

Karthekeyan Chandrasekaran, [László A. Végh](#), Santosh S. Vempala

The cutting plane method is polynomial for perfect matchings

Article (Accepted version)
(Refereed)

Original citation: Chandrasekaran, Karthekeyan, Végh, László A. and Vempala, Santosh S. (2015) *The cutting plane method is polynomial for perfect matchings*. [Mathematics of Operations Research](#), 41 (1). pp. 23-48. ISSN 0364-765X
DOI: [10.1287/moor.2015.0714](https://doi.org/10.1287/moor.2015.0714)

© 2016 [INFORMS](#)

This version available at: <http://eprints.lse.ac.uk/66123/>
Available in LSE Research Online: April 2016

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

The Cutting Plane Method is Polynomial for Perfect Matchings ^{*}

Karthekeyan Chandrasekaran [†] László A. Végh [‡] Santosh S. Vempala [§]

Abstract

The cutting plane approach to finding minimum-cost perfect matchings has been discussed by several authors over past decades [24, 16, 19, 28, 12], and its convergence has been an open question. We give a cutting plane algorithm that converges in polynomial-time using only Edmonds' blossom inequalities; it maintains half-integral intermediate LP solutions supported by a disjoint union of odd cycles and edges. Our main insight is a method to retain only a subset of the previously added cutting planes based on their dual values. This allows us to quickly find violated blossom inequalities and argue convergence by tracking the number of odd cycles in the support of intermediate solutions.

1 Introduction.

Integer programming is a powerful and widely used approach for modeling and solving discrete optimization problems [22, 25]. Not surprisingly, it is NP-complete and the fastest known algorithms are exponential in the number of variables (roughly $n^{O(n)}$ [18]). In spite of this intractability, integer programs of considerable sizes are routinely solved in practice. A popular approach is the cutting plane method, proposed by Dantzig, Fulkerson and Johnson [9] and pioneered by Gomory [13, 14, 15]. This approach can be summarized as follows:

1. Solve a linear programming relaxation (LP) of the given integer program (IP) to obtain a basic optimal solution x .
2. If x is integral, terminate. If x is not integral, find a linear inequality that is valid for the convex hull of all integer solutions but violated by x .
3. Add the inequality to the current LP, possibly drop some other inequalities and solve the resulting LP to obtain a basic optimal solution x . Go back to Step 2.

For the method to be efficient, we require the following: (a) an efficient procedure for finding a violated inequality (called a cutting plane), (b) convergence of the method to an integral solution using the efficient cut-generation procedure and (c) a bound on the number of iterations to convergence. Gomory gave the first efficient cut-generation procedure and showed that the cutting plane method implemented using his procedure always converges to an integral solution [15]. Today, there is a rich theory on the choice of cutting planes, both in general and for specific problems

^{*}This work was supported in part by NSF award AF0915903, and the second author was also supported by NSF Grant CCF-0914732. This work was done while the first two authors were affiliated with the College of Computing, Georgia Institute of Technology.

[†]School of Engineering and Applied Sciences, Harvard University. Email: karthe@seas.harvard.edu.

[‡]Department of Management, London School of Economics. Email: l.vegh@lse.ac.uk.

[§]College of Computing, Georgia Institute of Technology. Email: vempala@cc.gatech.edu.

of interest. This theory includes interesting families of cutting planes with efficient cut-generation procedures [13, 1, 6, 2, 7, 23, 3, 20, 27], valid inequalities, closure properties and a classification of the strength of inequalities based on their *rank* with respect to cut-generating procedures [8] (e.g., the Chvátal-Gomory rank [6]), and testifies to the power and generality of the cutting plane method.

To our knowledge, however, there are no polynomial bounds on the number of iterations to convergence of the cutting plane method even for specific problems using specific cut-generation procedures. The best bound for general 0-1 integer programs remains Gomory's bound of 2^n [15]. It is possible that such a bound can be significantly improved for IPs with small Chvátal-Gomory rank [12]. A more realistic possibility is that the approach is provably efficient for combinatorial optimization problems that are known to be solvable in polynomial time. An ideal candidate could be a problem that (i) has a polynomial-size IP-description (the LP-relaxation is polynomial-size), and (ii) the convex-hull of integer solutions has a polynomial-time separation oracle. Such a problem admits a polynomial-time algorithm via the Ellipsoid method [17]. Perhaps the first such interesting problem is minimum-cost perfect matching: *given a graph with costs on the edges, find a perfect matching of minimum total cost*. This is a very well-studied problem with efficient algorithms [19, 26].

A polyhedral characterization of the matching problem was discovered by Edmonds [10]. Basic solutions (extreme points of the polytope) of the following linear program correspond to perfect matchings of the graph.

$$\begin{aligned}
& \min \sum_{uv \in E} c(uv)x(uv) \\
& x(\delta(u)) = 1 \quad \forall u \in V \\
& x(\delta(S)) \geq 1 \quad \forall S \subsetneq V, |S| \text{ odd}, 3 \leq |S| \leq |V| - 3 \\
& x \geq 0
\end{aligned} \tag{P}$$

The relaxation with only the degree and nonnegativity constraints, known as the *bipartite relaxation*, suffices to characterize the convex-hull of perfect matchings in bipartite graphs, and serves as a natural starting relaxation. The inequalities corresponding to sets of odd cardinality greater than 1 are called *blossom* inequalities. These inequalities have Chvátal rank 1, i.e., applying one round of all possible Gomory cuts to the bipartite relaxation suffices to recover the perfect matching polytope of any graph [6]. Moreover, although the number of blossom inequalities is exponential in the size of the graph, for any point not in the perfect matching polytope, a violated (blossom) inequality can be found in polynomial time [24]. This suggests a natural cutting plane algorithm (Algorithm 1), proposed by Padberg and Rao [24] and discussed by Lovász and Plummer [19]. Experimental evidence suggesting that this method converges quickly was given by Grötschel and Holland [16], by Trick [28], and by Fischetti and Lodi [12]. It has been open to rigorously explain their findings. In this paper, we address the question of whether the method can be implemented to converge in polynomial time.

Known polynomial-time algorithms for minimum-cost perfect matching are variants of Edmonds' weighted matching algorithm [10]. A natural idea is to interpret the Edmonds' algorithm, which maintains a partial matching and shrinks and unshrinks odd sets, as a cutting plane algorithm, possibly by adding cuts corresponding to the shrunk sets in the iterations of Edmonds' algorithm. However, there seems to be no correspondence between LP solutions and partial matchings and shrunk sets in his algorithm. It is even possible that the initial bipartite relaxation already has an integer optimal solution, whereas Edmonds' algorithm proceeds by shrinking and deshrinking a long sequence of odd sets. So we take a different route.

Algorithm 1 Cutting plane method for perfect matching

1. Start by solving the bipartite relaxation.
 2. While the current solution has a fractional coordinate,
 - (a) Find a violated blossom inequality and add it to the LP.
 - (b) Solve the new LP.
-

The bipartite relaxation has the nice property that any basic solution is half-integral and the connected components of its support are either edges or odd cycles. This makes it particularly easy to find violated blossom inequalities—any odd component of the support gives one. This is also the simplest heuristic that is employed in the implementations [16, 28] for finding violated blossom inequalities. However, if we have a fractional solution in a later phase, there is no guarantee that we can find an odd connected component whose blossom inequality is violated, and therefore sophisticated and significantly slower separation methods are needed for finding cutting planes, e.g., the Padberg-Rao procedure [24]. Thus, it is natural to wonder if there is a choice of cutting planes that maintains half-integrality of intermediate LP optimal solutions.




Figure 1: An example where half-integrality is not preserved by the black-box LP solver.

At first sight, maintaining half-integrality using a black-box LP solver seems to be impossible. Figure 1 shows an example where the starting solution consists of two odd cycles. There is only one reasonable way to impose cuts, and it leads to a basic feasible solution that is not half-integral. We observe however, that in the example, the bipartite relaxation also has an integer optimal solution. The issue here seems to be the existence of multiple basic optimal solutions. To avoid such degeneracy, we will ensure that all linear systems that we encounter have unique optimal solutions.

This uniqueness is achieved by a simple deterministic perturbation of the integer cost function, which increases the input size polynomially. However, this perturbation is only a first step towards maintaining half-integrality of intermediate LP optima. As we will see presently, more careful cut retention and cut addition procedures are needed even to maintain half-integrality.

1.1 Main result.

To state our main result, we first recall the definition of a laminar family: A family \mathcal{F} of subsets of V is called *laminar*, if for any two sets $X, Y \in \mathcal{F}$, either $X \cap Y = \emptyset$ or $X \subseteq Y$ or $Y \subseteq X$.

