Fair Correlation Clustering in General Graphs

Roy Schwartz ✉
The Henry and Marilyn Taub Faculty of Computer Science, Technion, Haifa, Israel

Roded Zats ✉
The Henry and Marilyn Taub Faculty of Computer Science, Technion, Haifa, Israel

Abstract

We consider the family of Correlation Clustering optimization problems under fairness constraints. In Correlation Clustering we are given a graph whose every edge is labeled either with a $+$ or a $-$, and the goal is to find a clustering that agrees the most with the labels: $+$ edges within clusters and $-$ edges across clusters. The notion of fairness implies that there is no over, or under, representation of vertices in the clustering: every vertex has a color and the distribution of colors within each cluster is required to be the same as the distribution of colors in the input graph. Previously, approximation algorithms were known only for fair disagreement minimization in complete unweighted graphs. We prove the following: (1) there is no finite approximation for fair disagreement minimization in general graphs unless $P = NP$ (this hardness holds also for bicriteria algorithms); and (2) fair agreement maximization in general graphs admits a bicriteria approximation of ≈ 0.591 (an improved ≈ 0.609 true approximation is given for the special case of two uniformly distributed colors). Our algorithm is based on proving that the sticky Brownian motion rounding of [Abbasi Zadeh-Bansal-Guruganesh-Nikolov-Schwartz-Singh SODA'20] copes well with uncut edges.

2012 ACM Subject Classification Theory of computation → Facility location and clustering

Keywords and phrases Correlation Clustering, Approximation Algorithms, Semi-Definite Programming

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2022.37

Category APPROX

Funding This work has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement no. 852870-ERC-SUBMODULAR.

1 Introduction

`Correlation-Clustering` is a family of clustering optimization problems in which the goal is to cluster objects given pairwise similarity/dissimilarity information over the objects. In `Correlation-Clustering` we are given a graph $G = (V, E)$ equipped with edge weights $w : E \to \mathbb{R}_+$, whose vertices are the objects, and each edge $e = (u, v) \in E$ is labeled either with a $+$ or a $-$. A $+$ indicates similarity of u and v and a $-$ indicates dissimilarity of u and v. We denote by E^+ the collection of edges labeled with a $+$ and by E^- the collection of edges labeled with a $-$. The goal is to find a clustering $\mathcal{C}$, i.e., $\mathcal{C} = \{C_1, \ldots, C_l\}$ is a partition of V with no restriction on l, that agrees as much as possible with the labeling of the edges. A $+$ edge is in agreement if its endpoints are in the same cluster and a $-$ edge is in agreement if its endpoints are in different clusters. Two natural objectives have attracted much attention since the introduction of `Correlation-Clustering` close to two decades ago by Bansal, Blum and Chawla [12]. The first, denoted as `Max-Agreement`, is to maximize the total weight of edges that are in agreement:

$$\max_{\mathcal{C}} \left\{ \sum_{e=(u,v)\in E^+ : \mathcal{C}(u)=\mathcal{C}(v)} w(e) + \sum_{e=(u,v)\in E^- : \mathcal{C}(u)\neq\mathcal{C}(v)} w(e) \right\},$$

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022).
Editors: Amit Chakrabarti and Chaitanya Swamy; Article No. 37; pp. 37:1–37:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

where $\mathcal{C}(u)$ denotes the cluster in $\mathcal{C}$ that vertex u belongs to. The second, denoted as `Min-Disagreement`, is to minimize the total weight of edges that are in disagreement:

$$\min_{\mathcal{C}} \left\{ \sum_{e=(u,v)\in E^+:\mathcal{C}(u)\neq\mathcal{C}(v)} w(e) + \sum_{e=(u,v)\in E^-:\mathcal{C}(u)=\mathcal{C}(v)} w(e) \right\}.$$

`Correlation-Clustering` has attracted much attention [7, 9, 19, 23, 39, 5, 24, 25, 6, 18, 28], both from the theoretical and practical perspectives. From a theoretical perspective `Correlation-Clustering` captures some fundamental graph cut problems such as `Multicut` and `Multiway-Cut`. From a practical perspective, it has found numerous practical applications in a wide range of settings, e.g., image segmentation [41], cross lingual link detection [40], coreference resolution [35], to name a few (refer to the survey of Wirth [41] for additional details).

Chierichetti, Kumar, Lattanzi and Vassilvitskii [21] introduced the notion of fairness in clustering where they considered the k-`Center` and k-`Median` problems. Informally, in fair clustering problems, each vertex has a type and each cluster needs to contain not too many and not too few vertices from each type. In general, fairness in clustering has received much attention in recent years that goes beyond k-`Center` and k-`Median`, e.g., [2, 11, 14, 33, 38, 15, 1, 4] (refer to surveys [16, 20] for additional details). One of the main reasons for considering fairness in algorithms in general, and clustering in particular, arises from human-centric applications. The goal is to ensure that the solutions are not biased with respect to a sensitive feature such as gender or race. For example, clustering and learning algorithms used for college admissions, bank loans, job applications etc. might be biased [16, 20]. Thus, there is a lot of effort to develop fair clustering algorithms as seen in the literature referenced above.

In this work we consider fairness in `Correlation-Clustering`. Formally, each vertex v is associated with one of k given colors $\{1, \ldots, k\}$ and we denote v's color by $c(v)$. Additionally, we denote by V_i all vertices of color i, i.e., $V_i \triangleq \{u : c(u) = i\}$ for every $i = 1, \ldots, k$. We are also given the ratios of these colors in V, i.e., there exists $h \in \mathbb{N}$ such that V contains $h \cdot p_i$ vertices of the i^{th} color (where $p_1, \ldots, p_k \in \mathbb{N}^{\geq 1}$).[1] We denote these ratios by $p_1 : \ldots : p_k$.[2] The fairness constraint on the clustering $\mathcal{C}$ is that for every cluster in $\mathcal{C}$ and for every two colors i and j the ratio of the number of vertices in the cluster of color i to the number of vertices in the cluster of color j equals p_i/p_j. Hence, every cluster in $\mathcal{C}$ preserves the color ratios of the vertices in the input graph G. We denote the problem of `Min-Disagreement` with a fairness constraint as `Fair-Min-Disagreement` and the problem of `Max-Agreement` with a fairness constraint as `Fair-Max-Agreement`. Typically, all the above mentioned applications of fairness in clustering satisfy that $\sum_{i=1}^{k} p_i = o(n)$.

`Fair-Min-Disagreement` in complete unweighted graphs was considered by Ahmadi, Galhotra, Saha and Schwartz [1] and by Ahmadian, Epasto, Kumar and Mahdian [4]. For two colors and a ratio of $1:1$ [1] present an approximation of $(3\alpha + 4)$ where α is the best known approximation for `Min-Disagreement` in complete unweighted graphs ($\alpha = 2.06$ [19]). For two colors and a ratio of $1:p$ [4, 1] present an approximation of $O(p^2)$. For a general number of colors k and ratios $1 : p_2 : \ldots : p_k$ an approximation of $k^2 \cdot \max_{i=2,\ldots,k}\{p_i^2\}$ was also given by [4, 1], as well as relaxed bi-criteria guarantees. All the above results reduce the problem to `Min-Disagreement` (without any fairness requirements) by matching nodes of different colors and merging them.

[1] It is assumed without loss of generality that there does not exist a number $s > 1$ such that $p_i/s \in \mathbb{N}$ for all $i = 1, \ldots, k$.

[2] Note that $p_i/p_j = |V_i|/|V_j|$ for every $i, j = 1, \ldots, k$.

To the best of our knowledge, no approximation algorithms are known for general instances of `Fair-Min-Disagreement` as the above results of [1, 4] apply only to complete unweighted graphs. Additionally, to the best of our knowledge, no approximation algorithms are known for `Fair-Max-Agreement`.

1.1 Our Results and Techniques

We show that `Fair-Min-Disagreement` is hard to approximate within any finite approximation factor. Moreover, we prove that this hardness holds even for the special case of only two colors and a ratio of $1 : 1$. This is summarized in the following theorem.

▶ **Theorem 1.** *If `Fair-Min-Disagreement` with 2 colors and a ratio of $1 : 1$ admits a polynomial time approximation algorithm with a finite approximation guarantee, then $P = NP$.*

This hardness result is extended to bi-criteria algorithms, and it holds even for the special case of only three colors and ratios of $1 : 1 : 1$. We say that an algorithm for `Fair-Min-Disagreement` is a bi-criteria $(\alpha, 1 + \varepsilon)$-approximation if it outputs a clustering $\mathcal{C} = \{C_1, \dots, C_l\}$ that satisfies: (1) the cost of $\mathcal{C}$ is at most α times the cost of an optimal solution; and (2) for each $1 \le r \le l$ it holds that $|C_r \cap V_i|/|C_r \cap V_j| \le (1 + \varepsilon)p_i/p_j$ for every (ordered) pair of colors i and j. This is summarized in the following theorem.

▶ **Theorem 2.** *For every $\alpha \ge 1$ and $\varepsilon > 0$, if `Fair-Min-Disagreement` with 3 colors and ratios of $1 : 1 : 1$ admits a bi-criteria $(\alpha, 1 + \varepsilon)$ polynomial time approximation algorithm, then $P = NP$.*

Let us focus now on `Fair-Max-Agreement`. Obtaining an approximation of (roughly) $1/2$ is easy (see Section 2), and thus the challenge is improving it. In order to achieve such an improvement, we first notice that one can restrict attention to solutions that contain only two clusters. We prove that if one returns the best of: (1) an α-approximate fair two-cluster solution; and (2) a suitably chosen solution that is comprised of the smallest possible fair clusters (note that each such cluster contains exactly p_i vertices of color i), then we can obtain an approximation better than $1/2$ for `Fair-Max-Agreement` assuming α is sufficiently large. The resulting approximation for `Fair-Max-Agreement` depends on α, thus the bulk of the effort is focused on obtaining a good approximation for the problem where the output is restricted to having only two clusters.

