 5.5](#) to obtain an integral selection  $\hat{z}$  for the WPCS instance  $(\hat{\ell}, \hat{\mathcal{S}}, \hat{d}, \hat{y})$  with cost  $\hat{c}$ . Hence,  $D_{\hat{y}, \hat{z}}([\hat{i}]) \leq \frac{d_{\max}}{2}$  for  $\hat{i} \in [\hat{\ell}]$ , and  $\hat{c}^T \hat{z} \leq \hat{c}^T \hat{y}$ . Then we define

$$\bar{y}_i := \begin{cases} y_i & \text{if } i \in [\ell] \setminus F \\ y_i + 2^{-k+1} \cdot (\hat{z}_{f(i)} - \hat{y}_{f(i)}) & \text{if } i \in F. \end{cases}$$

Because  $\hat{z}$  is integral, the definition of  $F$  implies that  $\bar{y}$  is  $2^{-k+1}$ -integral. Moreover,  $\hat{c}^T \hat{z} \leq \hat{c}^T \hat{y}$  implies  $\hat{c}^T \bar{y} \leq \hat{c}^T y$ . Finally, we consider an index  $i \in [\ell]$ . If  $F \cap [i] = \emptyset$ , we have  $\bar{y}^d([i]) = y^d([i])$ . Otherwise, let  $i_F := \max\{f(j) : j \in F \cap [i]\}$ . Then

$$\left| y^d([i]) - \bar{y}^d([i]) \right| = 2^{-k+1} \left| \hat{y}^d([i_F]) - \hat{z}^d([i_F]) \right| = 2^{-k+1} \cdot D_{\hat{y}, \hat{z}}([i_F]) \leq 2^{-k} d_{\max}.$$

By [Lemma 5.5](#), we obtain  $\hat{z}$  in  $O(\hat{\ell})$ , where  $\hat{\ell} \leq \ell$  by definition. The subsequent definition of  $\bar{y}$  can clearly be done in  $O(\ell)$ .  $\square$

By applying [Lemma 5.6](#) repeatedly, we can round a  $2^{-k}$ -integral fractional selection to an integral one. However, we might be given a fractional selection  $y$  that is not  $2^{-k}$ -integral for any  $k \in \mathbb{Z}_{>0}$ . To handle this case, we will use the following simple observation.

**Lemma 5.7.** *Consider a WPCS instance  $(\ell, \mathcal{S}, d, y)$  with  $\mathcal{S}$  being non-interleaving, and let  $c \in \mathbb{Q}^\ell$ . Let  $\varepsilon > 0$  and  $k_\varepsilon := \lceil \log_2(\ell/\varepsilon) \rceil$ . Then, there exists a  $2^{-k_\varepsilon}$ -integral fractional selection  $\tilde{y}$  such that*

$$\begin{aligned} \left| y^d(U) - \tilde{y}^d(U) \right| &\leq \varepsilon \cdot d_{\max} \quad \text{for all } U \subseteq [\ell], \text{ and} \\ c^T \tilde{y} &\leq c^T y, \end{aligned}$$

Moreover, such a selection  $\tilde{y}$  can be computed in time  $O(\ell)$ .

*Proof.* For each  $S \in \mathcal{S}$ , let  $i_S \in S$  be such that its cost is minimum, i.e., such that  $c_{i_S} \leq c_j$  for all  $j \in S$ . Then, for all  $j \in S \setminus \{i_S\}$ , let  $\tilde{y}_j$  be the number obtained from rounding down  $y_j$  to the next integer multiple of  $2^{-k_\varepsilon}$ . We define  $\tilde{y}_{i_S} := 1 - \sum_{j \in S \setminus \{i_S\}} \tilde{y}_j$ . Then  $\tilde{y}$  is a fractional selection and, by the choice of the element  $i_S \in S$ , we have  $c^T \tilde{y} \leq c^T y$ . Moreover, for every set  $U \subseteq [\ell]$ , we have  $|y(U) - \tilde{y}(U)| \leq \ell \cdot 2^{-k_\varepsilon} \leq \varepsilon$  and thus  $|y^d(U) - \tilde{y}^d(U)| \leq \varepsilon \cdot d_{\max}$ .  $\square$

