# Fair Correlation Clustering in Forests 

Katrin Casel $\square$ (<br>Humboldt University of Berlin, Germany<br>Tobias Friedrich $\square$ (<br>Hasso Plattner Institute, University of Potsdam, Germany<br>Martin Schirneck $\square$ (0)<br>Faculty of Computer Science, University of Vienna, Austria<br>Simon Wietheger<br>Hasso Plattner Institute, University of Potsdam, Germany


#### Abstract

The study of algorithmic fairness received growing attention recently. This stems from the awareness that bias in the input data for machine learning systems may result in discriminatory outputs. For clustering tasks, one of the most central notions of fairness is the formalization by Chierichetti, Kumar, Lattanzi, and Vassilvitskii [NeurIPS 2017]. A clustering is said to be fair, if each cluster has the same distribution of manifestations of a sensitive attribute as the whole input set. This is motivated by various applications where the objects to be clustered have sensitive attributes that should not be over- or underrepresented. Most research on this version of fair clustering has focused on centriod-based objectives.

In contrast, we discuss the applicability of this fairness notion to Correlation Clustering. The existing literature on the resulting Fair Correlation Clustering problem either presents approximation algorithms with poor approximation guarantees or severely limits the possible distributions of the sensitive attribute (often only two manifestations with a $1: 1$ ratio are considered). Our goal is to understand if there is hope for better results in between these two extremes. To this end, we consider restricted graph classes which allow us to characterize the distributions of sensitive attributes for which this form of fairness is tractable from a complexity point of view.

While existing work on Fair Correlation Clustering gives approximation algorithms, we focus on exact solutions and investigate whether there are efficiently solvable instances. The unfair version of Correlation Clustering is trivial on forests, but adding fairness creates a surprisingly rich picture of complexities. We give an overview of the distributions and types of forests where Fair Correlation Clustering turns from tractable to intractable.

As the most surprising insight, we consider the fact that the cause of the hardness of Fair Correlation Clustering is not the strictness of the fairness condition. We lift most of our results to also hold for the relaxed version of the fairness condition. Instead, the source of hardness seems to be the distribution of the sensitive attribute. On the positive side, we identify some reasonable distributions that are indeed tractable. While this tractability is only shown for forests, it may open an avenue to design reasonable approximations for larger graph classes.


2012 ACM Subject Classification Theory of computation $\rightarrow$ Graph algorithms analysis; Social and professional topics $\rightarrow$ Computing / technology policy; Theory of computation $\rightarrow$ Dynamic programming

Keywords and phrases correlation clustering, disparate impact, fair clustering, relaxed fairness
Digital Object Identifier 10.4230/LIPIcs.FORC.2023.
Related Version https://arxiv.org/abs/2302.11295

## 1 Introduction

In the last decade, the notion of fairness in machine learning has increasingly attracted interest, see for example the review by Pessach and Schmueli [26]. Feldman, Friedler, Moeller,

© Katrin Casel, Tobias Friedrich, Martin Schirneck, and Simon Wietheger;
licensed under Creative Commons License CC-BY 4.0
4th Symposium on Foundations of Responsible Computing (FORC 2023).
Editor: Kunal Talwar; Article No. ; pp. :1-:12
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Scheidegger, and Venkatasubramanian [21] formalize fairness based on a US Supreme Court decision on disparate impact from 1971. It requires that sensitive attributes like gender or skin color should neither be explicitly considered in decision processes like hiring but also should the manifestations of sensitive attributes be proportionally distributed in all outcomes of the decision process. Feldman et al. formalize this notion for classification tasks. Chierichetti, Kumar, Lattanzi, and Vassilvitskii [15] adapt this concept for clustering tasks.

In this paper we employ the same disparate impact based understanding of fairness. Formally, the objects to be clustered have a color assigned to them that represents some sensitive attribute. Then, a clustering of these colored objects is called fair if for each cluster and each color the ratio of objects of that color in the cluster corresponds to the total ratio of vertices of that color. More precisely, a clustering is fair, if it partitions the set of objects into fair subsets.

- Definition 1 (Fair Subset). Let $U$ be a finite set of objects colored by a function $c: U \rightarrow[k]$ for some $k \in \mathbb{N}_{>0}$. Let $U_{i}=\{u \in U \mid c(u)=i\}$ be the set of objects of color $i$ for all $i \in[k]$. Then, a set $S \subseteq U$ is fair if and only if for all colors $i \in[k]$ we have $\frac{\left|S \cap U_{i}\right|}{|S|}=\frac{\left|U_{i}\right|}{|U|}$.

To understand how this notion of fairness affects clustering decisions, consider the following example. Imagine that an airport security wants to find clusters among the travelers to assign to each group a level of potential risk with corresponding anticipating measures. There are attributes like skin color that should not influence the assignment to a risk level. A bias in the data, however, may lead to some colors being over- or underrepresented in some clusters. Simply removing the skin color attribute from the data may not suffice as it may correlate with other attributes. Such problems are especially likely if one of the skin colors is far less represented in the data than others. A fair clustering finds the optimum clustering such that for each risk level the distribution of skin colors is fair, by requiring the distribution of each cluster to roughly match the distribution of skin colors among all travelers.

The seminal fair clustering paper by Chierichetti et al. [15] introduced this notion of fairness for clustering and studied it for the objectives $k$-center and $k$-median. Their work was extended by Bera, Chakrabarty, Flores, and Negahbani [9], who relax the fairness constraint in the sense of requiring upper and lower bounds on the representation of a color in each cluster. More precisely, they define the following generalization of fair sets.