First, we consider a case study with two colors and a ratio of $1 : 1$. For this case study problem, we prove that one can reduce the two-cluster problem to a cut maximization problem that captures both `Max-Bisection` and Max-$\frac{n}{2}$-Uncut (see Section 1.3 for the exact definitions) with *no* fairness constraints. Thus, we use machinery developed for `Max-Bisection` and Max-$\frac{n}{2}$-Uncut, that is based on rounding a Lassere SDP hierarchy relaxation (see [37, 10, 42]), to obtain the following theorem.

▶ **Theorem 3.** *`Fair-Max-Agreement` with two colors and a ratio of $1 : 1$ admits a polynomial time 0.609-approximation algorithm.*

When considering general instances, it is not clear if (or how) one can reduce the two-cluster problem to a problem that has no fairness constraints. Hence, a different approach is needed. We adopt the sticky Brownian motion approach of Abbasi-Zadeh, Bansal, Guruganesh, Nikolov, Schwartz and Singh [43], which was successfully used for approximating `Max-Cut` with side constraints (see Section 2.2 for the definition). In order to apply this approach to our problem, we prove that it can simultaneously handle both edges that cross between the two clusters and edges that do not cross between the two clusters (it is important to note that the work of [43] deals only with edges that cross the cut when considering `Max-Cut` with side constraints). However, this comes at a price of a slightly worse bi-criteria approximation when compared to the case study which has two colors and a ratio of 1 : 1.

We say that an algorithm for `Fair-Max-Agreement` is a bi-criteria (α, ε)-approximation if it outputs a clustering $\mathcal{C} = \{C_1, \dots, C_l\}$ that satisfies: (1) the value of $\mathcal{C}$ is at least α times the value of an optimal clustering; and (2) for every $1 \le j \le l$ there exists a $h_j \in \{1, \dots, n/\sum_{i=1}^{k} p_i\}$ such that $\big||C_j \cap V_i| - h_j \cdot p_i\big| \le \varepsilon n$ for all $1 \le i \le k$.

▶ **Theorem 4.** *`Fair-Max-Agreement` with $k \ge 2$ colors and ratios of $p_1 : \dots : p_k$ admits for every $0 < \varepsilon < 1 - \sum_{i=1}^{k} p_i/n$ a bicriteria $((0.591 - \varepsilon)(1 - \sum_{i=1}^{k} p_i/n), \varepsilon)$-approximation whose running time is $O(n^{poly(\log(k)/\varepsilon)})$.*

Recalling that $\sum_{i=1}^{k} p_i = o(n)$ is the typical case, the approximation in the above theorem is in fact $0.591 - \varepsilon - o(1)$. Moreover, if $k = O(1)$ then the running time of the algorithm is polynomial.

1.2 Related Work

`Correlation-Clustering` has received a lot of attention since its introduction by Bansal, Blum, and Chawla [12] close to two decades ago. The best known approximation algorithm `Min-Disagreement` in general graphs obtains a $O(\log n)$ approximation [24, 18]. For `Max-Agreement` the best known approximations are obtained by rounding the natural SDP relaxation and achieve a guarantee of 0.7666 [39] and 0.7664 [18]. For complete unweighted graphs `Max-Agreement` admits a PTAS [12] while `Min-Disagreement` has a long sequence of works [12, 19, 18, 6] where the current best known achieves an approximation of 2.06 [19].

Fairness in clustering has attracted much attention since the work of Chierichetti, Kumar, Lattanzi and Vassilvitskii [21] for k-`Center` and k-`Median`. It was followed by works on the same two problems [14, 15, 11], as well as k-`Means` [38, 27]. Moreover, [33] considered fairness in the context of spectral clustering. Related notions of fairness were also studied [3, 14]. Fairness in `Correlation-Clustering` was considered by [1, 4] and also extended to hierarchical clustering [2].

In this work we use Lasserre SDP hierarchy to formulate relaxations. The Lasserre hierarchy [34] has been used to develop approximation algorithms for numerous combinatorial optimization problems. Here we mention only few of the related works that directly relate to our problem. Focusing on `Max-Bisection`, Raghavendra and Tan [37] obtained a 0.85 approximation ratio using the Lasserre SDP hierarchy. Following their work, Austrin, Benabbas and Georgiou [10] improved this ratio to 0.8776 which almost matches the Goemans-Williamson approximation ratio for `Max-Cut` [29]. Wu, Du and Xu [42] considered other graph bisection maximization problems and generalized the algorithm of [10] and showed that `Max-`$\frac{n}{2}$`-Uncut` admits an approximation ratio of 0.8776.

1.3 Preliminaries

We denote a cut as $\mathcal{S} = \{S, V \setminus S\}$ where $\delta(S) = \{(u, v) \in E | (u \in S \wedge v \notin S) \vee (u \notin S \wedge v \in S)\}$ is the collection of edges crossing S, and $E(S) = \{(u, v) \in E | u, v \in S\}$ is the collection of edges that have both endpoints in S. For every $X \subseteq E$ we denote $w(X) = \sum_{e \in X} w(e)$ the sum of weights of edges in X. We similarly use the notations above for E^+ and E^-, i.e., $w^-(X) = w(X \cap E^-)$ and $w^+(X) = w(X \cap E^+)$. For the `Max-Agreement` objective and a clustering $\mathcal{C}$ we denote the weight of edges in agreement in $\mathcal{C}$ as $v(\mathcal{C})$ (alternatively, $v(\mathcal{C})$ is the value of the clustering $\mathcal{C}$). Additionally, we denote an optimal clustering by $\mathcal{C}^*$ and its value by $OPT = v(\mathcal{C}^*)$. Moreover, we denote by $v^+(\mathcal{C})$ and $v^-(\mathcal{C})$ the contribution of the $+$ and $-$ edges to $v(\mathcal{C})$, respectively. Let us now define the variant of the problem where the number of clusters is bounded.

▶ **Definition 5.** *`Fair-Max-Agreement` where the numbers of clusters in a solution is required to be at most r is denoted as `Fair-Max-Agreement[r]`.*

Note that `Fair-Max-Agreement[n]` is essentially `Fair-Max-Agreement` with no restriction on the number of clusters in the output. `Fair-Max-Agreement[2]` is related to `Max-Bisection` and `Max-`$\frac{n}{2}$`-Uncut` problems which are defined as follows. Given a graph $G = (V, E)$ the goal is to find a cut $S \subseteq V$ where $|S| = n/2$ such that $w(\delta(S))$ is maximized for `Max-Bisection` and $w(E(S)) + w(E(V \setminus S))$ is maximized for `Max-`$\frac{n}{2}$`-Uncut`.

In this work we use Lasserre SDP hierarchy relaxations, which contain vectors $\mathbf{v_S}$ for subsets $S \subseteq V$. We use the following abbreviated notations: $\mathbf{v_i} = \mathbf{v_{\{i\}}}$ for the singleton set $\{i\}$, $\mathbf{v_0} = \mathbf{v_\emptyset}$ for the empty set, $\mu_i = \mathbf{v_i} \cdot \mathbf{v_0}$ denotes the "marginal probability" of vertex i, and $\rho_{ij} = \mathbf{v_i} \cdot \mathbf{v_j}$ is the covariance between vertices i and j. Additionally, $\widetilde{\mathbf{w}}_\mathbf{i} = \mathbf{v_i} - \mu_i \mathbf{v_0}$ is the component of $\mathbf{v_i}$ in the linear subspace that is orthogonal to $\mathbf{v_0}$, and $\mathbf{w_i} = \widetilde{\mathbf{w}}_\mathbf{i} / \|\widetilde{\mathbf{w}}_\mathbf{i}\|_2$ is its normalized vector.

2 Algorithms for Fair-Max-Agreement

We split this section into two parts. In the first part we consider the case study with two colors and a uniform ratio of $1:1$. In the second part we consider general instances with $k \geq 2$ colors and ratios $p_1 : \ldots : p_k$.

2.1 Case Study – Two Colors with Ratio 1:1

2.1.1 A Simple (1/2)-Approximation

▶ **Observation 6.** *Let $G = (V_1 \cup V_2, E^+ \cup E^-)$ be an instance with two colors and a ratio of $1:1$. Let $f : V_1 \to V_2$ be a bijection such that $M^- \triangleq \{(u, f(u)) : u \in V_1\}$ minimizes $w(M^- \cap E^-)$. Then every clustering $\mathcal{C}$ satisfies $v^-(\mathcal{C}) \leq w(E^-) - w(M^- \cap E^-)$.* [3]

The proof of the following theorem, which is based on Observation 6, appears in Appendix A.

▶ **Theorem 7.** *There is a polynomial time $(1/2)$-approximation algorithm for `Fair-Max-Agreement` in general weighted graphs with two colors and a ratio of $1:1$.*

2.1.2 Beating the (1/2)-Approximation Ratio

The following lemma shows that there is a solution with only two clusters whose value is sufficiently large (a similar idea was used in, .*e.g.*, Charikar and Wirth [17]). Its proof appears in Appendix B.

[3] Intuitively, for every clustering $\mathcal{C}$ we create a matching M between V_1 and V_2 such that every matched pair of nodes appears in the same cluster in $\mathcal{C}$ (note that there can be more than one such matching). Thus, every clustering $\mathcal{C}$ must incur a loss due to the $-$ edges whose value is at least the total weight of $-$ edges in M, i.e., $w(M \cap E^-)$.

▶ **Lemma 8.** *For every clustering $\mathcal{C}$ there is a clustering $\mathcal{S} = \{S, \overline{S}\}$ satisfying: $v(\mathcal{S}) \geq v^+(\mathcal{C}) + \frac{1}{2}v^-(\mathcal{C})$.*

The following lemma reduces `Fair-Max-Agreement` to `Fair-Max-Agreement[2]` with bounded loss in the approximation factor (its proof appears in Appendix C).

▶ **Lemma 9.** *If there is an α-approximation algorithm for `Fair-Max-Agreement[2]` with two colors and a ratio of $1:1$, then there is a $(2\alpha)/(2+\alpha)$-approximation algorithm for `Fair-Max-Agreement` with two colors and a ratio of $1:1$.*

We note that if $\alpha > {}^2\!/_3$ then Lemma 9 implies an approximation better than ${}^1\!/_2$ for `Fair-Max-Agreement`. Therefore, we focus our attention now on presenting an approximation that is strictly better than ${}^2\!/_3$ for `Fair-Max-Agreement[2]` assuming two colors and a ratio of $1:1$. To achieve this goal we define the following optimization problem.