Next we define a perturbation to the cost function that will help avoid some degeneracies. Given an integer cost function $c : E \rightarrow \mathbb{Z}$ on the edges of a graph $G = (V, E)$, let us define the perturbation \tilde{c} by ordering the edges arbitrarily, and increasing the cost of edge i by $1/2^i$.

We are now ready to state our main theorem.

Theorem 1.1. *Let $G = (V, E)$ be a graph on n nodes with edge costs $c : E \rightarrow \mathbb{Z}$ and let \tilde{c} denote the perturbation of c . Then, there exists an implementation of the cutting plane method that finds the minimum \tilde{c} -cost perfect matching such that*

- (i) *every intermediate LP is defined by the bipartite relaxation constraints and a collection of blossom inequalities corresponding to a laminar family of odd subsets,*
- (ii) *every intermediate LP optimum is unique, half-integral, and the connected components of its support are either edges or odd cycles, and*
- (iii) *the total number of iterations to arrive at a minimum \tilde{c} -cost perfect matching is $O(n \log n)$.*

The collection of blossom inequalities used at each step can be identified by solving an LP of the same size as the current LP. Further, the minimum \tilde{c} -cost perfect matching is also a minimum c -cost perfect matching.

To our knowledge, this is the first polynomial bound on the convergence of a cutting plane method for matchings using a black-box LP solver. It is easy to verify that for an n -vertex graph, a laminar family of nontrivial odd sets may have at most $n/2$ members, hence every intermediate LP has at most $3n/2$ inequalities apart from the non-negativity constraints. This ensures that the intermediate LPs do not blow-up in size. While the LPs could be solved using a black-box LP solver, we also provide a combinatorial algorithm that could be used to solve them.

1.2 Related work.

We discovered after completing this work that a very similar question was addressed by Bunch in his thesis [4]. He proposed a cutting plane algorithm for the more general b -matching problem, maintaining a sequence of half-integral primal solutions. The intermediate LPs are solved using a primal-dual simplex method. Each new cut and simplex pivot is chosen carefully so that the primal/dual solution resulting after the pivot step can also be obtained through a combinatorial operation in an associated graph. So, the intermediate solutions correspond to the intermediate solutions of a combinatorial algorithm. This combinatorial algorithm is a variant of Miller-Pekny's combinatorial algorithm [21] that proceeds by maintaining a sequence of half-integral solutions and is known to terminate in polynomial-time. Consequently, the intermediate primal solutions of Bunch's algorithm are also half-integral and the algorithm terminates in polynomial time.

The main advantages of our algorithm compared to Bunch's work are the following:

- We present a purely cutting plane method, using a black-box LP solver. In contrast, the cut generation method in Bunch's algorithm, similar to Gomory cuts, relies heavily on the optimal simplex tableaux that is derived by the primal-dual simplex method; this mimics a certain combinatorial algorithm. We will also refer to a similar combinatorial algorithm, however, it will only be used in the analysis.
- We provide a simple and concise sufficient condition for the LP defined by the bipartite relaxation constraints and a subset of blossom inequalities to have a half-integral optimum (Lemma 3.1). No such insight is given in [4].

- Our treatment is substantially simpler, both regarding the algorithm and the proofs. We prove a bound $O(n \log n)$ on the number of cutting plane iterations, whereas no explicit bound is given in [4].

1.3 Cut selection via dual values.

Uniqueness of optimum LP solutions does not suffice to maintain half-integrality of optimal solutions upon adding any sequence of blossom inequalities. At any iteration, inequalities that are tight for the current optimal solution are natural candidates to be retained in the next iteration while the new inequalities are determined by odd cycles in the support of the current optimal solution. However, it turns out that keeping all tight inequalities does not maintain half-integrality. In fact, as shown in Fig. 1, even a laminar family of blossom inequalities is insufficient to guarantee the nice structural property on the intermediate LP optimum. Thus the new cuts have to be added carefully and it is also crucial that we choose carefully which older cuts to retain.

Our main algorithmic insight is that the choice of cuts for the next iteration can be determined by examining optimal dual solutions to the current LP — we retain those cuts whose dual values are strictly positive. Since there could be multiple dual optimal solutions, we use a restricted type of dual optimal solution (called *positively-critical dual* in this paper) that can be computed either by solving a single LP of the same size or via a combinatorial subroutine. We ensure that the set of cuts imposed in any LP are laminar and correspond to blossom inequalities. We remark that eliminating cutting planes that have zero dual values is common in implementations of the cutting plane algorithm.

1.4 Algorithm C-P-Matching.

All graphs in the paper will be undirected. For a graph $G = (V, E)$, and a subset $S \subseteq V$, let $\delta(S)$ denote the set of edges in E with exactly one endpoint in S , and $E[S]$ the set of edges in E with both endpoints inside E . For a node $u \in V$, $\delta(u)$ will be used to denote $\delta(\{u\})$, the set of edges incident to u . For a vector $x : E \rightarrow \mathbb{R}_+$, $\text{supp}(x)$ will denote its support, i.e., the set of edges $e \in E$ with $x(e) > 0$.

We now describe our cutting plane algorithm. Let $G = (V, E)$ be a graph, $c : E \rightarrow \mathbb{R}$ a cost function on the edges, and assume G has a perfect matching. The *bipartite relaxation* polytope and its dual are specified as follows.

$$\begin{array}{ll}
 \min \sum_{uv \in E} c(uv)x(uv) & \max \sum_{u \in V} \pi(u) \\
 x(\delta(u)) = 1 \quad \forall u \in V & \pi(u) + \pi(v) \leq c(uv) \quad \forall uv \in E \\
 x \geq 0 &
 \end{array}
 \quad (P_0(G, c)) \qquad (D_0(G, c))$$

We call a vector $x \in \mathbb{R}^E$ *proper-half-integral* if $x(e) \in \{0, 1/2, 1\}$ for every $e \in E$, and its support $\text{supp}(x)$ is a disjoint union of edges and odd cycles. The bipartite relaxation of any graph has the following well-known property.

Proposition 1.2. [Thm 30.2 in [26]] *Every basic feasible solution x of $P_0(G, c)$ is proper-half-integral.* □

Let \mathcal{O} be the set of all odd subsets of V of size at least 3, and let \mathcal{V} denote the set of one element subsets of V . For a family of odd sets $\mathcal{F} \subseteq \mathcal{O}$, consider the following pair of linear programs.

$$\begin{array}{ll}
\min \sum_{uv \in E} c(uv)x(uv) & (P_{\mathcal{F}}(G, c)) \\
x(\delta(u)) = 1 \quad \forall u \in \mathcal{V} & \\
x(\delta(S)) \geq 1 \quad \forall S \in \mathcal{F} & \\
x \geq 0 &
\end{array}
\qquad
\begin{array}{ll}
\max \sum_{S \in \mathcal{V} \cup \mathcal{F}} \Pi(S) & (D_{\mathcal{F}}(G, c)) \\
\sum_{S \in \mathcal{V} \cup \mathcal{F}: uv \in \delta(S)} \Pi(S) \leq c(uv) \quad \forall uv \in E & \\
\Pi(S) \geq 0 \quad \forall S \in \mathcal{F} &
\end{array}$$

Note that $P_{\emptyset}(G, c)$ is identical to $P_0(G, c)$, whereas $P_{\mathcal{O}}(G, c)$ is identical to (P). Every intermediate LP in our cutting plane algorithm will be $P_{\mathcal{F}}(G, c)$ for some laminar family \mathcal{F} . We will use $\Pi(v)$ to denote $\Pi(\{v\})$ for dual solutions.

Assume we are given a dual feasible solution Γ to $D_{\mathcal{F}}(G, c)$. We say that a dual optimal solution Π to $D_{\mathcal{F}}(G, c)$ is Γ -*extremal*, if it minimizes

$$h(\Pi, \Gamma) = \sum_{S \in \mathcal{V} \cup \mathcal{F}} \frac{|\Pi(S) - \Gamma(S)|}{|S|}$$

among all dual optimal solutions Π . A Γ -extremal dual optimal solution can be found by solving a single LP if we are provided with the primal optimal solution to $P_{\mathcal{F}}(G, c)$ (see Section 5.3).

Algorithm 2 gives our proposed cutting plane implementation. From the current set of cuts, we retain only those which have a positive value in an extremal dual optimal solution; let \mathcal{H}' denote this set of cuts. The new set of cuts \mathcal{H}'' correspond to odd cycles in the support of the current solution. However, in order to maintain laminarity of the cut family, we do not add the vertex sets of these cycles but instead their union with all the sets in \mathcal{H}' that they intersect. We will show that these unions are also odd sets and thus give blossom inequalities. It will follow from our analysis that each set in \mathcal{H}' intersects at most one new odd cycle, so the sets added in \mathcal{H}'' are disjoint. This step is illustrated in Fig. 2. (In the first iteration, there is no need to solve the dual LP as \mathcal{F} will be empty.)