- Definition 2 (Relaxed Fair Set). For a finite set $U$ and coloring $c: U \rightarrow[k]$ for some $k \in \mathbb{N}_{>0}$ let $p_{i}, q_{i} \in \mathbb{Q}$ with $0<p_{i} \leqslant \frac{\left|U_{i}\right|}{|U|} \leqslant q_{i}<1$ for all $i \in[k]$, where $U_{i}=\{u \in U \mid c(u)=i\}$. A set $S \subseteq U$ is relaxed fair with respect to $q_{i}$ and $p_{i}$ if and only if $p_{i} \leqslant \frac{\left|S \cap U_{i}\right|}{|S|} \leqslant q_{i}$ for all $i \in[k]$.

Following these results, this notion of (relaxed) fairness was extensively studied for centroidbased clustering objectives with many positive results.

For example, Bercea et al. [10] give bicreteira constant-factor approximations for facility location type problems like $k$-center and $k$-median. Bandyapadhyay, Fomin and Simonov [6] use the technique of fair coresets introduced by Schmidt, Schwiegelshohn, and Sohler [28] to give constant factor approximations for many centroid-based clustering objectives; among many other results, they give a polynomial-time approximation scheme (PTAS) for fair $k$-means and $k$-median in Euclidean space. Fairness for centroid-based objectives seems to be so well understood, that most research already considers more generalized settings, like streaming [28], or imperfect knowledge of group membership [20].

In comparison, there are few (positive) results for this fairness notion applied to graph clustering objectives. The most studied with respect to fairness among those is Correlation Clustering, arguably the most studied graph clustering objective. For Correlation

Clustering we are given a pairwise similarity measure for a set of objects and the aim is to find a clustering that minimizes the number of similar objects placed in separate clusters and the number of dissimilar objects placed in the same cluster. Formally, the input to Correlation Clustering is a graph $G=(V, E)$, and the goal is to find a partition $\mathcal{P}$ of $V$ that minimizes the Correlation Clustering cost defined as

$$
\begin{equation*}
\operatorname{cost}(G, \mathcal{P})=\left|\left\{\left.\{u, v\} \in\binom{V}{2} \backslash E \right\rvert\, \mathcal{P}[u]=\mathcal{P}[v]\right\}\right|+|\{\{u, v\} \in E \mid \mathcal{P}[u] \neq \mathcal{P}[v]\}| \tag{1}
\end{equation*}
$$

Fair Correlation Clustering then is the task to find a partition into fair sets that minimizes the Correlation Clustering cost. We emphasize that this is the complete, unweighted, min-disagree form of Correlation Clustering. (It is often called complete because every pair of objects is either similar or dissimilar but none is indifferent regarding the clustering. It is unweighted as the (dis)similarity between two vertices is binary. A pair of similar objects that are placed in separate clusters as well as a pair of dissimilar objects in the same cluster is called a disagreement, hence the naming of the min-disagree form.)

There are two papers that appear to have started studying Fair Correlation Clustering independently ${ }^{1}$. Ahmadian, Epasto, Kumar, and Mahdian [2] analyze settings where the fairness constraint is given by some $\alpha$ and require that the ratio of each color in each cluster is at most $\alpha$. For $\alpha=\frac{1}{2}$, which corresponds to our fairness definition if there are two colors in a ratio of $1: 1$, they obtain a 256 -approximation. For $\alpha=\frac{1}{k}$, where $k$ is the number of colors in the graph, they give a $16.48 k^{2}$-approximation. We note that all their variants are only equivalent to our fairness notion if there are $\alpha^{-1}$ colors that all occur equally often. Ahmadi, Galhotra, Saha, and Schwartz [1] give an $O\left(c^{2}\right)$-approximation algorithm for instances with two colors in a ratio of $1: c$. In the special case of a color ratio of $1: 1$, they obtain a $3 \beta+4$-approximation, given any $\beta$-approximation to unfair Correlation Clustering. With a more general color distribution, their approach also worsens drastically. For instances with $k$ colors in a ratio of $1: c_{2}: c_{3}: \ldots: c_{k}$ for positive integers $c_{i}$, they give an $\mathrm{O}\left(k^{2} \cdot \max _{2 \leqslant i \leqslant k} c_{i}\right)$-approximation for the strict, and an $\mathrm{O}\left(k^{2} \cdot \max _{2 \leqslant i \leqslant k} q_{i}\right)$-approximation for the relaxed setting ${ }^{2}$.

Following these two papers, Friggstad and Mousavi [23] provide an approximation to the $1: 1$ color ratio case with a factor of 6.18 . To the best of our knowledge, the most recent publication on Fair Correlation Clustering is by Ahmadian and Negahbani [3] who give approximations for Fair Correlation Clustering with a slightly different way of relaxing fairness. They give an approximation with ratio $\mathcal{O}\left(\varepsilon^{-1} k \max _{2 \leqslant i \leqslant k} c_{i}\right)$ for color distribution $1: c_{2}: c_{3}: \ldots: c_{k}$, where $\varepsilon$ relates to the amount of relaxation (roughly $q_{i}=(1+\epsilon) c_{i}$ for our definition of relaxed fairness).