▶ **Definition 10.** *The `Max-Agreement-Bisection` problem is defined as follows. Given an edge weighted graph $G = (V, E)$ equipped with non-negative edge weights $w : E \to \mathbb{R}_+$, where each edge is labeled either $+$ or $-$, the task is to partition the nodes into two clusters of equal size so as to maximize the overall agreement, i.e.,*

$$\max_{S \subseteq V : |S| = n/2} \left\{ w^-(\delta(S)) + w^+(E(S)) + w^+(E(\overline{S})) \right\}.$$

It is important to note that in `Max-Agreement-Bisection` there are no colors, therefore no fairness constraints. Nonetheless, relying on the fact that the number of colors is only two and the ratio is $1:1$, we present an approximation preserving reduction from `Fair-Max-Agreement[2]` to `Max-Agreement-Bisection`. This is summarized in the following lemma.

▶ **Lemma 11.** *`Fair-Max-Agreement[2]` with two colors and a ratio of $1:1$ has an approximation preserving reduction to `Max-Agreement-Bisection`.*

Proof. We are given an instance of `Fair-Max-Agreement[2]` with two colors and a ratio of $1:1$. I.e., a graph $G = (V_1 \cup V_2, E^+ \cup E^-)$ where $|V_1| = |V_2|$. We construct an instance for `Max-Agreement-Bisection` as follows. Consider the graph $\tilde{G} = (V_1 \cup V_2, \tilde{E}^+ \cup \tilde{E}^-)$ where

$$\tilde{E}^+ \triangleq \{(u,v) \in E^+ \mid c(u) = c(v)\} \cup \{(u,v) \in E^- \mid c(u) \neq c(v)\}$$
$$\tilde{E}^- \triangleq \{(u,v) \in E^- \mid c(u) = c(v)\} \cup \{(u,v) \in E^+ \mid c(u) \neq c(v)\}.$$

For every solution $\mathcal{S} = \{S, \overline{S}\}$ for `Max-Agreement-Bisection` we efficiently construct a solution $\mathcal{S}' = \{S', \overline{S'}\}$ for `Fair-Max-Agreement[2]`. Let $\mathcal{S} = \{S, \overline{S}\}$ be a solution to the former problem, we construct a clustering $\mathcal{S}' = \{S', \overline{S'}\}$ as follows: $S' = \{u \in S \mid c(u) = 1\} \cup \{u \in \overline{S} \mid c(u) = 2\}$. One can note that S' is obtained from S by swapping the side of the cut all vertices of color 2 reside in.

Note that every edge $e \in E$ which was in agreement in the solution $\{S, \overline{S}\}$ for `Max-Agreement-Bisection`, has a corresponding edge $\tilde{e} \in \tilde{E}$ which is in agreement in the solution $\{S', \overline{S'}\}$ for `Fair-Max-Agreement[2]`, and vice versa. Thus, $v(\mathcal{S}) = v(\mathcal{S}')$, i.e. the value of the solution remains the same. All that remains to prove is that $\mathcal{S}'$ satisfies the fairness constraints. First, one can note that $|V_1 \cap S| = {}^n\!/_2 - |V_2 \cap S|$ since $|S| = {}^n\!/_2$. Moreover, ${}^n\!/_2 - |V_2 \cap S| = |V_2 \cap \overline{S}|$ since $|V_2| = {}^n\!/_2$ (recall that the ratio is $1:1$ and there are n vertices on total). From the definition of S' we can infer that: $|V_2 \cap S'| = |V_2 \cap \overline{S}|$. This proves $|V_2 \cap S'| = |V_1 \cap S'|$, i.e., S' satisfies the fairness constraints (and therefore $\overline{S'}$ also satisfies the fairness constraints). This concludes the proof. ◀

We emphasize that the above approach of reducing `Fair-Max-Agreement[2]` to a graph bisection problem, heavily relies on the fact that there are only two colors with a ratio of 1 : 1 and it fails for a general instance. Thus, for general instances a different approach is required.

In order to cope with `Max-Agreement-Bisection` we apply the approach of Raghavendra and Tan [37], and the subsequent works of [10, 42], which is based on rounding a Lasserre SDP hierarchy relaxation.

Following Halperin and Zwick [31] and Han, Ye and Zhang [32], we present a general graph bisection problem. This problem is parametrized by four coefficients c_0, c_1, c_2, c_3 and is defined as follows (via a quadratic formulation):

$$\begin{array}{lll} \max & \sum_{e=(i,j)\in E} w(e)(c_0 + c_1 x_0 x_i + c_2 x_0 x_j + c_3 x_i x_j) & \\ \text{s.t.} & \sum_{i\in V} x_0 x_i = 0 & \\ & x_i^2 = 1 & 0 \le i \le n \end{array}$$

Note that in this problem, $x_i \in \{\pm 1\}$ for every $i \in V$, since the last constraint is $x_i^2 = 1$. Therefore, the first constraint, $\sum_{i\in V} x_0 x_i = 0$ is equivalent to the fact that exactly half the variables equal 1 and the other half equal -1.

The coefficients c_0, c_1, c_2, c_3 depend on the exact graph bisection problem which we aim to solve. For example, when considering `Max-Bisection` the coefficients are $c_0 = 1/2, c_1 = 0, c_2 = 0, c_3 = -1/2$. Additionally, when considering `Max-`$\frac{n}{2}$`-Uncut` the coefficients are $c_0 = 1/2, c_1 = 0, c_2 = 0, c_3 = 1/2$.

We note that `Max-Agreement-Bisection` resembles both `Max-Bisection` and `Max-`$\frac{n}{2}$`-Uncut` since: (1) in all three problems we aim to find a cut that contains exactly half of the vertices; and (2) the objective of `Max-Agreement-Bisection` can be seen as the sum of the objectives of `Max-Bisection` and `Max-`$\frac{n}{2}$`-Uncut` on two graphs over the same set of vertices V (the graph (V, E^-) corresponds to the `Max-Bisection` objective whereas (V, E^+) corresponds to the `Max-`$\frac{n}{2}$`-Uncut` objective). Therefore, the objective for `Max-Agreement-Bisection` can be formally written as follows:

$$\sum_{e=(i,j)\in E^+} w(e)\left(\frac{1}{2} + \frac{1}{2}x_i x_j\right) + \sum_{e=(i,j)\in E^-} w(e)\left(\frac{1}{2} - \frac{1}{2}x_i x_j\right).$$

Equivalently, `Max-Agreement-Bisection` can be seen as an extension of the general graph bisection problem in which the $c_0, \ldots, c_3$ coefficients are not uniform over the edges of the graph. Specifically, for + edges the coefficients are $c_0 = 1/2, c_1 = 0, c_2 = 0, c_3 = 1/2$ and for $-$ edges the coefficients are $c_0 = 1/2, c_1 = 0, c_2 = 0, c_3 = -1/2$.

Our main observation is that the algorithm and analysis of [10] for `Max-Bisection`, and the followup work of [42] for the general graph bisection problem described above, can be extended to `Max-Agreement-Bisection` with virtually no change in the analysis. We refer to Appendix D to the high level details as to why `Max-Agreement-Bisection` admits an approximation of 0.8776 (an approximation of 0.8776 is the guarantee [10] proved for `Max-Bisection` and [42] for `Max-`$\frac{n}{2}$`-Uncut`). This enables us to obtain the following Corollary.

▶ **Corollary 12.** *`Fair-Max-Agreement[2]` with two colors and a ratio of* 1 : 1 *is approximable in polynomial time to within a factor 0.8776.*

Proof of Theorem 3. Follows from Corollary 12 and Lemma 9. ◀

2.2 Approximating General Instances

For general instances (either $k > 2$ or non-uniform ratios) the approximation guarantees we provide are slightly worse than for instances with two colors and a ratio of $1:1$. The main use of Lemma 8 is that there is a good solution that has only two clusters. It is important to note that this lemma holds for any number of colors and any ratios, hence it also applies to general instances. However, Lemma 9, which reduces the problem to the two cluster variant, does not hold for a general instance with the exact same guarantee. The reason is that even for the case of uniform ratios $1 : \ldots : 1$ and k colors we are required to find a min cost perfect matching in a k-partite graph (where the cost of a hyperedge is the total weight of $-$ edges between nodes in the hyperedge). Hence, we no longer can find in polynomial time a clustering $\mathcal{C}$ which satisfies the condition $v^-(\mathcal{C}) \geq v^-(\mathcal{C}^*)$. To overcome the above difficulty we use a different approach in which we randomly choose a clustering that is based on a random k-partite matching between the colors. This approach incurs only a small loss in the approximation guarantee. Let us first describe the simple randomized approximation algorithm which obtains an approximation ratio of $1/2 - o(1)$ and then describe how to improve upon this ratio.

2.2.1 Simple $(1/2 - o(1))$-Approximation

Our random clustering algorithm is summarized in Algorithm 1.

Algorithm 1 Random k-Partite Matching.

```
Input: G = (V_1 ∪ ... ∪ V_k, E), {p_i}_{i=1}^k;
𝒞 ← ∅;
while V ≠ ∅ do
    C ← ∅;
    for i ← 1 to k do
        Let S_i be a uniform random set of p_i nodes from V_i;
        C ← C ∪ S_i;
        V_i ← V_i \ S_i;
    end
    𝒞 ← 𝒞 ∪ {C};
end
return 𝒞;
```

▶ **Observation 13.** *Let $G = (V, E)$ be a graph with $k \geq 2$ colors and ratios of $p_1 : p_2 : \ldots : p_k$. Let $\mathcal{C}$ be the output of Algorithm 1. Then $\mathbb{E}[v^-(\mathcal{C})] \geq (1 - (\sum_{i=1}^k p_i)/n) \cdot w(E^-)$.*

The following extends Theorem 7, which provided an approximation of $1/2$ for the case there are two colors and a ratio of $1:1$, to a general number of colors k and ratios $p_1 : \ldots : p_k$ while suffering a small loss of $(\sum_{i=1}^k p_i)/(2n)$ in the approximation guarantee. Note that if $\sum_{i=1}^k p_i = o(n)$ then this loss is at most $o(1)$. This is achieved by replacing the clustering that corresponds to the matching M^- of Observation 6 with the clustering $\mathcal{C}$ generated by Algorithm 1.