Figure 2: Adding laminar sets for new cycles.

1.5 Overview of the analysis.

The aim of our analysis is two-fold: to show that the algorithm maintains half-integrality and that it converges quickly.

Algorithm 2 Algorithm C-P-Matching

1. Let c be the cost function on the edges after perturbation (i.e., after ordering the edges arbitrarily and increasing the cost of edge i by $1/2^i$).
 2. **Initialization.** $\mathcal{F} = \emptyset$, $\Gamma \equiv 0$.
 3. **Repeat** until x is integral:
 - (a) **Solve LP.** Find an optimal solution x to $P_{\mathcal{F}}(G, c)$.
 - (b) **Choose old cutting planes.** Find a Γ -extremal dual optimal solution Π to $D_{\mathcal{F}}(G, c)$.
Let
$$\mathcal{H}' = \{S \in \mathcal{F} : \Pi(S) > 0\}.$$
 - (c) **Choose new cutting planes.** Let \mathcal{C} denote the set of odd cycles in $\text{supp}(x)$. For each $C \in \mathcal{C}$, define \hat{C} as the union of $V(C)$ and the maximal sets of \mathcal{H}' intersecting it. Let
$$\mathcal{H}'' = \{\hat{C} : C \in \mathcal{C}\}.$$
 - (d) Set the next $\mathcal{F} = \mathcal{H}' \cup \mathcal{H}''$ and $\Gamma = \Pi$.
 4. **Return** the minimum-cost perfect matching x .
-

Our analysis to show half-integrality is based on extending the notion of factor-criticality. Recall that a graph is *factor-critical* if deleting any node leaves the graph with a perfect matching. Factor-critical graphs play an important role in matching algorithms (for more background, we refer to the books [19] and [26]). As an important example, the sets contracted during the course of Edmonds' matching algorithms (both unweighted [11] and weighted [10]) are factor-critical subgraphs. We define a notion of factor-criticality for weighted graphs that also takes a laminar odd family \mathcal{F} into account.

To prove convergence, we use the number of odd cycles in the support of an optimal half-integral solution as a potential function. We first show $\text{odd}(x_{i+1}) \leq \text{odd}(x_i)$, where x_i, x_{i+1} are consecutive LP optimal solutions, and $\text{odd}(x)$ is the number of odd cycles in the support of x . We further show that the cuts added in iterations where $\text{odd}(x_i)$ does not decrease continue to be retained until $\text{odd}(x_i)$ decreases. Using the fact that the maximum size of a laminar family of nontrivial odd sets is $n/2$, we bound the number of iterations by $O(n \log n)$ (this idea is elaborated in the proof of Theorem 1.1, see Section 3).

The analysis of the potential function behavior is quite intricate. It proceeds by designing a *half-integral* combinatorial procedure for minimum-cost perfect matching, and arguing that the optimal solution to the extremal dual LP must correspond to the one found by this procedure. We emphasize that this procedure is used only in the analysis. (It could also be used as a combinatorial method to solve the intermediate LPs in place of a black-box LP solver.) A complete, stand-alone extension of the half-integral combinatorial procedure to obtain min-cost perfect matchings is given in [5].


2 Factor-critical sets.

In this section, we define factor-critical sets and factor-critical duals.

Let $H = (V, E)$ be a graph and $\mathcal{F} \subseteq \mathcal{O}$ be a laminar family of odd subsets of V . We say that an edge set $M \subseteq E$ is an \mathcal{F} -*matching*, if it is a matching, and for any $S \in \mathcal{F}$, $|M \cap \delta(S)| \leq 1$. For a set $U \subseteq V$, we call a set M of edges to be an (U, \mathcal{F}) -*perfect-matching*, if it is an \mathcal{F} -matching covering precisely the vertex set U .

A set $S \in \mathcal{F}$ is defined to be (H, \mathcal{F}) -*factor-critical* or \mathcal{F} -*factor-critical* in the graph H , if for every node $u \in S$, there exists an $(S \setminus \{u\}, \mathcal{F})$ -perfect-matching using the edges of H . For a laminar family \mathcal{F} and a feasible solution Π to $D_{\mathcal{F}}(G, c)$, let $G_{\Pi} = (V, E_{\Pi})$ denote the graph of *tight* edges, i.e., edges $e \in E$ for which the corresponding inequality in $D_{\mathcal{F}}(G, c)$ is satisfied as an equality by the solution Π . For simplicity we will say that a set $S \in \mathcal{F}$ is (Π, \mathcal{F}) -factor-critical if it is (G_{Π}, \mathcal{F}) -factor-critical, i.e., S is \mathcal{F} -factor-critical in G_{Π} . For a vertex $u \in S$, the corresponding matching M_u is called the Π -*critical-matching* for u . (If there are multiple such matchings, select M_u arbitrarily.) If \mathcal{F} is clear from the context, then we simply say S is Π -factor-critical.

Fig. 3 gives an example of an (H, \mathcal{F}) factor-critical set, with all tight edges indicated. The three sets in \mathcal{F} besides S are indicated with circles. For any vertex $u \in S$, deleting u leaves an $(S \setminus \{u\}, \mathcal{F})$ -perfect matching. However, when the edge e is removed, the graph on the vertex set S remains factor-critical, but the only perfect matching for $S \setminus \{u\}$ has two edges crossing T , a set in \mathcal{F} , and therefore S is not (H, \mathcal{F}) -factor-critical.



An $(S \setminus \{u\}, \mathcal{F})$ p.m.; S is (H, \mathcal{F}) -factor-critical. After deleting e , S is not (H, \mathcal{F}) -factor-critical even though S is factor-critical.

Figure 3: (H, \mathcal{F}) -factor-critical vs factor-critical

A feasible solution Π to $D_{\mathcal{F}}(G, c)$ is an \mathcal{F} -*critical dual*, if every $S \in \mathcal{F}$ is (Π, \mathcal{F}) -factor-critical, and $\Pi(T) > 0$ for every non-maximal set T in \mathcal{F} . A family $\mathcal{F} \subseteq \mathcal{O}$ is called a *critical family*, if \mathcal{F} is laminar, and there exists an \mathcal{F} -critical dual solution. This will be a significant notion: the set of cuts imposed in every iteration of the cutting plane algorithm will be a critical family. The following observation provides some context and motivation for these definitions.

Proposition 2.1. *Let \mathcal{F} be the set of contracted sets at some stage of Edmonds' matching algorithm. Then the corresponding dual solution Π in the algorithm is an \mathcal{F} -critical dual. \square*

We call Π to be an \mathcal{F} -*positively-critical dual*, if Π is a feasible solution to $D_{\mathcal{F}}(G, c)$, and every $S \in \mathcal{F}$ such that $\Pi(S) > 0$ is (Π, \mathcal{F}) -factor-critical. Clearly, every \mathcal{F} -critical dual is also an \mathcal{F} -positively-critical dual, but the converse is not true. The extremal dual optimal solutions found in every iteration of Algorithm C-P-Matching will be \mathcal{F} -positively-critical, where \mathcal{F} is the family of blossom inequalities imposed in that iteration.

The next lemma summarizes elementary properties of Π -critical matchings.

Lemma 2.2. *Let \mathcal{F} be a laminar odd family, Π be a feasible solution to $D_{\mathcal{F}}(G, c)$, and $S \in \mathcal{F}$ be a (Π, \mathcal{F}) -factor-critical set. For $u, v \in S$, let M_u, M_v be the Π -critical-matchings for u, v respectively.*

(i) *For every $T \in \mathcal{F}$ such that $T \subsetneq S$,*

$$|M_u \cap \delta(T)| = \begin{cases} 1 & \text{if } u \in S \setminus T, \\ 0 & \text{if } u \in T. \end{cases}$$

(ii) *Assume the symmetric difference of M_u and M_v contains a cycle C . Then $M_u \Delta C$ is also a Π -critical matching for u .*

Proof. Proof. (i) M_u is a perfect matching of $S \setminus \{u\}$, hence for every $T \subsetneq S$,

$$|M_u \cap \delta(T)| \equiv |T \setminus \{u\}| \pmod{2}.$$

By definition of M_u , $|M_u \cap \delta(T)| \leq 1$ for any $T \subsetneq S$, $T \in \mathcal{F}$, implying the claim.