All these results for Fair Correlation Clustering seem to converge towards considering the very restricted setting of two colors in a ratio of $1: 1$ in order to give some decent approximation ratio. In this paper, we want to understand if this is unavoidable, or if there is hope to find better results for other (possibly more realistic) color distributions. In order to isolate the role of fairness, we consider "easy" instances for Correlation Clustering, and study the increase in complexity when adding fairness constraints. Correlation Clustering without the fairness constraint is easily solved on forests. We find that Fair Correlation Clustering restricted to forests turns NP-hard very quickly, even when additionally assuming constant degree or diameter. Most surprisingly, this hardness essentially

[^0]also holds for relaxed fairness, showing that the hardness of the problem is not due to the strictness of the fairness definition.

On the positive side, we identify color distributions that allow for efficient algorithms. Not surprisingly, this includes ratio $1: 1$, and extends to a constant number of $k$ colors with distribution $c_{1}: c_{2}: c_{3}: \ldots: c_{k}$ for constants $c_{1}, \ldots, c_{k}$. Such distributions can be used to model sensitive attributes with a limited number of manifestation that are almost evenly distributed. Less expected, we also find tractability for, in a sense, the other extreme. We show that Fair Correlation Clustering on forests can be solved in polynomial time for two colors with ratio $1: c$ with $c$ being very large (linear in the number of overall vertices). Such a distribution can be used to model a scenario where a minority is drastically underrepresented and thus in dire need of fairness constraints. Although our results only hold for forests, we believe that they can offer a starting point for more general graph classes. We especially hope that our work sparks interest in the so far neglected distribution of ratio $1: c$ with $c$ being very large.

### 1.1 Related Work

The study of clustering objectives similar or identical to Correlation Clustering dates back to the 1960s [8, 27, 31]. Bansal, Blum, and Chawla [7] were the first to coin the term Correlation Clustering as a clustering objective. We note that it is also studied under the name Cluster Editing. The most general formulation of Correlation Clustering regarding weights considers two positive real values for each pair of vertices, the first to be added to the cost if the objects are placed in the same cluster and the second to be added if the objects are placed in separate clusters [4]. The recent book by Bonchi, García-Soriano, and Gullo [11] gives a broad overview of the current research on Correlation Clustering.

We focus on the particular variant that considers a complete graph with $\{-1,1\}$ edgeweights, and the min disagreement objective function. This version is APX-hard [13], implying in particular that there is no algorithm giving an arbitrarily good approximation unless $P=N P$. The best known approximation for Correlation Clustering is the very recent breakthrough by Cohen-Addad, Lee and Newman [16] who give a ratio of $(1.994+\epsilon)$.

We show that in forests, all clusters of an optimal Correlation Clustering solution have a fixed size. In such a case, Correlation Clustering is related to $k$-Balanced Partitioning. There, the task is to partition the graph into $k$ clusters of equal size while minimizing the number of edges that are cut by the partition. Feldmann and Foschini [22] study this problem on trees and their results have interesting parallels with ours.

Aside from the results on Fair Correlation Clustering already discussed above, we are only aware of three papers that consider a fairness notion close to the one of Chierichetti et al. [15] for a graph clustering objective. Schwartz and Zats [29] consider incomplete Fair Correlation Clustering with the max-agree objective function. Dinitz, Srinivasan, Tsepenekas, and Vullikanti [18] study Fair Disaster Containment, a graph cut problem involving fairness. Their problem is not directly a fair clustering problem since they only require one part of their partition (the saved part) to be fair. Ziko, Yuan, Granger, and Ayed [32] give a heuristic approach for fair clustering in general that however does not allow for theoretical guarantees on the quality of the solution.

## 2 Contribution

We now outline our findings on Fair Correlation Clustering. We start by giving several structural results that underpin our further investigations. Afterwards, we present


Figure 1 Example forest where a cluster of size 4 and two clusters of size 2 incur the same cost. With one cluster of size 4 (left), the inter-cluster cost is 0 and the intra-cluster cost is 4 . With two clusters of size 2 (right), both the inter-cluster and intra-cluster cost are 2 .
our algorithms and hardness results for certain graph classes and color ratios. We further show that the hardness of fair clustering does not stem from the requirement of the clusters exactly reproducing the color distribution of the whole graph. This section is concluded by a discussion of possible directions for further research.

### 2.1 Structural Insights

We outline here the structural insights that form the foundation of all our results. We first give a close connection between the cost of a clustering, the number of edges "cut" by a clustering, and the total number of edges in the graph. We refer to this number of "cut" edges as the inter-cluster cost as opposed to the number of non-edges inside clusters, which we call the intra-cluster cost. Formally, the intra- and inter-cluster cost are the first and second summand of the Correlation Clustering cost in Equation (1), respectively. The following lemma shows that minimizing the inter-cluster cost suffices to minimize the total cost if all clusters are of the same size. This significantly simplifies the algorithm development for Correlation Clustering.

- Lemma 3. Let $\mathcal{P}$ be a partition of the vertices of an m-edge graph $G$. Let $\chi$ denote the inter-cluster cost incurred by $\mathcal{P}$ on $G$. If all sets in the partition are of size $d$, then $\operatorname{cost}(\mathcal{P})=\frac{(d-1)}{2} n-m+2 \chi$. In particular, if $G$ is a tree, $\operatorname{cost}(\mathcal{P})=\frac{(d-3)}{2} n+2 \chi+1$.