▶ **Theorem 14.** *There is a randomized polynomial time $(1/2 - (\sum_{i=1}^k p_i)/(2n))$-approximation algorithm for `Fair-Max-Agreement` in general weighted graphs with $k \geq 2$ colors and ratios of $p_1 : p_2 : \ldots : p_k$.*

Proof. The algorithm chooses the best from the following two solutions: a single cluster containing all the nodes or the output of Algorithm 1. The value of the former clustering is the total weight of + edges, i.e., $w(E^+)$. On the other hand, the expected value of the latter clustering is at least $(1-(\sum_{i=1}^k p_i)/n)\cdot w(E^-)$ (see Observation 13). Let us denote by $\mathcal{C}_{\text{ALG}}$ the chosen clustering, i.e., $\mathcal{C}_{\text{ALG}} = \arg\max\{v(\mathcal{C}), v(\{V\})\}$. One can note that $\mathbb{E}[v(\mathcal{C}_{\text{ALG}})] \geq 1/2\cdot(w(E^+)+(1-(\sum_{i=1}^k p_i)/n)\cdot w(E^-)) \geq (1/2-(\sum_{i=1}^k p_i)/(2n))\cdot w(E) \geq (1/2-(\sum_{i=1}^k p_i)/(2n))\cdot OPT$. ◀

2.2.2 Beating the $(1/2 - o(1))$-Approximation Ratio

The following lemma shows that one can reduce `Fair-Max-Agreement` to `Fair-Max-Agreement[2]`, for general instances, with a small loss in the approximation guarantee (similarly to Lemma 9). If $\sum_{i=1}^k p_i = o(n)$ then the approximation guarantee of the following lemma equals $(2\alpha)/(2+\alpha)-o(1)$.

▶ **Lemma 15.** *If there is an α-approximation algorithm for `Fair-Max-Agreement[2]` with $k \geq 2$ colors and ratios of $p_1 : \ldots : p_k$, then there is a $(1-\frac{\sum_{i=1}^k p_i}{n})(\frac{1}{\alpha}(\frac{2+\alpha}{2}-\frac{\sum_{i=1}^k p_i}{n}))^{-1}$-approximation algorithm for `Fair-Max-Agreement` with $k \geq 2$ colors and ratios of $p_1 : \ldots : p_k$.*

Proof. The proof is similar to the proof of Lemma 9 except that instead of choosing the best of two solutions when one of them has a value of at least $v^-(\mathcal{C}^*)$ for some optimal clustering $\mathcal{C}^*$, we need to settle for a solution with value $(1-(\sum_{i=1}^k p_i)/n)\cdot v^-(\mathcal{C}^*)$ (swapping the clustering Observation 6 guarantees with the clustering Observation 13 guarantees). The rest of the calculations are similar. ◀

All that remains is to find a good approximation for `Fair-Max-Agreement[2]`. When considering an approximation better than $1/2$, we note that the approach taken for two colors and a ratio of $1:1$, which reduces `Fair-Max-Agreement[2]` to `Max-Agreement-Bisection`, does not work for general instances. Instead we take the approach of Abbasi-Zadeh, Bansal, Guruganesh, Nikolov, Schwartz and Singh [43] which presented the problem of `Max-Cut` with side constraints (denoted by `Max-Cut-Sc`) and an algorithm for it. In the `Max-Cut-Sc` problem we are given an n-vertex graph $G=(V,E)$, a collection $\mathcal{F}=\{F_1,\ldots,F_k\}$ of k subsets of V, and cardinality bounds $b_1,...,b_k \in \mathbb{N}$. The goal is to find a subset $S \subseteq V$ that maximizes the total weight of edges that cross S, subject to satisfying $|S \cap F_i| = b_i$ for all $1 \leq i \leq k$. Their algorithm uses a sticky Brownian motion for the rounding process of a suitable semi-definite relaxation for `Max-Cut-Sc`. In order to utilize this approach we: (1) present a generalization of `Max-Cut-Sc` which also handles uncut edges (this problem is denoted by `Max-Agreement-Sc`); and (2) prove that the rounding approach of [43] can handle uncut edges, i.e., + edges, with the same approximation guarantee of the cut edges.[4] Formally, the input for the `Max-Agreement-Sc` problem is the same as `Max-Cut-Sc`, with the addition that every edge is labeled either with a + or a −. The goal is to find a subset $S \subseteq V$ that maximizes $w^-(\delta(S)) + w^+(E(S)) + w^+(E(\overline{S}))$ subject to the same constraints as in `Max-Cut-Sc`.

[4] The algorithm of [43] for `Max-Cut-Sc` in fact takes the best out of two solutions: the Brownian motion rounding and randomized rounding. The latter algorithm is needed for the case that the value of the optimal solution is small. In our case we prove that the instance for which we need to solve `Max-Agreement-Sc` cannot have an optimal solution of small value (see Lemma 17), thus our algorithm just utilizes the Brownian motion approach.

Following the above discussion, we show how one can handle `Fair-Max-Agreement[2]` by solving a sequence of `Max-Agreement-Sc` instances with an appropriate choice of cardinality bounds. Specifically, given an instance of `Fair-Max-Agreement[2]` with k colors, we choose $\mathcal{F} = \{V_1, \ldots, V_k\}$ and all the cardinality bounds $b_i = h \cdot p_i$ for all $i = 1, \ldots, k$. The sequence of `Max-Agreement-Sc` is defined by enumerating over all values of h in the range $h = 1, \ldots, n/\sum_{i=1}^{k} p_i$. Finding a solution to this problem for all possible h values captures the fairness constraints. This is true since an optimal clustering of `Fair-Max-Agreement[2]` corresponds to some specific (unknown) value of h in the above range. Note that an instance as above to the `Max-Agreement-Sc` is always feasible for every value of h in the range between 1 and $n/\sum_{i=1}^{k} p_i$. This is summarized in Algorithm 2.

Algorithm 2 `Fair-Max-Agreement[2]` via `Max-Agreement-Sc`.

Input: $G = (V_1 \cup \ldots \cup V_k, E^+ \cup E^-)$, $\{p_i\}_{i=1}^{k}$;
$\mathcal{F} \leftarrow \{V_1, \ldots, V_k\}$;
for $h \leftarrow 1$ **to** $n/\sum_{i=1}^{k} p_i$ **do**
  $b_i \leftarrow h \cdot p_i \quad \forall 1 \leq i \leq k$;
  Let $\mathcal{C}_h$ be the sticky Brownian motion solution for `Max-Agreement-Sc` with bounds $\{b_i\}_{i=1}^{k}$;
end
return $\text{argmax}\{v(\mathcal{C}_h) : h = 1, \ldots, n/\sum_{i=1}^{k} p_i\}$;

The bulk of the effort is in solving `Max-Agreement-Sc` via the Brownian motion approach of [43]. Theorem 3 in [43] provides, for every $\varepsilon > 0$, an algorithm for `Max-Cut-Sc` whose running time is $O(n^{poly(\log(k)/\varepsilon)})$ which finds a solution $S \subseteq V$ with the following properties (in what follows S^* is an optimal solution for `Max-Cut-Sc`): (1) $\mathbb{E}[w^-(\delta(S))] \geq (0.843 - \varepsilon) \cdot w^-(\delta(S^*))$; and (2) for every color $i = 1, \ldots k$ $||S \cap V_i| - b_i| \leq \varepsilon n$. We prove that the rounding algorithm of [43] can also handle, in addition to the above two properties, the contribution of the + edges. Specifically, we show how one can easily change the SDP relaxation and then apply the rounding algorithm of [43] to also guarantee that

$$\mathbb{E}[w^+(E(S)) + w^+(E(\overline{S}))] \geq (0.843 - \varepsilon) \cdot (w^+(E(S^*)) + w^+(E(\overline{S^*}))),$$

where S^* here denotes an optimal solution to `Max-Agreement-Sc` for the correct choice of h.

First, let us start by formulating `Max-Agreement-Sc` as a quadratic optimization problem:

$$\begin{array}{lll} \max & \sum_{e=(i,j)\in E^-} w(e) \cdot (x_i - x_j)^2 + \sum_{e=(i,j)\in E^+} w(e) \cdot (1 - (x_i - x_j)^2) & \\ \text{s.t.} & \sum_{j \in F_i} x_j = b_i & i = 1, \ldots, k \\ & x_j \cdot (1 - x_j) = 0 & j = 1, \ldots, n \end{array}$$

Recall that in the above b_i equals $p_i \cdot h$ (for some value of h).