(ii) Let $M' = M_u \Delta C$. First observe that $u, v \notin V(C)$. Hence M' is a perfect matching on $S \setminus \{u\}$ using only tight edges w.r.t. Π . It remains to show that $|M' \cap \delta(T)| \leq 1$ for every $T \in \mathcal{F}$, $T \subsetneq S$. Let γ_u and γ_v denote the number of edges in $C \cap \delta(T)$ belonging to M_u and M_v , respectively. Since these are critical matchings, we have $\gamma_u, \gamma_v \leq 1$. On the other hand, since C is a cycle, $|C \cap \delta(T)|$ is even and hence $\gamma_u + \gamma_v = |C \cap \delta(T)|$ is even. These imply that $\gamma_u = \gamma_v$. The claim follows since $|M' \cap \delta(T)| = |M_u \cap \delta(T)| - \gamma_u + \gamma_v$. \square \square

The following corollary shows that (Π, \mathcal{F}) -factor-critical property of a set implies that all sets contained inside it are also (Π, \mathcal{F}) -factor-critical.

Corollary 2.3. *Let \mathcal{F} be a laminar family, Π be a feasible solution to $D_{\mathcal{F}}(G, c)$, and $S \in \mathcal{F}$ be a (Π, \mathcal{F}) -factor-critical set. Then, every set $T \subseteq S$, $T \in \mathcal{F}$ is also (Π, \mathcal{F}) -factor-critical.*

Proof. Proof. By definition, for each vertex $u \in S$, we have a matching M_u supported on the tight edges of Π such that (1) M_u is a perfect matching on $S \setminus \{u\}$ and (2) $|M_u \cap \delta(U)| \leq 1$ for all sets $U \subseteq S$, $U \in \mathcal{F}$.

Now, for any vertex $u \in T$, take $N_u = M_u \cap E[T]$. By Lemma 2.2, we have that $|M_u \cap \delta(T)| = 0$ and hence N_u is a perfect matching on $T \setminus \{u\}$. Further, for each set $U \subseteq T$, $U \in \mathcal{F}$, we have that $|N_u \cap \delta(U)| \leq |M_u \cap \delta(U)| \leq 1$. Thus, N_u is the required Π -critical-matching. \square \square

The next simple claim shows that being a critical family is a downwards monotone property.

Claim 2.4. *Let \mathcal{F} be a critical family, and $\mathcal{H} \subseteq \mathcal{F}$ a downwards closed subfamily, i.e., if $S, T \in \mathcal{F}$, $S \subseteq T$ and $T \in \mathcal{H}$, then $S \in \mathcal{H}$. Then \mathcal{H} is also a critical family.*

Proof. Proof. Let Π be a dual solution proving that \mathcal{F} is a critical family. Consider a downward closed laminar family $\mathcal{H} \subset \mathcal{F}$, and let Π' be obtained from Π by setting $\Pi'(T) = 0$ for $T \in \mathcal{F} \setminus \mathcal{H}$ and $\Pi' = \Pi$ for all other sets. Since we only lower Π values, Π' remains feasible. Take any set $S \in \mathcal{H}$. Then we need to show that for any $u \in S$, there is an $(S \setminus \{u\}, \mathcal{H})$ -perfect matching. But the tight edges of Π and Π' inside S are the same, since any such edge does not cross any set T for which $\Pi(T) \neq \Pi'(T)$. Thus, the $(S \setminus \{u\}, \mathcal{F})$ -perfect matching is also an $(S \setminus \{u\}, \mathcal{H})$ -perfect matching. \square \square

The following *uniqueness* property is used to guarantee the existence of a proper-half-integral solution in each step. We require that the cost function $c : E \rightarrow \mathbb{R}$ satisfies:

$$\text{For every critical family } \mathcal{F}, P_{\mathcal{F}}(G, c) \text{ has a unique optimal solution.} \quad (\star)$$

The next lemma shows that an arbitrary integer cost function can be perturbed to satisfy this property. The proof of is presented in Section 7.

Lemma 2.5. *Let $c : E \rightarrow \mathbb{Z}$ be an integer cost function, and \tilde{c} be its perturbation (order the edges arbitrarily, and increase the cost of edge i by $1/2^i$). Then \tilde{c} satisfies the uniqueness property (\star) .*

3 Outline of analysis and proof of the main theorem.

The proof of our main theorem is established via four main lemmas, stated below. Lemma 3.1 (proved in Section 4) and Lemma 3.2 (Section 5) establishing the half-integrality of intermediate solutions; Lemma 3.3 showing that the number of odd cycles in the support of intermediate solutions is nondecreasing, and finally Lemma 3.4 measuring progress in a sequence of steps when the number of odd cycles stays fixed (the latter two lemmas are proved in Section 6). The proof of Theorem 1.1 follows easily from these lemmas and is given here.

Lemma 3.1. *Let \mathcal{F} be a laminar odd family and assume $P_{\mathcal{F}}(G, c)$ has a unique optimal solution x . If $D_{\mathcal{F}}(G, c)$ has an \mathcal{F} -positively-critical dual optimal solution, then x is proper-half-integral.*

Lemma 3.1 is shown using a basic contraction operation. Let Π be an \mathcal{F} -positively-critical dual optimal solution for the laminar odd family \mathcal{F} . Then contracting every set $S \in \mathcal{F}$ with $\Pi(S) > 0$ preserves primal and dual optimal solutions for the contracted graph and corresponding primal and dual LPs. This is shown in Lemma 4.1. Moreover, for a unique primal optimal solution x to $P_{\mathcal{F}}(G, c)$, its image x' in the contracted graph is the unique optimal solution; if x' is proper-half-integral, then so is x . Lemma 3.1 then follows: we contract all maximal sets $S \in \mathcal{F}$ with $\Pi(S) > 0$. The image x' of the unique optimal solution x is the unique optimal solution to the bipartite relaxation in the contracted graph, and consequently, half-integral.

Such \mathcal{F} -positively-critical dual optimal solutions are hence quite helpful, but their existence is far from obvious. We next show that if \mathcal{F} is a critical family, then the extremal dual optimal solutions found by the algorithm are in fact \mathcal{F} -positively-critical dual optimal solutions.

Lemma 3.2. *Suppose that in an iteration of Algorithm C-P-Matching, \mathcal{F} is a critical family with Γ being an \mathcal{F} -critical feasible solution to $D_{\mathcal{F}}(G, c)$. Then a Γ -extremal dual optimal solution Π is an \mathcal{F} -positively-critical optimal solution to $D_{\mathcal{F}}(G, c)$. Moreover, the next set of cuts $\mathcal{H} = \mathcal{H}' \cup \mathcal{H}''$ is a critical family with Π being an \mathcal{H} -critical dual to $D_{\mathcal{H}}(G, c)$.*

Our goal then is to show that a critical family \mathcal{F} always admits an \mathcal{F} -positively-critical dual optimum; and that every extremal dual solution satisfies this property. We need a deeper understanding of the structure of dual optimal solutions. Section 5 is dedicated to this analysis. Let Γ be an \mathcal{F} -critical dual solution, and Π be an arbitrary dual optimal solution to $D_{\mathcal{F}}(G, c)$. Lemma 5.1 shows the following relation between Π and Γ inside sets $S \in \mathcal{F}$ that are tight for a primal optimal solution x : Let $\Gamma_S(u)$ and $\Pi_S(u)$ denote the sum of the dual values of sets containing u that are strictly contained inside S in solutions Γ and Π respectively. Then, every edge in $\text{supp}(x) \cap \delta(S)$ is incident to some node $u \in S$ that maximizes $\Gamma_S(u) - \Pi_S(u)$. Also, if $S \in \mathcal{F}$ is both Γ - and Π -factor-critical, then Γ and Π are identical inside S (Lemma 5.8).

If $\Pi(S) > 0$ but S is not Π -factor-critical, the above property (called *consistency* later) enables us to modify Π by moving towards Γ inside S , and decreasing $\Pi(S)$ so that optimality is

maintained. Thus, we either get that Π and Γ are identical inside S thereby making S to be Π -factor-critical or $\Pi(S) = 0$. A sequence of such operations converts an arbitrary dual optimal solution to an \mathcal{F} -positively-critical dual optimal one, leading to a combinatorial procedure to obtain positively-critical dual optimal solutions (Section 5.2). Moreover, such operations decrease the secondary objective value $h(\Pi, \Gamma)$ and thus show that every Γ -extremal dual optimum is also an \mathcal{F} -positively-critical dual optimum.

Lemmas 3.1 and 3.2 together guarantee that the unique primal optimal solutions obtained during the execution of the algorithm are proper-half-integral. In the second part of the proof of Theorem 1.1, we show convergence by considering the number of odd cycles, $\text{odd}(x)$, in the support of the current primal optimal solution x .