The condition that all clusters need to be of the same size seems rather restrictive at first. However, we prove in the following that in bipartite graphs and, in particular, in forests and trees there is always a minimum-cost fair clustering such that indeed all clusters are equally large. This property stems from how the fairness constraint acts on the distribution of colors and is therefore specific to Fair Correlation Clustering. It allows us to fully utilize Lemma 3 both for building reductions in NP-hardness proofs as well as for algorithmic approaches as we can restrict our attention to partitions with equal cluster sizes.

Consider two colors of ratio $1: 2$, then any fair cluster must contain at least 1 vertex of the first color and 2 vertices of the second color to fulfil the fairness requirement. We show that a minimum-cost clustering of a forest, due to the small number of edges, consists entirely of such minimal clusters. Every clustering with larger clusters incurs a higher cost.

- Lemma 4. Let $F$ be a forest with $k \geqslant 2$ colors in a ratio of $c_{1}: c_{2}: \ldots: c_{k}$ with $c_{i} \in \mathbb{N}_{>0}$ for all $i \in[k], \operatorname{gcd}\left(c_{1}, c_{2}, \ldots, c_{k}\right)=1$, and $\sum_{i=1}^{k} c_{i} \geqslant 3$. Then, all clusters of every minimum-cost fair clustering are of size $d=\sum_{i=1}^{k} c_{i}$.

Lemma 4 does not extend to two colors in a ratio of $1: 1$ as illustrated in Figure 1. This color distribution is the only case for forests where a partition with larger clusters can have the same (but no smaller) cost. We prove a slightly weaker statement than Lemma 4, namely, that there is always a minimum-cost fair clustering with minimal clusters. This property, in

Table 1 Running times of our algorithms for Fair Correlation Clustering on forests depending on the color ratio. Value $p$ is any rational such that $n / p-1$ is integral; $c_{1}, c_{2}, \ldots, c_{k}$ are coprime positive integers, possibly depending on $n$. Functions $f$ and $g$ are given in the full version.

| Color Ratio | $1: 1$ | $1: 2$ | $1:(n / p-1)$ | $c_{1}: c_{2}: \ldots: c_{k}$ |
| :--- | :---: | :---: | :---: | :---: |
| Running Time | $\mathrm{O}(n)$ | $\mathrm{O}\left(n^{6}\right)$ | $\mathrm{O}\left(n^{f(p)}\right)$ | $\mathrm{O}\left(n^{g\left(c_{1}, \ldots, c_{k}\right)}\right)$ |

turn, holds not only for forests but for every bipartite graph. Note that in general bipartite graphs there are more color ratios than only $1: 1$ that allow for these ambiguities.

- Lemma 5. Let $G=(A \cup B, E)$ be a bipartite graph with $k \geqslant 2$ colors in a ratio of $c_{1}: c_{2}: \ldots: c_{k}$ with $c_{i} \in \mathbb{N}_{>0}$ for all $i \in[k]$ and $\operatorname{gcd}\left(c_{1}, c_{2}, \ldots, c_{k}\right)=1$. Then, there is a minimum-cost fair clustering such that all its clusters are of size $d=\sum_{i=1}^{k} c_{i}$. Further, each minimum-cost fair clustering with larger clusters can be transformed into a minimum-cost fair clustering such that all clusters contain no more than d vertices in linear time.

In summary, the results above show that the ratio of the color classes is the key parameter determining the cluster size. If the input is a bipartite graph whose vertices are colored with $k$ colors in a ratio of $c_{1}: c_{2}: \cdots: c_{k}$, our results imply that without loosing optimality, solutions can be restricted to contain only clusters of size $d=\sum_{i=1}^{k} c_{i}$, each with exactly $c_{i}$ vertices of color $i$. Starting from these observations, we show in this work that the color ratio is also the key parameter determining the complexity of Fair Correlation Clustering. On the one hand, the simple structure of optimal solutions restricts the search space and enables polynomial-time algorithms, at least for some instances. Additionally, due to the fixed cluster size $d$, returning any fair clustering in a forest can only cause so many mistakes. In fact, this procedure yields an approximation factor decreasing in $d$ and converging to 1 as $d \rightarrow \infty$. Combining this with the fact that Fair Correlation Clustering can be solved in time increasing in $d$, see Table 1, allows for a PTAS in forests. On the other hand, these insights allow us to show hardness already for very restricted input classes. The technical part of most of the proofs consists of exploiting the connection between the clustering cost, total number of edges, and the number of edges cut by a clustering.

### 2.2 Tractable Instances

We start by discussing the algorithmic results. The simplest case is that of two colors, each one occurring equally often. We prove that for bipartite graphs with a color ratio $1: 1$ Fair Correlation Clustering is equivalent to the maximum bipartite matching problem, namely, between the vertices of different color. Via the standard reduction to computing maximum flows, this allows us to benefit from the recent breakthrough by Chen, Kyng, Liu, Peng, Probst Gutenberg, and Sachdeva [14]. It gives an algorithm running in time $m^{1+o(1)}$.