We denote the above quadratic problem by Q and the solutions to the ℓ-level Lasserre strengthening of the standard SDP relaxation of Q by $SoS(Q)$. A solution in $SoS(Q)$ can be represented by a collection of unit vectors $\mathbf{v_S}$ for all subsets $S \subseteq V$ ($|S| \leq \ell$). For completeness we write the ℓ-round SDP relaxation for the problem:

$$\max \quad \sum_{e=(i,j)\in E^-} w(e) \cdot \|\mathbf{v_i} - \mathbf{v_j}\|^2 + \sum_{e=(i,j)\in E^+} w(e) \cdot (1 - \|\mathbf{v_i} - \mathbf{v_j}\|^2)$$

s.t.	$\sum_{j\in F_i} \mathbf{v_0}\cdot\mathbf{v_j} = b_i$	$i = 1,\dots,k$
	$\mathbf{v_0}\cdot\mathbf{v_0} = 1$	
	$\mathbf{v_{S_1}}\cdot\mathbf{v_{S_2}} = \mathbf{v_{S_3}}\cdot\mathbf{v_{S_4}}$	$\forall S_1,S_2,S_3,S_4 \subseteq V$, $S_1\cup S_2 = S_3\cup S_4$ and $\lvert S_1\cup S_2\rvert \le \ell$
	$\mathbf{v_0}\cdot\mathbf{v_i} + \mathbf{v_j}\cdot\mathbf{v_0} - \mathbf{v_i}\mathbf{v_j} \le 1$	$1\le i,j\le n$
	$\mathbf{v_i}\cdot\mathbf{v_0} \ge \mathbf{v_i}\cdot\mathbf{v_j}$	$1\le i,j\le n$
	$\mathbf{v_i}\cdot\mathbf{v_j} \ge 0$	$1\le i,j\le n$

For completeness of presentation, let us now focus on defining the rounding algorithm of [43], as we require some of the notations in order to present the analysis of the uncut + edges. Recall that $\widetilde{\mathbf{w}}_\mathbf{i} \triangleq \mathbf{v_i} - \mu_i\mathbf{v_0}$ and $\mathbf{w_i} \triangleq \widetilde{\mathbf{w}}_\mathbf{i}/\|\widetilde{\mathbf{w}}_\mathbf{i}\|_2$. Let $\mathbf{W}$ and $\widetilde{\mathbf{W}}$ be the PSD correlation matrices defined by the above vectors, that is $\mathbf{W}_{ij} = \mathbf{w_i}\cdot\mathbf{w_j}$ and $\widetilde{\mathbf{W}}_{ij} = \widetilde{\mathbf{w}}_\mathbf{i}\cdot\widetilde{\mathbf{w}}_\mathbf{j}$, for every $1\le i,j\le n$. The following lemma is used to obtain the input vectors to the rounding algorithm, this lemma is based on [13] and [30] and appears as Lemma 10 in [43]. We refer the reader to [13, 30, 43] for its proof.

▶ **Lemma 16.** *Let $\varepsilon_0 \le 1, \ell \ge 1/\varepsilon_0^4+2$, for any solution in $SoS_\ell(Q)$ where $\ell \ge 1/\varepsilon_0^4+2$, there exists an efficiently computable solution in $SoS_{\ell-1/\varepsilon_0^4}(Q)$ such that $\sum_{i=1}^n\sum_{j=1}^n \widetilde{\mathbf{W}}_{ij}^2 \le \varepsilon_0^4 n^2$.*

Second, let us focus on the rounding algorithm. The input to the rounding algorithm is the vectors obtained by Lemma 16. To round the vectors, the algorithm performs a sticky Brownian motion inside the hypercube $[0,1]^n$, that is the random process $\{\mathbf{X}_t\}_{t\ge 0}$ which is defined as follows. The starting point of the random walk is $\mathbf{X}_0$ such that $(\mathbf{X}_0)_i = \mu_i$ for every $1\le i\le n$. Denote $\{\mathbf{B}_t\}_{t\ge 0}$ as the standard Brownian motion in $\mathbb{R}^n$. Let $\tau_1 = \inf\{t : \mathbf{X}_0 + \mathbf{W}^{1/2}\mathbf{B}_t \notin [0,1]^n\}$, then for all $0\le t\le\tau_1$: $\mathbf{X}_t = \mathbf{X}_0 + \mathbf{W}^{1/2}\mathbf{B}_t$. Let $A_t = \{i|(\mathbf{X}_t)_i \ne 0,1\}$ be the collection of active nodes at time t, and $F_t = \{\mathbf{x}\in[0,1]^n | x_i = (\mathbf{X}_t)_i, \forall i\notin A_t\}$. The covariance matrix $\mathbf{W}_t$ used for the random walk at time t is based on $\mathbf{W}$ and an entry in this matrix is not 0 only for the indices in A_t, i.e., $(\mathbf{W}_t)_{ij} = \mathbf{W}_{ij}$ if $i,j\in A_t$ (otherwise $(\mathbf{W}_t)_{ij} = 0$). After time τ_1 the random process is changed to $\mathbf{X}_t = \mathbf{X}_{\tau_1} + \mathbf{W}_{\tau_1}^{1/2}(\mathbf{B}_t - \mathbf{B}_{\tau_1})$, it is defined for $\tau_1\le t\le\tau_2$ where $\tau_2 = \inf\{t : \mathbf{X}_{\tau_1} + \mathbf{W}_{\tau_1}^{1/2}(\mathbf{B}_t - \mathbf{B}_{\tau_1}) \notin F_{\tau_1}\}$. In general, $\tau_i = \inf\{t : \mathbf{X}_{\tau_{i-1}} + \mathbf{W}_{\tau_{i-1}}^{1/2}(\mathbf{B}_t - \mathbf{B}_{\tau_{i-1}}) \notin F_{\tau_{i-1}}\}$ and when $\tau_{i-1}\le t\le\tau_i$ the process is defined as follows: $\mathbf{X}_t = \mathbf{X}_{\tau_{i-1}} + \mathbf{W}_{\tau_{i-1}}^{1/2}(\mathbf{B}_t - \mathbf{B}_{\tau_{i-1}})$. The algorithm does not terminate at time τ_n but it is stopped at a fixed pre-specified time τ (which is chosen to be $\Theta(\log(1/\varepsilon))$) and rounds to 1 the remaining nodes $i\in A_\tau$ with probability $(\mathbf{X}_\tau)_i$. The output cut $S\subseteq V$ contains all the nodes i for which $(\mathbf{X}_\tau)_i = 1$.

As previously mentioned, the algorithm for `Max-Cut-Sc` in [43] distinguishes between two cases. For instances with small optimal value a different approach was taken instead of the Brownian motion approach described above. However, since we use a sequence of `Max-Agreement-Sc` instances to solve `Fair-Max-Agreement[2]`, this case is not possible due to the following lemma. It states that the optimal value of `Fair-Max-Agreement[2]` is not small, hence for the correct choice of h the optimal value of `Max-Agreement-Sc` is also not small. Thus, we can focus solely on instances whose optimal value is not small.

▶ **Lemma 17.** *The optimal value of an instance $G = (V_1\cup\ldots\cup V_k, E^+\cup E^-)$ to `Fair-Max-Agreement[2]` is at least $(1/2 - \sum_{i=1}^k p_i/(2n))w(E)$.*

Proof. Let $\mathcal{C}$ be the output of Algorithm 1. The simple algorithm which outputs $\mathcal{S} = \{S,\overline{S}\}$ by placing all the nodes of each cluster $C\in\mathcal{C}$ together in S or $\overline{S}$ with probability $1/2$ (and independently over the different clusters in $\mathcal{C}$) results in a solution to `Fair-Max-Agreement[2]` with an expected value of $(1/2 - \sum_{i=1}^k p_i/(2n))w(E)$. This is true since $\mathbb{E}[v^+(\mathcal{S})] = 1/2\cdot w(E^+)$ and $\mathbb{E}[v^-(\mathcal{S})] = 1/2\cdot(1-\sum_{i=1}^k p_i/n)w(E^-)$ (the latter follows from Observation 13) ◀

The analysis of the Brownian motion based algorithm relies heavily on the following theorem which appears in [36]. Intuitively, this theorem captures the connection between diffusion processes and partial differential equations (see chapter 9 in [36]). We present it here since we require it for the analysis of the + edges.

▶ **Theorem 18** (Theorem 9 in [43] and Theorem 9.14 in [36])**.** *Given a domain $D = (0,1)^2 \subseteq \mathbb{R}^2$, suppose L is uniformly elliptic in D of the form*

$$L = \sum_{i,j=1}^{2} a_{ij}(x) \frac{\partial^2}{\partial x_i \partial x_j}$$

where $\frac{1}{2}\sigma\sigma^T = [a_{ij}]$ for $\sigma \in \mathbb{R}^{2\times 2}$. For $x \in D$ consider the process $X_t = X_0 + \sigma B_t$. Denote $\tau_D = inf\{t > 0; X_t \notin D\}$ (the stopping time of X_t). Let ϕ be a bounded continuous function on ∂D. Put $u(x) = E^x[\phi(X_{\tau_D})]$ where E^x denotes the expected value when $X_0 = x$. Then u solves the Dirichlet problem

1. *$Lu = 0$ in D.*
2. *$\lim_{x\to y} u(x) = \phi(y)$ for all regular $y \in \partial D$.*

Let us fix two nodes i and j and denote $\overline{\mathbf{X}}_t$ as the projection of the random process $\mathbf{X}_t$ to the coordinates i, j and let

$$\overline{\mathbf{W}} = \begin{pmatrix} 1 & \rho_{ij} \\ \rho_{ij} & 1 \end{pmatrix}.$$

In our analysis $\sigma = \overline{\mathbf{W}}^{1/2}$, i.e., the entries a_{ij} in the above theorem are the entries of $\overline{\mathbf{W}}$. When performing an edge-wise analysis we can consider the projection of $\mathbf{X}_t$ we described above. We note that the first guarantee in the following lemma is identical to [43], but is included for completeness. The novelty of the following theorem lies in the second guarantee.[5]

▶ **Lemma 19.** *Let $i, j \in V$ and $\mathbf{v_i}, \mathbf{v_j}$ the corresponding vectors in the SDP solution. It holds that*

1. $\Pr[\mathbf{X}_{\tau_n i} \neq \mathbf{X}_{\tau_n j}] \geq 0.843 \cdot \|\mathbf{v_i} - \mathbf{v_j}\|^2$.
2. $\Pr[\mathbf{X}_{\tau_n i} = \mathbf{X}_{\tau_n j}] \geq 0.843 \cdot (1 - \|\mathbf{v_i} - \mathbf{v_j}\|^2)$.

Proof. Guarantee 1 of the lemma is identical to Lemma 11 in [43], and thus its proof is omitted.