We now complete the proof of [Theorem 5.4](#). By applying once [Lemma 5.7](#) for some  $\varepsilon > 0$  and applying [Lemma 5.6](#) repeatedly  $k_\varepsilon = \lceil \log_2(\ell/\varepsilon) \rceil$  times, we obtain an integral selection  $z$  such that, for all  $i \in [\ell]$ ,

$$\left| y^d([i]) - z^d([i]) \right| \leq \varepsilon \cdot d_{\max} + \sum_{i=1}^{k_\varepsilon} 2^{-i} \cdot d_{\max} < (1 + \varepsilon) \cdot d_{\max}. \quad (5)$$

If we choose  $\varepsilon > 0$  small enough, i.e., such that

$$\varepsilon \cdot d_{\max} < \left| y^d([i]) - \tilde{z}^d([i]) \right| - d_{\max} \quad (6)$$

for all integral selections  $\tilde{z}$  and all  $i \in [\ell]$  for which  $\left| y^d([i]) - \tilde{z}^d([i]) \right| > d_{\max}$ , then (5) implies  $\left| y^d([i]) - z^d([i]) \right| \leq d_{\max}$ . Because the right-hand side of (6) can attain only finitely many values, such an  $\varepsilon > 0$  does indeed exist.

To obtain a polynomial time algorithm, we exploit rationality of the entries of the vectors  $y$  and  $d$ , and write  $d_i = \frac{K_i}{L_i}$  and  $y_i = \frac{M_i}{N_i}$  for integers  $K_i, L_i, M_i, N_i$  for all  $i \in [\ell]$ . Then we set

$$\varepsilon := \frac{1}{d_{\max}} \cdot \prod_{i \in [\ell]} \frac{1}{L_i \cdot N_i}.$$

Note that for this choice of  $\varepsilon$ , the number  $k_\varepsilon = \lceil \log_2(\ell/\varepsilon) \rceil$  of applications of [Lemma 5.6](#) is linearly bounded in the input size, which implies an overall runtime of  $O((\text{input}) \cdot \ell)$ . Moreover, (5) implies that, for every  $i \in [\ell]$ , we have

$$\left| y^d([i]) - z^d([i]) \right| - d_{\max} < \varepsilon \cdot d_{\max} = \prod_{i \in [\ell]} \frac{1}{L_i \cdot N_i}.$$

Because both  $\left| y^d([i]) - z^d([i]) \right|$  and  $d_{\max}$  are integer multiples of  $\prod_{i \in [\ell]} \frac{1}{L_i \cdot N_i}$ , this implies that we have  $\left| y^d([i]) - z^d([i]) \right| \leq d_{\max}$ , concluding the proof of [Theorem 5.4](#).

## 6 Concluding remarks

Recall that we required all numbers in our SSUF instance to be rational, so that we have finite input length, which puts us in the traditional computational model to talk about efficient algorithms. We remark that even when allowing arbitrary real numbers (for the given flow  $x$ , the demands  $d$ , and the costs  $c$ ), the existence of an unsplittable flow as in [Theorem 1.7](#) and [Theorem 1.8](#) follows from our results. This can be derived from [Theorem 1.7](#) and [Theorem 1.8](#) by a standard continuity argument.

Moreover, we highlight that the capacity violation  $d_{\max}$  of [Theorem 1.7](#), which shows [Conjecture 1.5](#) in the special case of planar graphs, is tight. This is the case even if there is only a single demand, as observed already in earlier work [[DGG99](#); [MS22](#)]. To see this, consider an instance with a single terminal  $t$  with demand  $k$  and a flow  $x$  that can be decomposed into  $k$  arc-disjoint paths with flow value 1 each. Then  $x(a) \leq 1$  for all  $a \in A$ , but any unsplittable flow  $\mathcal{P}$  will have value  $k = d_{\max}$  on some  $s$ - $t$  path. [[MS22](#)] gives a similar example showing that the lower bound on the unsplittable flow in [Conjecture 1.5](#) is tight, even in planar graphs.