The remaining results focus on forests as the input, see Table 1. It should not come as a surprise that our main algorithmic paradigm is dynamic programming. A textbook version finds a maximum matching in linear time in a forests, solving the $1: 1$ case. For general color ratios, we devise much more intricate dynamic programs. We use the color ratio $1: 2$ as an introductory example. The algorithm has two phases. In the first, we compute a list of candidate splittings that partition the forest into connected parts containing at most 1 blue and 2 red vertices each. In the second phase, we assemble the parts of each of the splittings to fair clusters and return the cheapest resulting clustering. The difficulty lies in

Table 2 Complexity of Fair Correlation Clustering on trees and general graphs depending on the diameter. The value $c$ is a positive integer, possibly depending on $n$.

| Diameter | Color Ratio | Trees | General Graphs |
| :---: | :---: | :---: | :---: |
| 2,3 | any | $\mathrm{O}(n)$ | NP-hard |
| $\geqslant 4$ | $1: c$ | NP-hard | NP-hard |

the two phases not being independent from each other. It is not enough to minimize the "cut" edges in the two phases separately. We prove that the costs incurred by the merging additionally depends on the number of parts of a certain type generated in the splittings. Tracking this along with the number of cuts results in a $\mathrm{O}\left(n^{6}\right)$-time algorithm. Note that we did not optimize the running time as long as it is polynomial.

We generalize this to $k$ colors in a ratio $c_{1}: c_{2}: \cdots: c_{k} .{ }^{3}$ We now have to consider all possible colorings of a partition of the vertices such that in each part the $i$-th color occurs at most $c_{i}$ times. While assembling the parts, we have to take care that the merged colorings remain compatible. The resulting running time is $\mathrm{O}\left(n^{g\left(c_{1}, \ldots, c_{k}\right)}\right)$ for some (explicit) polynomial $g$. Recall that, by Lemma 4, the minimum cluster size is $d=\sum_{i=1}^{k} c_{i}$. If this is a constant, then the dynamic program runs in polynomial time. If, however, the number of colors $k$ or some color's proportion grows with $n$, it becomes intractable. Equivalently, the running time gets worse if there are very large but sublinearly many clusters.

To mitigate this effect, we give a complementary algorithm at least for forests with two colors. Namely, consider the color ratio $1: \frac{n}{p}-1$. Then, an optimal solution has $p$ clusters each of size $d=n / p$. The key observation is that the forest contains $p$ vertices of the color with fewer occurrences, say, blue, and any fair clustering isolates the blue vertices. This can be done by cutting at most $p-1$ edges and results in a collection of (sub-)trees where each one has at most one blue vertex. To obtain the clustering, we split the trees with red excess vertices and distribute those among the remaining parts. We track the costs of all the $\mathrm{O}\left(n^{\text {poly }(p)}\right)$ many cut-sets and rearrangements to compute the one of minimum cost. In total, the algorithm runs in time $\mathrm{O}\left(n^{f(p)}\right)$ for some polynomial in $p$. In summary, we find that if the number of clusters $p$ is constant, then the running time is polynomial. Considering in particular an integral color ratio $1: c,{ }^{4}$, we find tractability for forests if $c=\mathrm{O}(1)$ or $c=\Omega(n)$. We will show next that Fair Correlation Clustering with this kind of a color ratio is NP-hard already on trees, hence the hardness must emerge somewhere for intermediate $c$.

### 2.3 A Dichotomy for Bounded Diameter

Table 2 shows the complexity of Fair Correlation Clustering on graphs with bounded diameter. We obtain a dichotomy for trees with two colors with ratio $1: c$. If the diameter is at most 3, an optimal clustering is computable in $\mathrm{O}(n)$ time, but for diameter at least 4, the problem becomes NP-hard. In fact, the linear-time algorithm extends to trees with an arbitrary number of colors in any ratio.

The main result in that direction is the hardness of Fair Correlation Clustering already on trees with diameter at least 4 and two colors of ratio $1: c$. This is proven by a reduction from the strongly NP-hard 3-Partition problem. There, we are given positive

[^1]

Figure 2 The tree with diameter 4 in the reduction from 3-Partition to Fair Correlation Clustering.
integers $a_{1}, \ldots, a_{\ell}$ where $\ell$ is a multiple of 3 and there exists some $B$ with $\sum_{i=1}^{\ell} a_{i}=B \cdot \frac{\ell}{3}$. The task is to partition the numbers $a_{i}$ into triples such that each one of those sums to $B$. The problem remains NP-hard if all the $a_{i}$ are strictly between $B / 4$ and $B / 2$, ensuring that, if some subset of the numbers sums to $B$, it contains exactly three elements.

We model this problem as an instance of Fair Correlation Clustering as illustrated in Figure 2. We build $\ell$ stars, where the $i$-th one consists of $a_{i}$ red vertices, and a single star of $\ell / 3$ blue vertices. The centers of the blue star and all the red stars are connected. The color ratio in the resulting instance is $1: B$. Lemma 4 then implies that there is a minimum-costs clustering with $\ell / 3$ clusters, each with a single blue vertex and $B$ red ones. We then apply Lemma 3 to show that this cost is below a certain threshold if and only if each cluster consist of exactly three red stars (and an arbitrary blue vertex), solving 3-Partition.

### 2.4 Maximum Degree

The reduction above results in a tree with a low diameter but arbitrarily high maximum degree. We have to adapt our reductions to show hardness also for bounded degrees. The results are summarized in Table 3. If the Fair Correlation Clustering instance is not required to be connected, we can represent 3-Partition with a forest of trees with maximum degree 2 , that is, a forest of paths. The input numbers are modeled by paths with $a_{i}$ vertices. The forest also contains $\ell / 3$ isolated blue vertices, which again implies that an optimal fair clustering must have $\ell / 3$ clusters each with $B$ red vertices. By defining a sufficiently small cost threshold, we ensure that the fair clustering has cost below it if and only if none of the path-edges are "cut" by the clustering, corresponding to a partition of the $a_{i}$.