Let us focus on guarantee 2 above. Let us denote $\theta_{ij} = \arccos(\mathbf{w_i} \cdot \mathbf{w_j})$. Recall that $\|\mathbf{v_i}\|^2 = \mu_i$ and $\mathbf{v_i} = \mu_i \cdot \mathbf{v_0} + \sqrt{\mu_i - \mu_i^2} \cdot \mathbf{w_i}$. Therefore, the contribution of the + edges to the objective of the relaxation can be re-written as follows:

$$1 - \|\mathbf{v_i} - \mathbf{v_j}\|^2 = 1 - (\mu_i + \mu_j - 2 \cdot \mu_i \cdot \mu_j + 2\cos(\theta_{ij}) \cdot \sqrt{(\mu_i - \mu_i^2) \cdot (\mu_j - \mu_j^2)}).$$

For simplicity we denote $x = \mu_i, y = \mu_j, \theta = \theta_{ij}$ and the expression above as $SDP(x, y, \theta)$. Observe that the probability that the edge (i, j) is uncut equals the probability that the Brownian motion $\overline{\mathbf{X}}_t$ is absorbed in $(0, 0)$ or $(1, 1)$. Denote $u_\theta(x, y)$ as the probability of ending in $(0, 0)$ or $(1, 1)$ conditioned on starting the walk at point (x, y):

$$u_\theta(x, y) = \Pr[(\mathbf{X}_{\tau_n i} = \mathbf{X}_{\tau_n j}) | (\mathbf{X}_0)_i = x, (\mathbf{X}_0)_j = y].$$

[5] We mention that one can derive Lemma 19 via rotational symmetry of the boundary conditions of $\partial[0, 1]^2$ for both cut and uncut edges, and similar rotational symmetry of the contribution to the SDP relaxation of both cut and uncut edges.

Observe that the boundary condition on $u_\theta(x,y)$ is the following: $u_\theta(x,y) = 1-(x+y-2xy) \quad \forall (x,y) \in \partial[0,1]^2$. Following Theorem 18 u_θ is the unique solution to the Dirichlet problem:

$$\frac{\partial^2 u_\theta}{\partial^2 x} + \frac{\partial^2 u_\theta}{\partial^2 y} + 2\cos(\theta)\frac{\partial^2 u_\theta}{\partial x \partial y} = 0 \quad \forall (x,y) \in Int[0,1]^2$$

$$u_\theta(x,y) = 1-(x+y-2xy) \qquad \forall (x,y) \in \partial[0,1]^2$$

The problem above can be numerically solved for any configuration (x,y,θ). Therefore, the approximation ratio for uncut edges is $\min_{(x,y,\theta)\in F} \frac{u_\theta(x,y)}{SDP(x,y,\theta)}$ where F is the collection of all feasible configurations. Specifically, $(x,y,\theta) \in F$ if it satisfies the triangle inequalities which are derived from the ℓ-round SDP relaxation (see Appendix D Lemma 11 [43]). The numerical calculation via adaptation of the code used in [43] results in an approximation ratio of 0.843 for the uncut edges. ◀

Lemmas 19 and 17 are sufficient to extend the proof of Theorem 3 of [43] to `Fair-Max-Agreement[2]`, this is summarized in the following theorem.

▶ **Theorem 20.** *There exists a $O(n^{poly(\log(k)/\varepsilon)})$-time algorithm for `Fair-Max-Agreement[2]`, which for an instance $G = (V_1 \cup \ldots \cup V_k, E)$ outputs a $(0.843-\varepsilon, \varepsilon)$-approximation with high probability.*

Proof of Theorem 4. Follows from Theorem 20 and Lemma 15. ◀

3 Hardness of Fair-Min-Disagreement

In this section we present the hardness results for `Fair-Min-Disagreement`. First we prove Theorem 1.

Proof of Theorem 1. We present a reduction from the `3-Partition` problem, as defined in [8, 26]. In `3-Partition` we are given $n = 3\ell$ integer numbers $a_1, a_2, ..., a_n$ and a threshold A such that $\frac{A}{4} < a_i < \frac{A}{2}$ and $\sum_{i=1}^{n} a_i = \ell A$ (where $a_1, \ldots, a_n$ and A are polynomial in n). The goal is to decide if the numbers can be partitioned into triplets such that each triplet sums up to exactly A. This problem is known to be strongly NP-complete [26].

Given an instance of the `3-Partition` problem we construct a graph for the `Fair-Min-Disagreement` problem as follows (we denote the two colors by red and blue). For each number a_i construct a clique with a_i red nodes, the edges in this clique are all labeled with $+$. Additionally, construct ℓ cliques where each of them contains A blue nodes and the edges within such a clique are all labeled with $+$. For every pair of blue nodes which are not in the same clique, place an edge between them which is labeled with $-$. This finishes the definition of our instance for `Fair-Min-Disagreement`.

We claim that there is a solution to the given `3-Partition` instance if and only if there is a clustering of the `Fair-Min-Disagreement` instance whose cost is zero. Given a solution to the `3-Partition` instance we can construct a clustering of zero cost as follows. For each triplet $a_{i_1}, a_{i_2}, a_{i_3}$ in the solution for `3-Partition` (recall that $a_{i_1} + a_{i_2} + a_{i_3} = A$), define a cluster which contains the three red cliques corresponding to the numbers $a_{i_1}, a_{i_2}, a_{i_3}$ and a single blue clique of size A. One can note that this is a valid, i.e., fair, clustering since the number of red and blue nodes is equal in all clusters. Furthermore, there are no unclustered nodes since $\sum_{i=1}^{n} a_i = \ell A$. The cost of this clustering is zero since: (1) all cliques, either red or blue, are contained as a whole in a single cluster, and thus all $+$ edges are in agreement; and (2) every cluster contains exactly a single blue clique, and thus all $-$ edges are also in agreement.

Given a clustering of cost zero we prove that one can partition the numbers to triplets such that the sum of each triplet is exactly A. Note that each clique, either red or blue, in the graph is contained as a whole in a single cluster, otherwise there is a $+$ edges that is in disagreement which stands in contradiction to the fact that the clustering has zero cost. Moreover, each cluster contains exactly a single blue clique as a whole. The reason is that there cannot be no blue cliques in the cluster (if this occurs then the cluster has no blue nodes at all and this contradicts the fact the clustering is fair) and there cannot be two or more blue cliques in the cluster (if this occurs the cluster contains a $-$ edge and this contradicts the fact the clustering has zero cost). Thus, the number of blue nodes in the cluster is A. Since the clustering is fair the number of red nodes in the cluster is also A. Recall that every number a_i satisfies that $\frac{A}{4} < a_i < \frac{A}{2}$. Hence, the cluster must contain exactly three red cliques that correspond to three numbers that sum up exactly to A. Therefore, the triplets we define as a solution to `3-Partition` are those that correspond to the three red cliques in each cluster. ◀

Let us now prove that a bi-criteria approximation is also not possible unless $P = NP$.

Proof of Theorem 2. We present a reduction from `Triangle-Partition`, which is known to be NP-hard [22]. In this problem the goal is to decide whether there is a set of node-disjoint triangles in a tripartite graph which covers all the nodes of the given tripartite graph. Note that without loss of generality one can assume that each of the three parts of the tripartite graph contains the same number of nodes. Otherwise, it is clear that the input graph cannot have all its nodes covered by node-disjoint triangles.

Given an instance $G = (A \cup B \cup C, E)$ to `Triangle-Partition` we construct a graph $G' = (A \cup B \cup C, E')$ for `Fair-Min-Disagreement` as follows. Each part of the three parts of G is given a unique color, i.e., $V_1 = A$, $V_2 = B$, and $V_3 = C$. Define the edges in G' as follows: $E'^{-} \triangleq \{(u, v)|(u, v) \notin E\}$ and $E'^{+} \triangleq \emptyset$. This finishes the definition of our instance for `Fair-Min-Disagreement`.

We claim that there is a solution to `Triangle-Partition` if and only if there is a solution $\mathcal{C} = \{C_1, \ldots, C_l\}$ to `Fair-Min-Disagreement` whose cost is zero and it satisfies that for every $1 \leq r \leq l$: $|C_r \cap V_i|/|C_r \cap V_j| \leq (1 + \varepsilon)$ for every (ordered) pair of colors i and j.

Given a solution to `Triangle-Partition` we can construct a solution to `Fair-Min--Disagreement` by setting every triangle to be a different cluster. The nodes in each triangle are connected by edges in E. Therefore, there are no $-$ edges between these nodes in E'. Since there are no $+$ edges in E', we can conclude that the cost of this solution for `Fair-Min-Disagreement` equals zero. Moreover, each cluster in the solution for `Fair-Min--Disagreement` contains exactly one node from each of the three colors. Hence, we proved the existence of the desired solution for `Fair-Min-Disagreement`.

Let $\mathcal{C} = \{C_1, C_2, ..., C_l\}$ be a solution to `Fair-Min-Disagreement` that has zero cost and satisfies that for every $1 \leq r \leq l$: $|C_r \cap V_i|/|C_r \cap V_j| \leq (1 + \varepsilon)$ for every (ordered) pair of colors i and j. Note that for every $1 \leq r \leq l$ and $i = 1, 2, 3$: $|C_r \cap V_i| \leq 1$. The reason for that is that two nodes of the same color are connected with a $-$ edge in E'. Since $\mathcal{C}$ has zero cost, any nodes of the same color cannot be in the same cluster. Moreover, for every $1 \leq r \leq l$ and $i = 1, 2, 3$: $|C_r \cap V_i| > 0$. The reason for that is that if there is a (non-empty) cluster C_r and a color i for which $|C_r \cap V_i| = 0$, then C_r contains at least one node of color j, $j \neq i$. For this ordered pair of colors (j and i) the condition on $\mathcal{C}$ is violated. Therefore, every cluster C_r contains exactly one node from every color, i.e., one node from every part of G. Because the cost of the clustering is zero, there are no $-$ edges from E' inside each cluster which means that it forms a triangle in G. Clearly, these triangles are node-disjoint and contain all nodes in $A \cup B \cup C$ since $\mathcal{C}$ is a partition of $A \cup B \cup C$. This finishes the proof. ◀

References

1 Saba Ahmadi, Sainyam Galhotra, Barna Saha, and Roy Schwartz. Fair correlation clustering, 2020. doi:10.48550/ARXIV.2002.03508.