In order to prove [Conjecture 1.4](#) for planar instances, one would need to strengthen [Theorem 1.8](#) by improving the bound on the deviation  $|x(a) - \text{flow}_{\mathcal{P}}(a)|$  of the unsplittable flow  $\mathcal{P}$  and the given flow  $x$  from  $2d_{\max}$  to  $d_{\max}$ . We expand on two natural strategies toward this and provide examples showing why these strategies do not work.

One natural strategy would be to strengthen the interval-discrepancy bound of [Theorem 2.11](#) from  $2d_{\max}$  to  $d_{\max}$ . In the proof of [Theorem 2.11](#) we showed that we can achieve a discrepancy bound of  $d_{\max}$  for prefix intervals  $[i]$  with  $i \in [\ell]$  (see [Theorem 5.4](#)). One might hope to strengthen this to a bound of  $\frac{d_{\max}}{2}$ , as we did in the setting without costs (see [Theorem 5.1](#)). However, in the setting with costs such a strengthening is impossible, as the following example shows. Let  $\varepsilon \in (0, 1]$  and consider the WPCS instance with  $\ell = 2$ ,  $\mathcal{S} = \{\{1, 2\}\}$ , demand  $d_{\{1, 2\}} = 1$ , and fractional selection  $y$  with  $y_1 = \varepsilon$  and  $y_2 = 1 - \varepsilon$ . If the cost vector  $c$  fulfills  $c_1 = 0$  and  $c_2 = 1$ , then the only integral selection  $z$  with  $c^T z \leq c^T y = 1 - \varepsilon$  is the selection  $z$  with  $z_1 = 1$  and  $z_2 = 0$ . Then  $D_{y,z}([1]) = 1 - \varepsilon$ . Because  $d_{\max} = 1$ , and  $\varepsilon$  can be chosen arbitrarily close to zero, this shows tightness of [Theorem 5.4](#).

Another possible strategy to prove [Conjecture 1.4](#) for planar SSUF by strengthening [Theorem 2.11](#) would be to use the reduction to the half-integral setting described in [Section 5.2](#), which

comes at the cost of loosing a factor of two in the discrepancy bound. One might then hope to achieve a  $(y, z)$ -interval-discrepancy bound of  $\frac{d_{\max}}{2}$  for the half-integral special case. However, this is in general impossible as the following example shows. Let  $\ell := 6$  and  $\mathcal{S} := \{\{1, 6\}, \{2, 5\}, \{3, 4\}\}$ . Note that  $\mathcal{S}$  is non-interleaving. Moreover, we define  $d_{\{1,6\}} := 2$ ,  $d_{\{2,5\}} := 1$ ,  $d_{\{3,4\}} := 2$  and  $y_i := \frac{1}{2}$  for all  $i \in [\ell]$ . Then

$$y_1^d = 1 \quad y_2^d = 0.5 \quad y_3^d = 1 \quad y_4^d = 1 \quad y_5^d = 0.5 \quad y_6^d = 1$$

and  $d_{\max} = 2$ . Now consider an integral selection  $z$ . By symmetry, we may assume  $z_1 = 1$  and  $z_6 = 0$ . If  $z_3 = 1$  and  $z_4 = 0$ , we have

$$z_1^d = 2 \quad z_2^d \geq 0 \quad z_3^d = 2 \quad z_4^d = 0 \quad z_5^d \leq 1 \quad z_6^d = 0$$

and thus  $D_{y,z}(\{1, 2, 3\}) \geq 1.5 > \frac{d_{\max}}{2}$ . Otherwise, we have  $z_3 = 0$  and  $z_4 = 1$ . Then we either have

$$z_1^d = 2 \quad z_2^d = 1 \quad z_3^d = 0 \quad z_4^d = 2 \quad z_5^d = 0 \quad z_6^d = 0$$

or

$$z_1^d = 2 \quad z_2^d = 0 \quad z_3^d = 0 \quad z_4^d = 2 \quad z_5^d = 1 \quad z_6^d = 0.$$