There is nothing special about paths, we can arbitrarily restrict the shape of the trees, as long it is possible to form such a tree with any given number of vertices. However, the argument crucially relies on the absence of edges between the $a_{i}$-trees and does not transfer to connected graphs. This is due to the close relation between inter-cluster costs and the number of edges, see Lemma 3. The complexity of Fair Correlation Clustering on a single path with a color ratio 1:c remains open. Notwithstanding, we show hardness for trees in two closely related settings: keeping the ratio $1: c$ but raising the maximum degree to 5 , or having a single path with $n / 2$ colors where each color is shared by exactly 2 vertices.

For the case of maximum degree 5 and two colors with ratio $1: c$, we can again build on the 3-Partition machinery. The construction is inspired by how Feldmann and Foschini [22] used the problem to show hardness of computing so-called $k$-balanced partitions. We adapt it to our setting in which the vertices are colored and the clusters need to be fair.

For the single path with $n / 2$ colors, we reduce from (the 1-regular 2-colored variant of) the Paint Shop Problem for Words [19]. There, a word is given in which every symbol

Table 3 Hardness of Fair Correlation Clustering on trees and forests depending on the maximum degree. The value $c$ is a positive integer, possibly depending on $n$. The complexity for paths (trees with maximum degree 2 ) with color ratio $1: c$ is open.

| Max. Degree | Color Ratio | Trees | Forests |
| :---: | :---: | :---: | :---: |
| 2 | $1: c$ |  | NP-hard |
| $\geqslant 2$ | $n / 2$ colors, | NP-hard | NP-hard |
| $\geqslant 5$ | $1: c$ | NP-hard | NP-hard |

appears exactly twice. The task is to assign the values 0 and 1 to the letters ${ }^{5}$ such that, for each symbol, exactly one occurrence receives a 1 , and the number of blocks of consecutive 0 s or 1 s is minimized. In the translation to Fair Correlation Clustering, we represent the word as a path and the symbols as colors. To remain fair, there must be two clusters containing exactly one vertex of each color, translating back to a $0 / 1$-assignment to the word.

### 2.5 Relaxed Fairness

One could think that the hardness of Fair Correlation Clustering already for classes of trees and forests has its origin in the strict fairness condition. After all, the color ratio in each cluster must precisely mirror that of the whole graph. This impression is deceptive. Instead, we lift most of our hardness results to Relaxed Fair Correlation Clustering considering the relaxed fairness of Bera et al. [9]. Recall Definition 2. It prescribes two rationals $p_{i}$ and $q_{i}$ for each color $i$ and allows, the proportion of $i$-colored elements in any cluster to be in the interval $\left[p_{i}, q_{i}\right]$, instead of being precisely $c_{i} / d$, where $d=\sum_{j=1}^{k} c_{j}$.

The main conceptual idea is that, in some settings, the minimum-cost solution under a relaxed fairness constraint is exactly fair. We show this for the cases in which we reduce from 3-Partition. In particular, Relaxed Fair Correlation Clustering with a color ratio of $1: c$ is NP-hard on trees with diameter 4 and forests of paths, respectively. Furthermore, the transferal of hardness is immediate for the case of a single path with $n / 2$ colors and exactly 2 vertices of each color. Any relaxation of fairness still requires one vertex of each color in every cluster, keeping the equivalence to the Paint Shop Problem for Words.

In contrast, algorithmic results are more difficult to extend if there are relaxedly fair solutions that have lower cost than any exactly fair one. We then no longer know the cardinality of the clusters in an optimal solution. As a proof of concept, we show that a slight adaption of our dynamic program for two colors in a ratio of $1: 1$ still works for what we call $\alpha$-relaxed fairness. ${ }^{6}$ There, the lower fairness ratio is $p_{i}=\alpha \cdot \frac{c_{i}}{d}$ and the upper one is $q_{i}=\frac{1}{\alpha} \cdot \frac{c_{i}}{d}$ for some parameter $\alpha \in(0,1)$. We give an upper bound on the necessary cluster size depending on $\alpha$, which is enough to find a good splitting of the forest. Naturally, the running time now also depends on $\alpha$, but is of the form $O\left(n^{h(1 / \alpha)}\right)$ for some polynomial $h$. In particular, we get an polynomial-time algorithm for constant $\alpha$. The proof of correctness consists of an exhaustive case distinction already for the simple case of $1: 1$. We are confident that this can be extended to more general color ratios, but did not attempt it in this work.

[^2]
### 2.6 Summary and Outlook

We show that Fair Correlation Clustering on trees, and thereby forests, is NP-hard. It remains so on trees of constant degree or diameter, and-for certain color distributions-it is also NP-hard on paths. On the other hand, we give a polynomial-time algorithm if the minimum size $d$ of a fair cluster is constant. We also provide an efficient algorithm for the color ratio $1: c$ if the total number of clusters is constant, corresponding to $c \in \Theta(n)$. For our main algorithms and hardness results, we prove that they still hold when the fairness constraint is relaxed, so the hardness is not due to the strict fairness definition. Ultimately, we hope that the insights gained from these proofs as well as our proposed algorithms prove helpful to the future development of algorithms to solve Fair Correlation Clustering on more general graphs. In particular, fairness with color ratio $1: c$ with $c$ being very large seems to be an interesting and potentially tractable type of distribution for future study.