Lemma 3.3. *Assume the cost function c satisfies (\star) . Then $\text{odd}(x)$ is non-increasing during the execution of Algorithm C-P-Matching.*

We observe that similar to Lemma 3.1, the above Lemma 3.3 is also true if we choose an arbitrary \mathcal{F} -positively-critical dual optimal solution Π in each iteration of the algorithm. To show that the number of cycles has to strictly decrease within a polynomial number of iterations, we need the more specific choice of extremal duals.

Lemma 3.4. *Assume the cost function c satisfies (\star) and that $\text{odd}(x)$ does not decrease between iterations i and j , for some $i < j$. Let \mathcal{F}_k be the set of blossom inequalities imposed in the k 'th iteration and $\mathcal{H}_k'' = \mathcal{F}_k \setminus \mathcal{F}_{k-1}$ be the subset of new inequalities in this iteration. Then,*

$$\bigcup_{k=i+1}^j \mathcal{H}_k'' \subseteq \mathcal{F}_{j+1}.$$

We prove this progress by coupling intermediate primal and dual solutions with the solutions of a *Half-integral Matching* procedure that we design for this purpose. This procedure is a variation of Edmonds' primal-dual weighted matching algorithm and reveals the structure of the intermediate LP solutions. An extension of this procedure as described in [5] leads to an algorithm for finding min-cost integral perfect matching. Unlike Edmonds' algorithm, which maintains an integral matching and extends the matching to cover all vertices, the algorithm described in [5] maintains a proper-half-integral perfect matching.

The main theorem can be proved using the above lemmas.

Proof of Theorem 1.1. We use Algorithm C-P-Matching (Algorithm 2) for the perturbed cost function. By Lemma 2.5, this satisfies (\star) . Let i denote the index of the iteration. We prove by induction on i that every intermediate solution x_i is proper-half-integral and thus (i) follows immediately by the choice of the algorithm. The proper-half-integral property holds for the initial solution x_0 by Proposition 1.2. The induction step follows by Lemmas 3.1 and 3.2 and the uniqueness property. Further, by Lemma 3.3, the number of odd cycles in the support does not increase.

Assume the number of cycles in the i 'th phase is ℓ , and we have the same number of odd cycles ℓ in a later iteration j . For $i \leq k \leq j$, the set \mathcal{H}_k'' always contains ℓ cuts (since $|\mathcal{H}_k''|$ is equal to the number of odd cycles in x_k), and thus the number of cuts added is at least $\ell(j-i)$. By Lemma 3.4, all cuts in $\bigcup_{k=i+1}^j \mathcal{H}_k''$ are imposed in the family \mathcal{F}_{j+1} . Since \mathcal{F}_{j+1} is a laminar odd family, it can contain at most $n/2$ subsets, and therefore $j-i \leq n/2\ell$. Consequently, the number of cycles must decrease from ℓ to $\ell-1$ within $n/2\ell$ iterations. Since $\text{odd}(x_0) \leq n/3$, the number of iterations is at most $O(n \log n)$.

Finally, we show that optimal solution returned by the algorithm using \tilde{c} is also optimal for the original cost function. Let M be the optimal matching returned by \tilde{c} , and assume for a

contradiction that there exists a different perfect matching M' with $c(M') < c(M)$. Since c is integral, it means $c(M') \leq c(M) - 1$. In the perturbation, since $c(e) < \tilde{c}(e)$ for every $e \in E$, we have $c(M) < \tilde{c}(M)$, and since $\sum_{e \in E} (\tilde{c}(e) - c(e)) < 1$, we have $\tilde{c}(M') < c(M') + 1$. This gives $\tilde{c}(M') < c(M') + 1 \leq c(M) < \tilde{c}(M)$, a contradiction to the optimality of M for \tilde{c} . \square

4 Contractions and half-integrality.

We define an important contraction operation and derive some fundamental properties. Let \mathcal{F} be a laminar odd family, let Π be a feasible solution to $D_{\mathcal{F}}(G, c)$, and let $S \in \mathcal{F}$ be a (Π, \mathcal{F}) -factor-critical set. Let us define

$$\Pi_S(u) := \sum_{T \in \mathcal{V} \cup \mathcal{F}: T \subseteq S, u \in T} \Pi(T)$$

to be the total dual contribution of sets inside S containing u .

By contracting S w.r.t. Π , we mean the following: Let $G' = (V', E')$ be the contracted graph on node set $V' = (V \setminus S) \cup \{s\}$, s representing the contraction of S . For each $u \in V$, we denote u' to be the image of u , i.e., if $u \in S$, then $u' = s$, otherwise $u' = u$. Let \mathcal{V}' denote the set of one-element subsets of V' . For a set $T \subseteq V$, let T' denote its contracted image. Let \mathcal{F}' be the set of nonsingular images of the sets of \mathcal{F} , that is, $T' \in \mathcal{F}'$ if $T \in \mathcal{F}$, and $T \setminus S \neq \emptyset$. Let E' contain all edges $uv \in E$ with $u, v \notin S$ and for every edge uv with $u \in S, v \in V \setminus S$ add an edge $u'v$. If this creates parallel edges, keep only one from each bundle. Let us define the image Π' of Π to be $\Pi'(T') = \Pi(T)$ for every $T' \in \mathcal{V}' \cup \mathcal{F}'$. For every edge $u'v' \in E'$, let $x'(u'v')$ be equal to the sum of x values of the pre-images of $u'v'$. (That is, if $u, v \notin S$, then $x'(uv) = x(uv)$, and let $x'(sv) = \sum_{u \in S} x(uv)$.) Define the new edge costs by

$$\begin{aligned} c'(uv) &= c(uv), & \text{for } u, v \in V \setminus S, \\ c'(sv) &= \min_{u \in S} c(uv) - \Pi_S(u), & \text{for } v \in V \setminus S. \end{aligned}$$

We refer the reader to Figure 4 for an example of the contraction operation.

Lemma 4.1. *Let \mathcal{F} be a laminar odd family, let x be an optimal solution to $P_{\mathcal{F}}(G, c)$, and let Π be a feasible solution to $D_{\mathcal{F}}(G, c)$. Let $S \in \mathcal{F}$ be a (Π, \mathcal{F}) -factor-critical set and let G', c', \mathcal{F}' denote the graph, costs and laminar family respectively obtained by contracting S w.r.t. Π ; let x', Π' be the images of x, Π , respectively. Then the following hold.*

- (i) Π' is a feasible solution to $D_{\mathcal{F}'}(G', c')$. Furthermore, if a set $T \in \mathcal{F}$, $T \setminus S \neq \emptyset$ is (Π, \mathcal{F}) -factor-critical, then its image T' is (Π', \mathcal{F}') -factor-critical.
- (ii) Suppose Π is an optimal solution to $D_{\mathcal{F}}(G, c)$ and $x(\delta(S)) = 1$. Then x' is an optimal solution to $P_{\mathcal{F}'}(G', c')$ and Π' is optimal to $D_{\mathcal{F}'}(G', c')$.
- (iii) Suppose x is the unique optimal solution to $P_{\mathcal{F}}(G, c)$, and Π is an optimal solution to $D_{\mathcal{F}}(G, c)$. Then x' is the unique optimal solution to $P_{\mathcal{F}'}(G', c')$. Moreover, x' is proper-half-integral if and only if x is proper-half-integral. If x' is proper-half-integral, then $\text{odd}(x) = \text{odd}(x')$ and $\text{supp}(x) \cap E[S]$ consists of a disjoint union of edges and an even path. Further, assume C' is an odd cycle in $\text{supp}(x')$ and let T be the pre-image of $V(C')$ in G . Then, $\text{supp}(x) \cap E[T]$ consists of an odd cycle and matching edges.

Proof. (i) For feasibility, it is sufficient to verify

$$\sum_{T' \in \mathcal{V}' \cup \mathcal{F}': u'v' \in \delta(T')} \Pi'(T') \leq c'(u'v') \quad \forall u'v' \in E'.$$




Figure 4: Contraction operation: image of dual and new cost function

If $u, v \neq s$, this is immediate from feasibility of Π to $D_{\mathcal{F}}(G, c)$. Consider an edge $sv' \in E(G')$. Let uv be the pre-image of this edge that gives $c'(s'v) = c(uv) - \Pi_S(u)$.

$$\sum_{T' \in \mathcal{V} \cup \mathcal{F}': sv' \in \delta(T')} \Pi'(T') = \Pi(S) + \sum_{T \in \mathcal{F}: uv \in \delta(T), T \setminus S \neq \emptyset} \Pi(T) \leq c(uv) - \Pi_S(u) = c'(sv').$$

We also observe that $u'v'$ is tight in G' w.r.t Π' if and only if the pre-image uv is tight in G w.r.t Π .