In the first case,  $D_{y,z}(\{1, 2\}) = 1.5 > \frac{d_{\max}}{2}$  and in the second case  $D_{y,z}(\{4, 5\}) = 1.5 > \frac{d_{\max}}{2}$ . We conclude that in general it is indeed impossible to achieve a  $(y, z)$ -interval discrepancy of at most  $\frac{d_{\max}}{2}$  on half-integral instances.

## A Hardness of SSUF with lower bounds on cyclic graphs

We start by observing that in case of general (cyclic) graphs, [Conjecture 1.4](#) and even [Conjecture 1.5](#) cannot hold as stated, because the desired unsplittable flow might not exist.

**Observation A.1.** *Given some constant  $\lambda > 0$ , there is an SSUF instance  $(G = (V, A), s, T, d, x)$  such that there exists no unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with*


$$x(a) - \lambda d_{\max} \leq \text{flow}_{\mathcal{P}}(a) \quad \text{for all } a \in A.$$

Consider the graph depicted in [Figure 11](#) with one commodity of demand  $d$ . The fractional flow  $x$  on all black, thin arcs is  $\frac{d}{2}$ , while it is  $d(\lambda + 1)$  on all blue thick arcs. Note that this flow is a convex combination of two walks. Still, in this case there is no path, and actually not even a walk, that can send positive flow along both triangles, which means that, for any unsplittable flow, some lower bounds are violated heavily.

Thus, when considering general graphs, the nature of the problem substantially changes. Instead of proving existence of certain unsplittable flows under the assumption that a splittable flow exists, we consider a different variant of the problem where we are given lower bounds  $(\ell(a))_{a \in A}$  instead.

In acyclic graphs we can use [Theorem 1.7](#) to either decide that there exists no unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with  $\ell(a) \leq \text{flow}_{\mathcal{P}}(a)$  for all  $a \in A$ , or find an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with  $\ell(a) - d_{\max} \leq \text{flow}_{\mathcal{P}}(a)$ . To this end we simply check if there exists a splittable flow satisfying the lower bounds  $(\ell(a))_{a \in A}$ . If no such splittable flow exists, there is also no unsplittable flow satisfying the lower bounds. Otherwise, we obtain a splittable flow  $x$  satisfying the given lower bounds, and we can apply [Theorem 1.7](#) to  $x$  to find an unsplittable flow that violates the lower bounds by at most  $d_{\max}$ .

While [Theorem 1.7](#) cannot be extended to cyclic graphs, as we observed above, one might still hope for an efficient algorithm that either decides that no unsplittable flow satisfying the lower



**Figure 11:** The figure shows an example, where no unsplittable flow can satisfy any  $O(d_{\max})$  violation of the lower bounds. Vertex  $s$  is the source and  $t$  is the terminal. The fractional flow splits half-half between the up and down option and then cycles often along the blue, thick triangles. The unsplittable flow can only choose one of the two options and thus violates the lower bounds heavily.

bounds exists, or finds one that violates the lower bounds by at most  $O(d_{\max})$ . We next show that this is impossible for general (not necessarily planar) instances unless  $P = NP$ .