As first steps to generalize our results, we give a polynomial-time approximation scheme (PTAS) for Fair Correlation Clustering on forests. This further motivates to study approximation algorithms on more general classes of graphs. Another avenue for future research could be that Lemma 5, bounding the cluster size of optimal solutions, extends also to bipartite graphs. This may prove helpful in developing exact algorithms for bipartite graphs with other color ratios than $1: 1$. Regarding further graph classes, we suspect that tractability will first have to be examined for the standard (unfair) Correlation Clustering before considering additional fairness constraints.

Parameterized algorithms are yet another approach to solving more general instances. When looking at the decision version of Fair Correlation Clustering, our results can be cast as an XP-algorithm when the problem is parameterized by the cluster size $d$, for it can be solved in time $\mathrm{O}\left(n^{g(d)}\right)$ for some function $g$. Similarly, we get an XP-algorithm for the number of clusters as parameter. We wonder whether Fair Correlation Clustering can be placed in the class FPT of fixed-parameter tractable problems for any interesting structural parameters. This would require a running time of, e.g., $g(d) \cdot \operatorname{poly}(n)$. There are FPT-algorithms for Cluster Editing parameterized by the cost of the solution [12]. Possibly, future research might provide similar results for the fair variant as well. A natural extension of our dynamic programming approach could potentially lead to an algorithm parameterizing by the treewidth of the input graph. Such a solution would be surprising, however, since to the best of our knowledge even for normal, unfair Correlation Clustering ${ }^{7}$ and for the related Max Dense Graph Partition [17] no treewidth approaches are known.

Finally, it is interesting how Fair Correlation Clustering behaves on paths. While we obtain NP-hardness for a particular color distribution from the Paint Shop Problem For Words, the question of whether Fair Correlation Clustering on paths with for example two colors in a ratio of $1: c$ is efficiently solvable or not is left open. However, we believe that this question is rather answered by the study of the related (discrete) Necklace Splitting problem, see the work of Alon and West [5]. There, the desired cardinality of every color class is explicitly given, and it is non-constructively shown that there always exists a split of the necklace with the number of cuts meeting the obvious lower bound. A constructive splitting procedure may yield some insights for Fair Correlation Clustering on paths.

[^3]— References
1 Saba Ahmadi, Sainyam Galhotra, Barna Saha, and Roy Schwartz. Fair correlation clustering. CoRR, arXiv:2002.03508, 2020. ArXiv preprint. URL: https://arxiv.org/abs/2002.03508.
2 Sara Ahmadian, Alessandro Epasto, Ravi Kumar, and Mohammad Mahdian. Fair correlation clustering. In Proceedings of the 23rd Conference on Artificial Intelligence and Statistics (AISTATS), page 4195-4205, 2020. URL: https://proceedings.mlr.press/v108/ahmadian20a. html.
3 Sara Ahmadian and Maryam Negahbani. Improved approximation for fair correlation clustering. CoRR, abs/2206.05050, 2022. doi:10.48550/arXiv.2206. 05050.
4 Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: Ranking and clustering. Journal of the ACM, 55(5):23:1-23:27, 2008. doi:10.1145/1411509. 1411513.

5 Noga Alon and Douglas B. West. The Borsuk-Ulam theorem and bisection of necklaces. Proceedings of the American Mathematical Society, 98(4):623-628, 1986. doi:10.2307/2045739.
6 Sayan Bandyapadhyay, Fedor V. Fomin, and Kirill Simonov. On coresets for fair clustering in metric and euclidean spaces and their applications. In Proceedings of the 48 th International Colloquium on Automata, Languages, and Programming (ICALP), pages 23:1-23:15, 2021. doi:10.4230/LIPIcs.ICALP.2021.23.
7 Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. Machine Learning, 56(1-3):89-113, 2004. doi:10.1023/В:МАСН. 0000033116 .57574.95.
8 Amir Ben-Dor, Ron Shamir, and Zohar Yakhini. Clustering gene expression patterns. Journal of Computational Biology, 6(3-4):281-297, 1999. doi:10.1089/106652799318274.
9 Suman K. Bera, Deeparnab Chakrabarty, Nicolas J. Flores, and Maryam Negahbani. Fair algorithms for clustering. In Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS), page 4954-4965, 2019.
10 Ioana Oriana Bercea, Martin Groß, Samir Khuller, Aounon Kumar, Clemens Rösner, Daniel R. Schmidt, and Melanie Schmidt. On the cost of essentially fair clusterings. In Proceedings of the 2019 Conference on Approximation for Combinatorial Optimization Problems and the 2019 Conference on Randomization in Computation (APPROX/RANDOM), volume 145 of LIPIcs, pages 18:1-18:22, 2019.
11 Francesco Bonchi, David García-Soriano, and Francesco Gullo. Correlation Clustering. Morgan \& Claypool Publishers, 2022. doi:10.2200/S01163ED1V01Y202201DMK019.
12 Sebastian Böcker and Jan Baumbach. Cluster editing. In Proceedings of the 9th Conference on Computability in Europe (CiE), page 33-44, 2013. doi:10.1007/978-3-642-39053-1_5.
13 Moses Charikar, Venkatesan Guruswami, and Anthony Wirth. Clustering with qualitative information. Journal of Computer and System Sciences, 71(3):360-383, 2005. doi:10.1016/j. jcss.2004.10.012.
14 Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum Flow and Minimum-Cost Flow in Almost-Linear Time. In Proceedings of the 63rd Symposium on Foundations of Computer Science (FOCS), pages 612-623, 2022. doi:10.1109/FOCS54457.2022.00064.
15 Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Fair clustering through fairlets. In Proceedings of the 31st Conference on Neural Information Processing Systems (NeurIPS), page 5036-5044, 2017.
16 Vincent Cohen-Addad, Euiwoong Lee, and Alantha Newman. Correlation Clustering with Sherali-Adams. In Proceedings of the 63rd Symposium on Foundations of Computer Science (FOCS), pages 651-661. IEEE, 2022. doi:10.1109/FOCS54457.2022.00068.
17 Julien Darlay, Nadia Brauner, and Julien Moncel. Dense and sparse graph partition. Discrete Applied Mathematics, 160(16):2389-2396, 2012. doi:10.1016/j.dam.2012.06.004.
18 Michael Dinitz, Aravind Srinivasan, Leonidas Tsepenekas, and Anil Vullikanti. Fair disaster containment via graph-cut problems. In Proceedings of the 25th Conference on Artificial