Let $T \in \mathcal{F}$ be a (Π, \mathcal{F}) -factor-critical set with $T \setminus S \neq \emptyset$. It is sufficient to verify that T' is (Π', \mathcal{F}') -factor-critical whenever T contains S . Let $u' \in T'$ be the image of $u \in V$, and consider the image M' of the Π -critical-matching M_u . Every edge in M' is tight with respect to Π' . Further, M' is a matching, since $|M_u \cap \delta(S)| \leq 1$ must hold. Let $Z' \subsetneq T'$, $Z' \in \mathcal{F}'$ and let Z be the pre-image of Z' . If $u' \neq s$, then $|M' \cap \delta(Z')| = |M_u \cap \delta(Z)| \leq 1$ and since $|M_u \cap \delta(S)| = 1$ by Lemma 2.2, the matching M' is a $(T' \setminus \{u'\}, \mathcal{F}')$ -perfect-matching. If $u' = s$, then $u \in S$. By Lemma 2.2, $M_u \cap \delta(S) = \emptyset$ and hence, M' misses s . Also, $|M_u \cap \delta(Z)| \leq 1$ implies $|M' \cap \delta(Z')| \leq 1$ and hence M' is a $(T' \setminus \{s\}, \mathcal{F}')$ -perfect-matching.

(ii) Since $x(\delta(S)) = 1$, we have $x'(\delta(v)) = 1$ for every $v \in V'$. It is straightforward to verify that $x'(\delta(T')) \geq 1$ for every $T' \in \mathcal{F}'$ with equality if $x(\delta(T)) = 1$. Thus, x' is feasible to $P_{\mathcal{F}'}(G', c')$. Optimality follows as x' and Π' satisfy complementary slackness, using that the image of tight edges is tight, as shown by the argument for part (i).

(iii) For uniqueness, consider an arbitrary optimal solution \tilde{y} to $P_{\mathcal{F}'}(G', c')$. We partition the edge set $\delta_{E'}(s)$ into subsets $J(u)$ for $u \in S$ as follows. For every $sv \in \delta_{E'}(s)$, pick the node $u \in S$ that minimizes $c(uv) - \Pi_S(u)$ (that is, $c(sv) = c(uv) - \Pi_S(u)$), and let $sv \in J(u)$; if there are multiple such nodes, assign one arbitrarily. Define $\alpha_u := \tilde{y}(J(u))$. It is straightforward by the

contraction that $\sum_{u \in S} \alpha_u = \tilde{y}(\delta(S)) = 1$. Let M_u be the Π -critical matching for u in S . Take $w := \sum_{u \in S} \alpha_u \chi(M_u)$, where $\chi(M_u)$ is the indicator vector of M_u and let

$$y(uv) := \begin{cases} \tilde{y}(uv) & \text{if } u, v \in V \setminus S, \\ \tilde{y}(sv) & \text{if } u \in S, v \in V \setminus S, sv \in J(u), \\ w(uv) & \text{if } u, v \in S, \\ 0 & \text{otherwise,} \end{cases}$$

Now it is easy to verify that y is a feasible solution to $P_{\mathcal{F}}(G, c)$, and its image in G' is $y' = \tilde{y}$. Moreover, since \tilde{y} was chosen as an optimal solution, it must satisfy complementary slackness with Π' , which is a dual optimal solution for $D_{\mathcal{F}}(G', c')$. We now verify complementary slackness between y and Π .

If $y(uv) > 0$ for an edge uv , then we have one of three possibilities: (i) $u, v \in S$, in which case, uv is a tight edge wrt Π by definition. (ii) $u, v \in V \setminus S$ and $\tilde{y}(uv) > 0$ in which case, uv is a tight edge wrt Π' which implies that uv is a tight edge wrt Π . (iii) $u \in S, v \in V \setminus S, sv \in J(u)$, and $\tilde{y}(sv) > 0$ in which case, sv is a tight edge wrt Π' . Therefore

$$\Pi(S) + \sum_{T \in \mathcal{F}, v \in T, T \setminus S \neq \emptyset} \Pi(T) = c'(sv) = c(uv) - \Pi_S(u).$$

The last equality is because $sv \in J(u)$. The above equation implies that uv is a tight edge w.r.t. Π .

Moreover, if $\Pi(T) > 0$ for $T \in \mathcal{F}, T \setminus S \neq \emptyset$, then it is easy to verify that $y(\delta(T)) = \tilde{y}(\delta(T)) = 1$. It remains to check that subsets $T \in \mathcal{F}, T \subsetneq S$ with $\Pi(T) > 0$ satisfy $y(\delta(T)) = 1$. For this, we note that

$$y(\delta(T)) = \sum_{u \in T, v \in V \setminus S, sv \in J(u)} \tilde{y}(sv) + \sum_{u \in T, v \in S \setminus T} w(uv).$$

Since w is a weighted sum of $(S \setminus \{u\}, \mathcal{F})$ -perfect matchings for $u \in S$, we have

$$\sum_{u \in T, v \in S \setminus T} w(uv) = \sum_{u \in S} \alpha_u |M_u \cap \delta(T)| = \sum_{u \in S \setminus T} \alpha_u.$$

The last equality above is by Lemma 2.2 (i). Moreover, by definition of α_u , we have

$$\sum_{u \in T, v \in V \setminus S, sv \in J(u)} \tilde{y}(sv) = \sum_{u \in T} \tilde{y}(J(u)) = \sum_{u \in T} \alpha_u.$$

Using the above equalities, we have that $y(\delta(T)) = \sum_{u \in S} \alpha_u = 1$. Hence, by complementary slackness, y is an optimal solution to $P_{\mathcal{F}}(G, c)$ and by uniqueness, $y = x$. This implies that $x' = y' = \tilde{y}$, proving the uniqueness of optimal solutions in $P_{\mathcal{F}'}(G', c')$.

The above argument also shows that x must be identical to w inside S . Suppose x' is proper-half-integral. First, assume s is covered by a matching edge in x' . Then $\alpha_u = 1$ for some $u \in S$ and $\alpha_v = 0$ for every $v \neq u$. Consequently, $w = \chi(M_u)$ is a perfect matching on $S \setminus \{u\}$. Next, assume s is incident to an odd cycle in x' . Then $\alpha_{u_1} = \alpha_{u_2} = 1/2$ for some nodes $u_1, u_2 \in S$, and $w = \frac{1}{2}(\chi(M_{u_1}) + \chi(M_{u_2}))$. We claim that M_{u_1} and M_{u_2} must both be unique. Indeed, the above argument constructed y from \tilde{y} using arbitrary choices of M_{u_1} and M_{u_2} , and concluded that $x = y$. Hence if either of them were not unique, that contradicts the uniqueness of x . Accordingly, by Lemma 2.2(ii), the symmetric difference of M_{u_1} and M_{u_2} does not contain any even cycles. Hence, $\text{supp}(w)$ contains an even path between u_1 and u_2 , and some matching edges. Consequently, x is proper-half-integral. The above argument immediately shows the following.

Claim 4.2. *Let C' be an odd (even) cycle such that $x'(e) = 1/2$ for every $e \in C'$ in $\text{supp}(x')$ and let T be the pre-image of the set $V(C')$ in G . Then, $\text{supp}(x) \cap E[T]$ consists of an odd (even) cycle C and a (possibly empty) set M of edges such that $x(e) = 1/2 \forall e \in C$ and $x(e) = 1 \forall e \in M$.* \square

Next, we prove that if x is proper-half-integral, then so is x' . It is clear that x' being the image of x is half-integral. If x' is not proper-half-integral, then $\text{supp}(x')$ contains an even $1/2$ -cycle, and thus by Claim 4.2, $\text{supp}(x)$ must also contain an even cycle, contradicting that it was proper-half-integral.

The above arguments also show $\text{supp}(x) \cap E[S]$ consists of a disjoint union of edges and an even path, $\text{odd}(x) = \text{odd}(x')$, and finally, if C' is an odd cycle in $\text{supp}(x')$, then Claim 4.2 provides the required structure for x inside T . \square

Iteratively applying the lemma from the innermost to the outermost sets in \mathcal{F} , we obtain the following corollary.