To this end, we reduce from the 2-Vertex-Disjoint Paths problem (2-VDP), which is well-known to be NP-complete [FHW80]. In a 2-VDP instance, we are given a directed graph  $G = (V, A)$  and two pairs of vertices  $(p_1, q_1)$  and  $(p_2, q_2)$ . The task is to decide whether there are two vertex-disjoint paths  $Q_1$  and  $Q_2$  from  $p_1$  to  $q_1$  and from  $p_2$  to  $q_2$ , respectively. Given a 2-VDP instance  $I = (G, (p_1, q_1), (p_2, q_2))$ , together with a desired number  $k \in \mathbb{Z}_{\geq 1}$  of terminals, we create an auxiliary graph  $\bar{G} = (\bar{V}, \bar{A})$  together with a source  $s \in V$ ,  $k$  terminals  $T \subseteq \bar{V}$ , and demands  $d$  as follows:

- $\bar{G}$  is obtained from  $G$  by adding an arc  $(q_1, p_2)$ , and by adding  $k$  additional vertices  $T = \{t_1, \dots, t_k\}$ , which will be the terminals, together with arcs  $(q_2, t_i)$  for all  $i \in [k]$ ;
- all demands are unit, i.e.,  $d_t = 1$  for all  $t \in T$ ;
- $s = p_1$  is the source.

See Figure 12 for an illustration of this construction. We denote by  $\bar{I} = (\bar{G}, s, T, d)$  the resulting single-source unsplittable flow instance without fractional flow vector  $x$ , as our statements here do not depend on a specific fractional flow or even its existence.

We then have the following immediate relation between the instance  $I$  and  $\bar{I}$ , which exhibits a gap in terms of how much flow can be sent over the arc  $(q_1, p_2)$  in  $\bar{I}$ , depending on whether  $I$  is feasible.

**Proposition A.2.** *Let  $I = (G, (p_1, q_1), (p_2, q_2))$  be a 2-VDP instance, and  $\bar{I} = (\bar{G}, s, T, d)$  be the corresponding SSUF instance (without fractional flow  $x$ ). Then we have the following implications:*


- $I$  is a yes instance  $\implies \exists$  unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  for  $\bar{I}$  with  $\text{flow}_{\mathcal{P}}((q_1, p_2)) = k$ .
- $I$  is a no instance  $\implies \nexists$  unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  for  $\bar{I}$  with  $\text{flow}_{\mathcal{P}}((q_1, p_2)) > 0$ .

*Proof.* If  $I$  is a yes instance, there is an  $s$ - $q_2$  path  $P$  in  $\bar{G}$  using the arc  $(q_1, p_2)$ . By setting each  $P^t$  for  $t \in T$  to be the path obtained by appending the arc  $(q_2, t)$  to  $P$ , we get the desired result.

We prove the contrapositive of the second statement. Hence, if there is an unsplittable flow  $\mathcal{P}$  with  $\text{flow}_{\mathcal{P}}((q_1, p_2)) > 0$ , then at least one path  $P \in \mathcal{P}$  must use the arc  $(q_1, p_2)$ . The path  $P$  contains a  $p_1$ - $q_1$  path and a  $p_2$ - $q_2$  path that are vertex-disjoint, which implies that  $I$  is a yes instance.  $\square$

There are different ways to interpret Proposition A.2. One natural way is that, in (cyclic) single-source unsplittable flow problems, it is hard to distinguish between the case where there is





**Figure 12:** Given a 2-VDP instance  $I = (G, (p_1, q_1), (p_2, q_2))$ , we obtain the corresponding graph  $\bar{G}$  from  $G$  by adding an arc  $(q_1, p_2)$ , adding terminals  $t_1, \dots, t_k$ , and adding for each terminal  $t_i$  an arc  $(q_2, t_i)$ . Moreover, the vertex  $p_1$  is the source. The idea is that there are two vertex-disjoint paths (as depicted in gray) if and only if there is a path from  $s$  to a terminal using the arc  $(q_1, p_2)$ .

an unsplittable flow fulfilling all lower bounds and the case where there is no unsplittable flow even if we allow for a heavy violation of the lower bounds.