2 Sara Ahmadian, Alessandro Epasto, Marina Knittel, Ravi Kumar, Mohammad Mahdian, Benjamin Moseley, Philip Pham, Sergei Vassilvitskii, and Yuyan Wang. Fair hierarchical clustering. *Advances in Neural Information Processing Systems*, 33:21050–21060, 2020.

3 Sara Ahmadian, Alessandro Epasto, Ravi Kumar, and Mohammad Mahdian. Clustering without over-representation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '19, pages 267–275, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3292500.3330987.

4 Sara Ahmadian, Alessandro Epasto, Ravi Kumar, and Mohammad Mahdian. Fair correlation clustering. *ArXiv*, abs/2002.02274, 2020.

5 Nir Ailon, Noa Avigdor-Elgrabli, Edo Liberty, and Anke van Zuylen. Improved approximation algorithms for bipartite correlation clustering. *SIAM Journal on Computing*, 41(5):1110–1121, 2012. doi:10.1137/110848712.

6 Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: Ranking and clustering. *J. ACM*, 55(5), November 2008. doi:10.1145/1411509.1411513.

7 Noga Amit. *The bicluster graph editing problem.* PhD thesis, Tel Aviv University, 2004.

8 Konstantin Andreev and Harald Räcke. Balanced graph partitioning. *Theory of Computing Systems*, 39:929–939, November 2006. doi:10.1007/s00224-006-1350-7.

9 Megasthenis Asteris, Anastasios Kyrillidis, Dimitris Papailiopoulos, and Alexandros Dimakis. Bipartite correlation clustering: Maximizing agreements. In *Artificial Intelligence and Statistics*, pages 121–129, 2016.

10 Per Austrin, Siavosh Benabbas, and Konstantinos Georgiou. Better balance by being biased: A 0.8776-approximation for max bisection. *ACM Transactions on Algorithms*, 13:1–27, October 2016. doi:10.1145/2907052.

11 Arturs Backurs, Piotr Indyk, Krzysztof Onak, Baruch Schieber, Ali Vakilian, and Tal Wagner. Scalable fair clustering. In *International Conference on Machine Learning*, pages 405–413. PMLR, 2019.

12 Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. In *Machine Learning*, pages 238–247, 2002.

13 Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. *Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS*, April 2011. doi:10.1109/FOCS.2011.95.

14 Suman Bera, Deeparnab Chakrabarty, Nicolas Flores, and Maryam Negahbani. Fair algorithms for clustering. *Advances in Neural Information Processing Systems*, 32, 2019.

15 Ioana O. Bercea, Samir Khuller, Clemens Rösner, Melanie Schmidt, Martin Groß, Aounon Kumar, and Daniel R. Schmidt. On the cost of essentially fair clusterings. In Dimitris Achlioptas and Laszlo A. Vegh, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2019*, Leibniz International Proceedings in Informatics, LIPIcs. Schloss Dagstuhl- Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, September 2019. doi:10.4230/LIPIcs.APPROX-RANDOM.2019.18.

16 Simon Caton and Christian Haas. Fairness in machine learning: A survey. *arXiv preprint arXiv:2010.04053*, 2020.

17 M. Charikar and A. Wirth. Maximizing quadratic programs: extending grothendieck's inequality. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 54–60, 2004. doi:10.1109/FOCS.2004.39.

18 Moses Charikar, Venkatesan Guruswami, and Anthony Wirth. Clustering with qualitative information. *Journal of Computer and System Sciences*, 71(3):360–383, 2005. Learning Theory 2003. doi:10.1016/j.jcss.2004.10.012.

19 Shuchi Chawla, Konstantin Makarychev, Tselil Schramm, and Grigory Yaroslavtsev. Near optimal lp rounding algorithm for correlationclustering on complete and complete k-partite graphs. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, pages 219–228, New York, NY, USA, 2015. Association for Computing Machinery. doi:10.1145/2746539.2746604.

20 Anshuman Chhabra, Karina Masalkovaitė, and Prasant Mohapatra. An overview of fairness in clustering. *IEEE Access*, 9:130698–130720, 2021. doi:10.1109/ACCESS.2021.3114099.

21 Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Fair clustering through fairlets. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pages 5036–5044, Red Hook, NY, USA, 2017. Curran Associates Inc.

22 Charles J. Colbourn. The complexity of completing partial latin squares. *Discrete Applied Mathematics*, 8(1):25–30, 1984. doi:10.1016/0166-218X(84)90075-1.

23 Tom Coleman, James Saunderson, and Anthony Wirth. A local-search 2-approximation for 2-correlation-clustering. In *European Symposium on Algorithms*, ESA '08, pages 308–319, 2008.

24 Erik D. Demaine, Dotan Emanuel, Amos Fiat, and Nicole Immorlica. Correlation clustering in general weighted graphs. *Theoretical Computer Science*, 361(2):172–187, 2006. Approximation and Online Algorithms. doi:10.1016/j.tcs.2006.05.008.

25 Dotan Emanuel and Amos Fiat. Correlation clustering – minimizing disagreements on arbitrary weighted graphs. In Giuseppe Di Battista and Uri Zwick, editors, *Algorithms - ESA 2003*, pages 208–220, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.

26 Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman and Co., USA, 1990.

27 Mehrdad Ghadiri, Samira Samadi, and Santosh Vempala. Socially fair k-means clustering. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, pages 438–448, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3442188.3445906.

28 Ioannis Giotis and Venkatesan Guruswami. Correlation clustering with a fixed number of clusters. *Theory of Computing*, 2, May 2005. doi:10.1145/1109557.1109686.

29 Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, November 1995. doi:10.1145/227683.227684.

30 Venkatesan Guruswami and Ali Kemal Sinop. Lasserre hierarchy, higher eigenvalues, and approximation schemes for quadratic integer programming with PSD objectives. *CoRR*, abs/1104.4746, 2011. arXiv:1104.4746.

31 Eran Halperin and Uri Zwick. A unified framework for obtaining improved approximation algorithms for maximum graph bisection problems. In *Proceedings of the 8th International IPCO Conference on Integer Programming and Combinatorial Optimization*, pages 210–225, Berlin, Heidelberg, 2001. Springer-Verlag.

32 Qiaoming Han, Yinyu Ye, and Jiawei Zhang. An improved rounding method and semidefinite programming relaxation for graph partition. *Mathematical Programming*, 92:509–535, 2002.

33 Matthäus Kleindessner, Samira Samadi, Pranjal Awasthi, and Jamie Morgenstern. Guarantees for spectral clustering with fairness constraints. In *International Conference on Machine Learning*, pages 3458–3467. PMLR, 2019.

34 Jean B. Lasserre. An explicit exact sdp relaxation for nonlinear 0-1 programs. In Karen Aardal and Bert Gerards, editors, *Integer Programming and Combinatorial Optimization*, pages 293–303, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.

35 Andrew McCallum and Ben Wellner. Toward conditional models of identity uncertainty with application to proper noun coreference. In *Proceedings of the IJCAI-2003 Workshop on Information Integration on the Web*, pages 79–86, August 2003.

36 Bernt Øksendal. *Stochastic Differential Equations*. Springer Berlin Heidelberg, 2003. doi:10.1007/978-3-642-14394-6.

37 Prasad Raghavendra and Ning Tan. Approximating csps with global cardinality constraints using sdp hierarchies. In *SODA*, 2012.

38 Melanie Schmidt, Chris Schwiegelshohn, and Christian Sohler. Fair coresets and streaming algorithms for fair k-means. In *Approximation and Online Algorithms: 17th International Workshop, WAOA 2019, Munich, Germany, September 12–13, 2019, Revised Selected Papers*, pages 232–251, Berlin, Heidelberg, 2019. Springer-Verlag. doi:10.1007/978-3-030-39479-0_16.

39 Chaitanya Swamy. Correlation clustering: Maximizing agreements via semidefinite programming. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '04, pages 526–527, USA, 2004. Society for Industrial and Applied Mathematics.

40 Jurgen Van Gael and Xiaojin Zhu. Correlation clustering for crosslingual link detection. In *IJCAI*, pages 1744–1749, 2007.

41 Anthony Wirth. Correlation clustering. In Claude Sammut and Geoffrey I. Webb, editors, *Encyclopedia of Machine Learning*, pages 227–231. Springer US, Boston, MA, 2010. doi:10.1007/978-0-387-30164-8_176.

42 Chenchen Wu, Donglei Du, and Dachuan Xu. An improved semidefinite programming hierarchies rounding approximation algorithm for maximum graph bisection problems. In Ding-Zhu Du and Guochuan Zhang, editors, *Computing and Combinatorics*, pages 304–315, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

43 Sepehr Abbasi Zadeh, Nikhil Bansal, Guru Guruganesh, Aleksandar Nikolov, Roy Schwartz, and Mohit Singh. Sticky brownian rounding and its applications to constraint satisfaction problems. *ArXiv*, abs/1812.07769, 2020.