Corollary 4.3. *Assume x is the optimal solution to $P_{\mathcal{F}}(G, c)$ and there exists an \mathcal{F} -positively-critical dual optimal solution Π . Let \hat{G}, \hat{c} be the graph, and cost obtained by contracting all maximal sets $S \in \mathcal{F}$ with $\Pi(S) > 0$ w.r.t. Π , and let \hat{x} be the image of x in \hat{G} .*

- (i) \hat{x} and $\hat{\Pi}$ are the optimal solutions to the bipartite relaxation $P_0(\hat{G}, \hat{c})$ and $D_0(\hat{G}, \hat{c})$ respectively.
- (ii) If x is the unique optimal solution to $P_{\mathcal{F}}(G, c)$, then \hat{x} is the unique optimal solution to $P_0(\hat{G}, \hat{c})$. If \hat{x} is proper-half-integral, then x is also proper-half-integral.

Proof of Lemma 3.1. Let Π be an \mathcal{F} -positively-critical dual optimal solution, and let x be the unique optimal solution to $P_{\mathcal{F}}(G, c)$. Contract all maximal sets $S \in \mathcal{F}$ with $\Pi(S) > 0$, obtaining the graph \hat{G} and cost \hat{c} . Let \hat{x} be the image of x in \hat{G} . By Corollary 4.3(ii), \hat{x} is unique optimal solution to $P_0(\hat{G}, \hat{c})$. By Proposition 1.2, \hat{x} is proper-half-integral and hence by Corollary 4.3(ii), x is also proper-half-integral. \square

5 Structure of dual solutions.

In this section, we derive two properties of positively-critical dual optimal solutions: (1) an optimal solution Ψ to $D_{\mathcal{F}}(G, c)$ can be transformed into an \mathcal{F} -positively-critical dual optimal solution if \mathcal{F} is a critical family (Section 5.2) and (2) a Γ -extremal dual optimal solution to $D_{\mathcal{F}}(G, c)$ as obtained in the algorithm is also an \mathcal{F} -positively-critical dual optimal solution (Section 5.3). In Section 5.1, we first show some lemmas characterizing arbitrary dual optimal solutions.

5.1 Consistency of dual solutions.

Assume $\mathcal{F} \subseteq \mathcal{O}$ is a critical family, with Π being an \mathcal{F} -critical dual solution, and let Ψ be an arbitrary dual optimal solution to $D_{\mathcal{F}}(G, c)$. Note that optimality of Π is not assumed. Let x be an optimal solution to $P_{\mathcal{F}}(G, c)$; we do not make the uniqueness assumption (\star) in this section. We shall describe structural properties of Ψ compared to Π ; in particular, we show that if we contract a Π -factor-critical set S , the images of x and Ψ will be primal and dual optimal solutions in the contracted graph.

Consider a set $S \in \mathcal{F}$. We say that the dual solutions Π and Ψ are *identical* inside S , if $\Pi(T) = \Psi(T)$ for every set $T \subsetneq S$, $T \in \mathcal{F} \cup \mathcal{V}$. We defined $\Pi_S(u)$ in the previous section; we also

use this notation for Ψ , namely, let $\Psi_S(u) := \sum_{T \in \mathcal{V} \cup \mathcal{F}: T \subseteq S, u \in T} \Psi(T)$ for $u \in S$. Let us now define

$$\Delta_{\Pi, \Psi}(S) := \max_{u \in S} (\Pi_S(u) - \Psi_S(u)).$$

We say that Ψ is *consistent* with Π inside S , if $\Pi_S(u) - \Psi_S(u) = \Delta_{\Pi, \Psi}(S)$ holds for every $u \in S$ that is incident to an edge $uv \in \delta(S) \cap \text{supp}(x)$. The main goal of this subsection is to prove the following lemma.

Lemma 5.1. *Let $\mathcal{F} \subseteq \mathcal{O}$, Π be a feasible solution to $D_{\mathcal{F}}(G, c)$, and let $S \in \mathcal{F}$ such that S is (Π, \mathcal{F}) -factor-critical and $\Pi(T) > 0$ for every subset $T \subsetneq S, T \in \mathcal{F}$. Let Ψ be an optimal solution to $D_{\mathcal{F}}(G, c)$, and assume there exists an optimal solution x to $P_{\mathcal{F}}(G, c)$ with $x(\delta(S)) = 1$. Then Ψ is consistent with Π inside S . Further, $\Delta_{\Pi, \Psi}(S) \geq 0$ for all such sets S .*

Consistency is important as it enables us to preserve optimality when contracting a set $S \in \mathcal{F}$ w.r.t. Π . Assume Ψ is consistent with Π inside S , and $x(\delta(S)) = 1$. Let us contract S w.r.t. Π to obtain G' and c' as defined in Section 4. Define

$$\Psi'(T') = \begin{cases} \Psi(T) & \text{if } T' \in (\mathcal{F}' \cup \mathcal{V}') \setminus \{s\}, \\ \Psi(S) - \Delta_{\Pi, \Psi}(S) & \text{if } T' = \{s\} \end{cases}$$

Lemma 5.2. *Let $\mathcal{F} \subseteq \mathcal{O}$, Π be a feasible solution to $D_{\mathcal{F}}(G, c)$, and let $S \in \mathcal{F}$ such that S is (Π, \mathcal{F}) -factor-critical and $\Pi(T) > 0$ for every subset $T \subsetneq S, T \in \mathcal{F}$. Let Ψ be an optimal solution to $D_{\mathcal{F}}(G, c)$, and x be an optimal solution to $P_{\mathcal{F}}(G, c)$ with $x(\delta(S)) = 1$. Suppose that Ψ is consistent with Π inside S . Let G', c', \mathcal{F}' denote the graph, costs and laminar family obtained by contraction. Then the image x' of x is an optimal solution to $P_{\mathcal{F}'}(G', c')$, and Ψ' (as defined above) is an optimal solution to $D_{\mathcal{F}'}(G', c')$.*

Proof. Feasibility of x' follows as in the proof of Lemma 4.1(ii). For the feasibility of Ψ' , we have to verify $\sum_{T' \in \mathcal{V}' \cup \mathcal{F}': uv \in \delta(T')} \Psi'(T') \leq c'(uv)$ for every edge $uv \in E(G')$. This follows immediately whenever $u, v \neq s$, since Ψ is a feasible solution for $D_{\mathcal{F}}(G, c)$. Consider an edge $sv \in E(G')$, and pick $u \in S$ such that $uv \in \delta(S)$ and it minimizes $c(uv) - \Pi_S(u)$. Let $\Delta = \Delta_{\Pi, \Psi}(S)$.

$$\begin{aligned} c(uv) &\geq \sum_{T \in \mathcal{V} \cup \mathcal{F}: uv \in \delta(T)} \Psi(T) \\ &= \Psi_S(u) + \Psi(S) + \sum_{T \in \mathcal{F}: uv \in \delta(T), T \setminus S \neq \emptyset} \Psi(T) \\ &= \Psi_S(u) + \Delta + \sum_{T' \in \mathcal{V}' \cup \mathcal{F}': sv \in \delta(T')} \Psi'(T'). \end{aligned}$$

In the last equality, we used the definition $\Psi'(s) = \Psi(S) - \Delta$. Therefore, using $\Pi_S(u) \leq \Psi_S(u) + \Delta$, we obtain

$$\sum_{T' \in \mathcal{V}' \cup \mathcal{F}': sv \in \delta(T')} \Psi'(T') \leq c(uv) - \Psi_S(u) - \Delta \leq c(uv) - \Pi_S(u) = c'(sv). \quad (1)$$

The last equality follows by the choice of u . Thus, Ψ' is a feasible solution to $D_{\mathcal{F}'}(G', c')$. To show optimality, we verify complementary slackness for x' and Ψ' . If $x'(uv) > 0$ for $u, v \neq s$, then $x(uv) > 0$. Thus, the tightness of the constraint for uv w.r.t. Ψ' in $D_{\mathcal{F}'}(G', c')$ follows from the tightness of the constraint w.r.t. Ψ in $D_{\mathcal{F}}(G, c)$. Suppose $x'(sv) > 0$ for an edge $sv \in E(G')$. Let $uv \in E(G)$ be a pre-image of sv with $x(uv) > 0$. We claim that (1) holds with equality

throughout. The first inequality is tight since uv is tight w.r.t. Ψ , and the second is tight since $\Pi_S(u) - \Psi_S(u) = \Delta$ by the consistency property. For the last equality, we need to show that u is a node in S that minimizes $c(uv) - \Pi_S(u)$. Consider an arbitrary $z \in S$ with $zv \in \delta(S)$. Then $c(zv) - \Pi_S(z) \geq c(zv) - \Psi_S(z) - \Delta$ by the definition of Δ . Since uv is tight w.r.t. Ψ , it follows that $c(zv) - \Psi_S(z) - \Delta \geq c(uv) - \Psi_S(u) - \Delta = c(uv) - \Pi_S(u)$. Hence $c(uv) - \Pi_S(u) = c'(sv)$ follows. Finally, if $\Psi'(T') > 0$ for some $T' \in \mathcal{F}'$, then $\Psi(T) > 0$ and hence $x(\delta(T)) = 1$, implying $x'(\delta(T')) = 1$. \square