**Corollary A.3.** *Let  $k \in \mathbb{Z}_{\geq 1}$  and let  $\lambda \in \mathbb{R}_{\geq 0}$  such that  $\lambda < k$  be constants. Then it is NP-hard to distinguish between the following two cases for SSUF instances on general (cyclic) graphs with  $k$  terminals:*

- (i) *There is an unsplittable flow satisfying all lower bounds, i.e., an unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with  $\ell(a) \leq \text{flow}_{\mathcal{P}}(a)$  for all  $a \in A$ .*
- (ii) *There is no unsplittable flow satisfying all lower bounds with a violation of up to  $\lambda d_{\max}$ , i.e., there is no unsplittable flow  $\mathcal{P} = \{P^t\}_{t \in T}$  with  $\ell(a) - \lambda d_{\max} \leq \text{flow}_{\mathcal{P}}(a)$  for all  $a \in A$ .*

*Proof.* The hardness follows by considering SSUF instances  $\bar{I} = (\bar{G}, s, T, d)$  (without fractional flow vector  $x$ ) stemming from 2-VDP instances  $I = (G, (p_1, q_1), (p_2, q_2))$ . Moreover, we assign a lower bound of  $\ell((q_1, p_2)) = k$  to the arc  $(q_1, p_2)$  and lower bounds of 0 to all other arcs. Then the two cases (i) and (ii) of Corollary A.3 correspond to (a) and (b) of Proposition A.2, respectively. Hence, they allow for distinguishing yes from no instances of 2-VDP, which is NP-hard.  $\square$

Note that the reduction and statement only hold true for general (cyclic) graphs because 2-VDP is polynomial-time solvable in acyclic graphs.

## References

- [AMO93] R. K. Ahuja, T. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Upper Saddle River, NJ, USA: Prentice Hall, 1993. ISBN: 0-13-617549-X.
- [BK15] N. Bansal and J. Kulkarni. “Minimizing flow-time on unrelated machines”. In: *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*. 2015, pp. 851–860. DOI: [10.1145/2746539.2746601](https://doi.org/10.1145/2746539.2746601).
- [BKS05] G. Baier, E. Köhler, and M. Skutella. “The  $k$ -Splittable Flow Problem”. In: *Algorithmica* 42 (2005), pp. 231–248. DOI: [10.1007/s00453-005-1167-9](https://doi.org/10.1007/s00453-005-1167-9).