A Proof of Theorem 7

Proof. The algorithm chooses the best from the following two solutions: a single cluster containing all nodes and a solution with all clusters of size two that correspond to M^- from Observation 6 (we note M^- can be computed efficiently by finding a minimum cost perfect matching in a bipartite graph). The former solution has value of at least $w(E^+)$, whereas the latter solution has value of at least $w(E^-) - w(M^- \cap E^-)$. Following observation 6, $OPT \leq w(E^+) + w(E^-) - w(M^- \cap E^-)$. If $w(E^+) > w(E^-) - w(M \cap E^-)$ we note that the single cluster solution has value of at least $1/2 \cdot OPT$. Otherwise the solution with all clusters of size two, that correspond to M^-, has value of at least $1/2 \cdot OPT$. Hence, the above algorithm achieves an approximation of $1/2$. ◀

B Proof of Lemma 8

Proof. Given a clustering $\mathcal{C} = \{C_1, \ldots, C_l\}$ we can construct a clustering that has only two clusters $\mathcal{S} = \{S, \overline{S}\}$ as follows. For every C_i, with a uniform probability (and independently over the clusters) we place all the nodes of C_i either in S or in $\overline{S}$. Note that all $+$ edges that are in agreement in $\mathcal{C}$ always remain in agreement in $\mathcal{S}$, thus $v^+(\mathcal{S}) \geq v^+(\mathcal{C})$. Moreover, the probability of every $-$ edge that is in agreement in $\mathcal{C}$ to still be in agreement in $\mathcal{S}$ is exactly $1/2$. Therefore, $\mathbb{E}[v^-(\mathcal{S})] \geq \frac{1}{2} v^-(\mathcal{C})$. Hence, we can conclude that $\mathbb{E}[v(\mathcal{S})] \geq v^+(\mathcal{C}) + \frac{1}{2} v^-(\mathcal{C})$ (so there exists a cluster $\mathcal{S}$ with a value of at least $v^+(\mathcal{C}) + \frac{1}{2} v^-(\mathcal{C})$). This finishes the proof. ◀

C Proof of Lemma 9

Proof. We are given an instance of `Fair-Max-Agreement` with two colors and a ratio of $1:1$. Let $\mathcal{C}^* = \{C_1^*, \ldots, C_l^*\}$ be an optimal clustering for this instance. Following observation 6 we can output the clustering M^- induces, and we note that the value of this clustering is at least $v^-(\mathcal{C}^*)$. Following Lemma 8 when applied for $\mathcal{C}^*$ the given α-approximation algorithm can be used to obtain a clustering $\mathcal{S} = \{S, \overline{S}\}$ with value $v(\mathcal{S}) \geq \alpha \cdot (v^+(\mathcal{C}^*) + \frac{1}{2}v^-(\mathcal{C}^*))$. Therefore, choosing the best of the above two clusterings, we can output a solution whose value is at least $\max\{\alpha \cdot (v^+(\mathcal{C}^*) + \frac{1}{2}v^-(\mathcal{C}^*)), v^-(\mathcal{C}^*)\}$ (we denote this value by y). Now we show that for $0 < \alpha < 1$ it holds that $y \geq \frac{2\alpha}{2+\alpha} \cdot v(\mathcal{C}^*)$.

The first case is when $y = v^-(\mathcal{C}^*)$ and the second case is when $y = \alpha \cdot v^+(\mathcal{C}^*) + \frac{1}{2}\alpha \cdot v^-(\mathcal{C}^*)$. Let us focus on the first case, and note that assuming $y = v^-(\mathcal{C}^*)$, the definition of y implies $v^+(\mathcal{C}^*) \leq (1/\alpha - 1/2)v^-(\mathcal{C}^*)$. This in turn implies that:

$$v(\mathcal{C}^*) = v^+(\mathcal{C}^*) + v^-(\mathcal{C}^*) \leq v^-(\mathcal{C}^*)\,(1 + (1/\alpha - 1/2)) = (2+\alpha)/(2\alpha) \cdot y.$$

This concludes the proof for the first case. Let us now focus on the second case, and note that assuming $y = \alpha \cdot v^+(\mathcal{C}^*) + \frac{1}{2}\alpha \cdot v^-(\mathcal{C}^*)$, the definition of y implies $v^-(\mathcal{C}^*) \leq (2\alpha)/(2-\alpha) \cdot v^+(\mathcal{C}^*)$. This in turn implies that:

$$\begin{aligned} v(\mathcal{C}^*) = v^+(\mathcal{C}^*) + v^-(\mathcal{C}^*) &= v^+(\mathcal{C}^*) + (2+\alpha)/4 \cdot v^-(\mathcal{C}^*) + (2-\alpha)/4 \cdot v^-(\mathcal{C}^*) \\ &\leq v^+(\mathcal{C}^*) + (2+\alpha)/4 \cdot v^-(\mathcal{C}^*) + (2-\alpha)/4 \cdot (2\alpha)/(2-\alpha) \cdot v^+(\mathcal{C}^*) \\ &= \frac{2+\alpha}{2\alpha}\left(\alpha \cdot v^+(\mathcal{C}^*) + \frac{\alpha}{2} \cdot v^-(\mathcal{C}^*)\right) = \frac{2+\alpha}{2\alpha} \cdot y. \end{aligned}$$

This concludes the proof for the second case. ◀

D Approximating Max-Agreement-Bisection

We claim that one can use the algorithm of Wu, Du and Xu [42], who built upon the work of Austrin, Benabbas and Georgiou [10] for `Max-Bisection`, to obtain a good approximation for `Max-Agreement-Bisection`. The algorithms of [10, 42] both perform the following three phases ([10] for the `Max-Bisection` problem and [42] for the general graph bisection problem). In the first phase the following ℓ-round Lasserre SDP relaxation is solved:

$$\begin{array}{lll} \max & \sum_{e=(i,j)\in E^+} w(e)(1/2 + 1/2\langle \mathbf{v_i}, \mathbf{v_j}\rangle)) + \sum_{e=(i,j)\in E^-} w(e)(1/2 - 1/2\langle \mathbf{v_i}, \mathbf{v_j}\rangle)) & \\ \text{s.t.} & \langle \mathbf{v}_\emptyset, \sum_{i\in V} \mathbf{v}_{\mathbf{S}\triangle\{\mathbf{i}\}}\rangle = 0 & \forall S \subseteq V, |S| < \ell \\ & \langle \mathbf{v_{S_1}}, \mathbf{v_{S_2}}\rangle = \langle \mathbf{v_{S_3}}, \mathbf{v_{S_4}}\rangle & \forall S_1, S_2, S_3, S_4 \subseteq V, \\ & & |S_1|, |S_2|, |S_3|, |S_4| \leq \ell, \\ & & S_1 \triangle S_2 = S_3 \triangle S_4 \\ & \langle \mathbf{v}_\emptyset, \mathbf{v}_\emptyset\rangle = 1 & \end{array}$$

Let us denote by $\{\mathbf{v_S}^*\}_{S\subseteq V, |S|<\ell}$ an optimal solution to the above relaxation. The following theorem shows how one can extract vectors $\{\mathbf{v_i}\}_{i=0}^n$, from $\{\mathbf{v_S}^*\}_{S\subseteq V, |S|<\ell}$, such that the value of the objective does not deteriorate much and the vectors $\{\mathbf{v_i}\}_{i=0}^n$ have low correlation. Before formally stating the above, we introduce the following notation:

$$SDPVal(\{\mathbf{v_i}\}) \triangleq \sum_{e=(i,j)\in E^+} w(e)(1/2 + 1/2\langle \mathbf{v_i}, \mathbf{v_j}\rangle)) + \sum_{e=(i,j)\in E^-} w(e)(1/2 - 1/2\langle \mathbf{v_i}, \mathbf{v_j}\rangle)).$$

When reading the following theorem, the reader should recall the definitions of μ_i, $\rho_{i,j}$, and $\mathbf{w}_i$, given in Section 1.3.

▶ **Theorem 21** (Theorem 3.1 in [10], Theorem 2 in [42]). *There is an algorithm which given a graph $G = (V, E)$ and $t \in \mathbb{N}^{\geq 1}$ outputs a set of vectors $\{\mathbf{v_i}\}_{i=0}^n$ in time $n^{O(t)}$ such that:*

1. $SDPVal(\{\mathbf{v_i}\}) \geq SDPVal(\{\mathbf{v_i}^*\}) - 10t^{-\frac{1}{2}}$.
2. $\sum_{i=1}^n \langle \mathbf{v_0}, \mathbf{v_i} \rangle = 0$.
3. *The following triangle inequalities are satisfied for every* $1 \leq i, j \leq n$*:*

$$\mu_i + \mu_j + \rho_{ij} \geq -1, \mu_i - \mu_j - \rho_{ij} \geq -1$$

$$-\mu_i + \mu_j - \rho_{ij} \geq -1, -\mu_i - \mu_j + \rho_{ij} \geq -1$$

4. $\mathbb{E}_{i,j \in V}[|\langle \mathbf{w_i}, \mathbf{w_j} \rangle|] \leq t^{-\frac{1}{4}}$.

We note that the above theorem was proved in [10] for the objective of `Max-Bisection` and in [42] for the objective of `Max-`$\frac{n}{2}$`-Uncut` (both heavily rely on Raghavendra and Tan [37]). However, one can note that the same proof holds for our definition of $SDPVal(\{\mathbf{v}_i\})$ for `Max-Agreement-Bisection`.

In the second phase the rounding algorithm of [10] uses $\{\mathbf{v_i}\}_{i=0}^n$ to extract a cut $\tilde{\mathcal{S}} = \{\tilde{S}, V \setminus \tilde{S}\}$. This rounding algorithm has the following properties: (1) the rounding does not depend on the coefficients c_0, c_1, c_2, c_3; and (2) the analysis is performed edge-wise, i.e., the ratio of the probability of an edge being satisfied by the rounding algorithm to the contribution of the same edge to the value of the relaxation is lower bounded. We note that this cut might not be a bisection, i.e., $|\tilde{S}|$ might not equal $n/2$, thus corrections must be made. The following lemma is immediate from Lemma 4 in [42] and Lemma 3.2 in [10], where the former is for the `Max-`$\frac{n}{2}$`-Uncut` objective and the latter for the `Max-Bisection` objective.

▶ **Lemma 22.** *(following Lemma 4 in [42] and Lemma 3.2 in [10])*
$\mathbb{E}[v(\tilde{\mathcal{S}})] \geq \alpha_0 \cdot SDPVal(\{\mathbf{v_i}\}_{i=0}^n)$*, where* $\alpha_0 \geq 0.8776$.

The last phase is a size adjusting phase in which a subset of vertices from the larger side of the cut $\tilde{\mathcal{S}}$ is moved to the smaller side of the cut in order to create a bisection. This is performed either by choosing a random subset (as is done in [10]), or equivalently, greedily (as is done [42]). This phase incurs an additive loss of $o(1)$ in the approximation guarantee. We can choose any of the above two options. The following lemma summarizes the approximation guarantee for `Max-Agreement-Bisection`, its proof follows from Theorem 21 and Lemma 22 similarly to [10, 42].

▶ **Lemma 23.** *`Max-Agreement-Bisection` is approximable in polynomial time to within a factor 0.8776.*