Intelligence and Statistics (AISTATS), page 6321-6333, 2022. URL: https://proceedings. mlr.press/v151/dinitz22a.html.
19 Thomas Epping, Winfried Hochstättler, and Peter Oertel. Complexity results on a paint shop problem. Discrete Applied Mathematics, 136:2-3:217-226, 2004. doi:10.1016/S0166-218X (03) 00442-6.
20 Seyed A. Esmaeili, Brian Brubach, Leonidas Tsepenekas, and John P. Dickerson. Probabilistic fair clustering. In Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS), page 12743-12755, 2020.
21 Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and removing disparate impact. In Proceedings of the 21th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), page 259-268, 2015. doi:10.1145/2783258.2783311.
22 Andreas E. Feldmann and Luca Foschini. Balanced partitions of trees and applications. Algorithmica, 71(2):354-376, 2015. doi:10.1007/s00453-013-9802-3.
23 Zachary Friggstad and Ramin Mousavi. Fair correlation clustering with global and local guarantees. In Proceedings of the 2021 Workshop on Algorithms and Data Structures (WADS), page 414-427, 2021. doi:10.1007/978-3-030-83508-8_30.
24 Jonggyu Jang and Hyun Jong Yang. $\alpha$-Fairness-maximizing user association in energyconstrained small cell networks. IEEE Transactions on Wireless Communications, 21(9):74437459, 2022. doi:10.1109/TWC.2022.3158694.
25 Suchi Kumari and Anurag Singh. Fair end-to-end window-based congestion control in timevarying data communication networks. International Journal of Communication Systems, 32(11), 2019. doi:10.1002/dac. 3986.
26 Dana Pessach and Erez Shmueli. A review on fairness in machine learning. ACM Computing Surveys, 55(3):51:1-51:44, 2022. doi:10.1145/3494672.
27 Simon Régnier. Sur quelques aspects mathématiques des problèmes de classification automatique. Mathématiques et Sciences Humaines, 82:31-44, 1983.
28 Melanie Schmidt, Chris Schwiegelshohn, and Christian Sohler. Fair coresets and streaming algorithms for fair $k$-means. In Proceedings of the $1^{7}$ th Workshop on Approximation and Online Algorithms (WAOA), page 232-251, 2020. doi:10.1007/978-3-030-39479-0_16.
29 Roy Schwartz and Roded Zats. Fair correlation clustering in general graphs. In Proceedings of the 2022 Conference on Approximation for Combinatorial Optimization Problems and the 2022 Conference on Randomization in Computation (APPROX/RANDOM), pages 37:1-37:19, 2022. doi:10.4230/LIPIcs.APPROX/RANDOM. 2022.37.

30 Xiao Xin. An FPT algorithm for the correlation clustering problem. Key Engineering Materials, 474-476:924-927, 2011. doi:10.4028/www.scientific.net/KEM.474-476.924.
31 Charles T. Zahn, Jr. Approximating symmetric relations by equivalence relations. Journal of the Society for Industrial and Applied Mathematics, 12(4):840-847, 1964. doi:10.1137/0112071.
32 Imtiaz Masud Ziko, Jing Yuan, Eric Granger, and Ismail Ben Ayed. Variational fair clustering. In Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI), page 11202-11209, 2021. doi:10.1609/aaai.v35i12.17336.


[^0]:    ${ }^{1}$ Confusingly, they both carry the title Fair Correlation Clustering.
    2 Their theorem states they achieve an $\mathrm{O}\left(\max _{2 \leqslant i \leqslant k} q_{i}\right)$-approximation but when looking at the proof it seems they have accidentally forgotten the $k^{2}$ factor.

[^1]:    ${ }^{3}$ The $c_{i}$ are coprime, but they are not necessarily constants with respect to $n$.
    ${ }^{4}$ In a color ratio $1: c, c$ is not necessarily a constant, but ratios like $2: 5$ are not covered.

[^2]:    ${ }^{5}$ The original formulation [19] assigns colors, aligning better with the paint shop analogy. We change the exposition here in order to avoid confusion with the colors in the fairness sense.
    6 This should not be confused with the notion of $\alpha$-fairness in resource allocation [24, 25].

[^3]:    7 In more detail, no algorithm for complete Correlation Clustering has been proposed. Xin [30] gives a treewidth algorithm for incomplete Correlation Clustering for the treewidth of the graph of all positively and negatively labeled edges.