- [BRS22] N. Bansal, L. Rohwedder, and O. Svensson. “Flow time scheduling and prefix becfiala”. In: *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*. 2022, pp. 331–342. DOI: [10.1145/3519935.3520077](https://doi.org/10.1145/3519935.3520077).
- [CNAO85] N. Chiba, T. Nishizeki, S. Abe, and T. Ozawa. “A Linear Algorithm for Embedding Planar Graphs Using PQ-Trees”. In: *Journal of Computer and System Sciences* 30 (1985), pp. 54–76. DOI: [10.1016/0022-0000\(85\)90004-2](https://doi.org/10.1016/0022-0000(85)90004-2).
- [DGG99] Y. Dinitz, N. Garg, and M. X. Goemans. “On The Single-Source Unsplittable Flow Problem”. In: *Combinatorica* 19.1 (1999), pp. 17–41. DOI: [10.1007/s004930050043](https://doi.org/10.1007/s004930050043).
- [DK05] J. Du and S. G. Kolliopoulos. “Implementing Approximation Algorithms for the Single-Source Unsplittable Flow Problem”. In: *Journal on Experimental Algorithmics* 10.2 (2005). DOI: [10.1145/1064546.1180614](https://doi.org/10.1145/1064546.1180614).
- [FHW80] S. Fortune, J. Hopcroft, and J. Wyllie. “The directed subgraph homeomorphism problem”. In: *Theoretical Computer Science* 10.2 (1980), pp. 111–121. DOI: [10.1016/0304-3975\(80\)90009-2](https://doi.org/10.1016/0304-3975(80)90009-2).
- [GMW22] F. Grandoni, T. Mömke, and A. Wiese. “A PTAS for unsplittable flow on a path”. In: *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*. 2022, pp. 289–302. DOI: [10.1145/3519935.3519959](https://doi.org/10.1145/3519935.3519959).
- [HT74] J. Hopcroft and R. Tarjan. “Efficient Planarity Testing”. In: *Journal of the ACM* 21.4 (1974), pp. 549–568. DOI: [10.1145/321850.321852](https://doi.org/10.1145/321850.321852).
- [Kle96a] J. M. Kleinberg. “Approximation Algorithms for Disjoint Paths Problems”. PhD thesis. Massachusetts Institute of Technology, 1996.
- [Kle96b] J. M. Kleinberg. “Single-Source Unsplittable Flow”. In: *Proceedings of the 37th Annual Symposium on Foundations of Computer Science (FOCS)*. 1996, pp. 68–77. DOI: [10.1109/SFCS.1996.548465](https://doi.org/10.1109/SFCS.1996.548465).
- [Kol05] S. G. Kolliopoulos. “Minimum-Cost Single-Source 2-Splittable Flow”. In: *Information Processing Letters* 94 (2005), pp. 15–18. DOI: [10.1016/j.ipl.2004.12.009](https://doi.org/10.1016/j.ipl.2004.12.009).
- [Kol07] S. G. Kolliopoulos. “Edge-Disjoint Paths and Unsplittable Flow”. In: *Handbook of Approximation Algorithms and Metaheuristics*. 2007. Chap. V, 57-1 to 57–16. ISBN: 1-58488-550-5.
- [KS02] S. G. Kolliopoulos and C. Stein. “Approximation Algorithms for Single-Source Unsplittable Flow”. In: *SIAM Journal on Computing* 31.3 (2002), pp. 919–946. DOI: [10.1137/S0097539799355314](https://doi.org/10.1137/S0097539799355314).
- [KSS08] R. Koch, M. Skutella, and I. Spenke. “Maximum  $k$ -splittable  $s$ - $t$ -flows”. In: *Theory of Computing Systems* 43 (2008), pp. 56–66. DOI: [10.1007/s00224-007-9068-8](https://doi.org/10.1007/s00224-007-9068-8).
- [LST90] J. Lenstra, D. B. Shmoys, and É. Tardos. “Approximation algorithms for scheduling unrelated parallel machines”. In: *Mathematical programming* 46 (1990), pp. 259–271. DOI: [10.1007/BF01585745](https://doi.org/10.1007/BF01585745).
- [LSV86] L. Lovász, J. Spencer, and K. Vesztergombi. “Discrepancy of Set-Systems and Matrices”. In: *European Journal of Combinatorics* 7 (1986), pp. 151–160. DOI: [10.1016/S0195-6698\(86\)80041-5](https://doi.org/10.1016/S0195-6698(86)80041-5).
- [MS22] S. Morell and M. Skutella. “Single source unsplittable flows with arc-wise lower and upper bounds”. In: *Mathematical Programming* 192 (2022), pp. 477–496. DOI: [10.1007/s10107-021-01704-4](https://doi.org/10.1007/s10107-021-01704-4).

- [MSS07] M. Martens, F. Salazar, and M. Skutella. “Convex Combinations of Single Source Unsplittable Flows”. In: *Proceedings of the 15th Annual European Symposium on Algorithms (ESA)*. Vol. 4698/2007. Lecture Notes in Computer Science. 2007, pp. 395–406. DOI: [10.1007/978-3-540-75520-3\\_36](https://doi.org/10.1007/978-3-540-75520-3_36).
- [Sku02] M. Skutella. “Approximating the single source unsplittable min-cost flow problem”. In: *Mathematical Programming* 91.3 (2002), pp. 493–514. DOI: [10.1007/s101070100260](https://doi.org/10.1007/s101070100260).
- [SS09] F. Salazar and M. Skutella. “Single-Source  $k$ -Splittable Min-Cost Flows”. In: *Operations Research Letters* 37 (2009). DOI: [10.1016/j.orl.2008.12.004](https://doi.org/10.1016/j.orl.2008.12.004).
- [Wil19] D. P. Williamson. *Network Flow Algorithms*. Cambridge University Press, 2019. ISBN: 9781316888568. DOI: [10.1017/9781316888568](https://doi.org/10.1017/9781316888568).