Lemma 5.3. *Let $\mathcal{F} \subseteq \mathcal{O}$, Π be a feasible solution to $D_{\mathcal{F}}(G, c)$. Let $S \in \mathcal{F}$ such that S is (Π, \mathcal{F}) -factor-critical, and $\Pi(T) > 0$ for every subset $T \subsetneq S, T \in \mathcal{F}$. Let x be an optimal solution to $P_{\mathcal{F}}(G, c)$. If $x(\delta(S)) = 1$, then all edges in $\text{supp}(x) \cap E[S]$ are tight w.r.t. Π , and $x(\delta(T)) = 1$ for every $T \subsetneq S, T \in \mathcal{F}$.*

Proof. Let $\alpha_u = x(\delta(u, V \setminus S))$ for each $u \in S$, and for each $T \subseteq S, T \in \mathcal{F}$, let $\alpha(T) = \sum_{u \in T} \alpha_u = x(\delta(T, V \setminus S))$. Note that $\alpha(S) = x(\delta(S)) = 1$. Let us consider the following pair of linear programs.

$$\begin{array}{ll} \min \sum_{uv \in E[S]} c(uv)z(uv) & (P_{\mathcal{F}}[S]) \\ z(\delta(u)) = 1 - \alpha_u \quad \forall u \in S \\ z(\delta(T)) \geq 1 - \alpha(T) \quad \forall T \subsetneq S, T \in \mathcal{F} \\ z(uv) \geq 0 \quad \forall uv \in E[S] \end{array} \quad \begin{array}{ll} \max \sum_{T \subsetneq S, T \in \mathcal{V} \cup \mathcal{F}} (1 - \alpha(T))\Gamma(T) & (D_{\mathcal{F}}[S]) \\ \sum_{\substack{T \subsetneq S, T \in \mathcal{V} \cup \mathcal{F} \\ uv \in \delta(T)}} \Gamma(T) \leq c(uv) \quad \forall uv \in E[S] \\ \Gamma(Z) \geq 0 \quad \forall Z \subsetneq T, Z \in \mathcal{F} \end{array}$$

For a feasible solution z to $P_{\mathcal{F}}[S]$, let x^z denote the solution obtained by replacing $x(uv)$ by $z(uv)$ for edges uv inside S , that is,

$$x^z(e) = \begin{cases} x(e) & \text{if } e \in \delta(S) \cup E[V \setminus S], \\ z(e) & \text{if } e \in E[S]. \end{cases}$$

Claim 5.4. *The restriction of x inside S is feasible to $P_{\mathcal{F}}[S]$, and for every feasible solution z to $P_{\mathcal{F}}[S]$, x^z is a feasible solution to $P_{\mathcal{F}}(G, c)$. Consequently, z is an optimal solution to $P_{\mathcal{F}}[S]$ if and only if x^z is an optimal solution to $P_{\mathcal{F}}(G, c)$.*

Proof. The first part is obvious. For feasibility of x^z , if $u \notin S$ then $x^z(u) = x(u) = 1$. If $u \in S$, then $x^z(u) = z(u) + x(\delta(u, V \setminus S)) = 1 - \alpha_u + \alpha_u = 1$. Similarly, if $T \in \mathcal{F}, T \setminus S \neq \emptyset$, then $x^z(\delta(T)) = x(T) \geq 1$. If $T \subseteq S$, then $x^z(\delta(T)) = z(\delta(T)) + x(\delta(T, V \setminus S)) \geq 1 - \alpha(T) + \alpha(T) = 1$.

Optimality follows since $c^T x^z = \sum_{uv \in E[S]} c(uv)z(uv) + \sum_{uv \in E \setminus E[S]} c(uv)x(uv)$. \square

Claim 5.5. *Let $\bar{\Pi}$ denote the restriction of Π inside S , that is, $\bar{\Pi}(T) = \Pi(T)$ for every $T \in \mathcal{V} \cup \mathcal{F}, T \subsetneq S$. Then $\bar{\Pi}$ is an optimal solution to $D_{\mathcal{F}}[S]$.*

Proof. Since $S \in \mathcal{F}$ is (Π, \mathcal{F}) -factor-critical, we have a Π -critical-matching M_u inside S for each $u \in S$. Let $z = \sum_{u \in S} \alpha_u \chi(M_u)$. The claim follows by showing that z is feasible to $P_{\mathcal{F}}[S]$ and that z and $\bar{\Pi}$ satisfy complementary slackness.

The degree constraint $z(\delta(u)) = 1 - \alpha_u$ is straightforward. By Lemma 2.2(i), if $T \subsetneq S, T \in \mathcal{F}$, then $z(\delta(T)) = \sum_{u \in S \setminus T} \alpha_u = 1 - \alpha(T)$. The feasibility of Π to $D_{\mathcal{F}}(G, c)$ immediately shows feasibility of $\bar{\Pi}$ to $D_{\mathcal{F}}[S]$.

Complementary slackness also follows since by definition, all M_u 's use only tight edges w.r.t. Π (equivalently, w.r.t. $\bar{\Pi}$). Also, for every odd set $T \subsetneq S, T \in \mathcal{F}$, we have that $z(\delta(T)) = 1 - \alpha(T)$ as verified above. Thus, all odd set constraints are tight in the primal. \square

By Claim 5.4, the solution obtained by restricting x to $E[S]$ must be optimal to $P_{\mathcal{F}}[S]$ and thus satisfies complementary slackness with $\bar{\Pi}$. Consequently, every edge in $E[S] \cap \text{supp}(x)$ must be tight w.r.t. $\bar{\Pi}$, and equivalently, w.r.t. Π . By the statement of the Lemma, every set $T \subsetneq S$, $T \in \mathcal{F}$ satisfies $\bar{\Pi}(T) = \Pi(T) > 0$. Thus, complementary slackness gives $x(\delta(T)) = 1$. \square

We need one more claim to prove Lemma 5.1.

Claim 5.6. *Let $S \in \mathcal{F}$ be an inclusionwise minimal set of \mathcal{F} . Let Λ and Γ be feasible solutions to $D_{\mathcal{F}}(G, c)$, and suppose S is (Λ, \mathcal{F}) -factor-critical. Then,*

$$\Delta_{\Lambda, \Gamma}(S) := \max_{u \in S} (\Lambda_S(u) - \Gamma_S(u)) = \max_{u \in S} |\Lambda_S(u) - \Gamma_S(u)|.$$

Further, if $\Delta_{\Lambda, \Gamma}(S) > 0$, define

$$\begin{aligned} A^+ &:= \{u \in S : \Gamma(u) = \Lambda(u) + \Delta_{\Lambda, \Gamma}(S)\}, \\ A^- &:= \{u \in S : \Gamma(u) = \Lambda(u) - \Delta_{\Lambda, \Gamma}(S)\}. \end{aligned}$$

Then $|A^-| > |A^+|$.

Proof. Let $\Delta = \max_{u \in S} |\Lambda_S(u) - \Gamma_S(u)|$; note that $\Delta \geq \Delta_{\Lambda, \Gamma}(S)$ by definition. If $\Delta = 0$, then $\Delta_{\Lambda, \Gamma}(S) = 0$ also follows, and thus the claim holds. In the rest of the proof, we assume $\Delta > 0$. Let us define the sets A^- and A^+ with Δ instead of $\Delta_{\Lambda, \Gamma}(S)$. Since S is (Λ, \mathcal{F}) -factor-critical, for every $a \in S$, there exists an $(S \setminus \{a\}, \mathcal{F})$ perfect matching M_a using only tight edges w.r.t. Λ , i.e., $M_a \subseteq \{uv : \Lambda(u) + \Lambda(v) = c(uv)\}$ by the minimality of S . Further, by feasibility of Γ , we have $\Gamma(u) + \Gamma(v) \leq c(uv)$ on every $uv \in M_a$. Thus, if $u \in A^+$, then $v \in A^-$ for every $uv \in M_a$. Since $\Delta > 0$, we have $A^+ \cup A^- \neq \emptyset$ and therefore $A^- \neq \emptyset$, and consequently, $\Delta = \Delta_{\Lambda, \Gamma}(S)$. Now pick $a \in A^-$ and consider M_a . This perfect matching M_a matches each node in A^+ to a node in A^- . Thus, $|A^-| > |A^+|$. \square

Proof of Lemma 5.1. We prove by induction on $|V|$, and subject to that, on $|S|$. Let us define $\Delta := \Delta_{\Pi, \Psi}(S)$. By the statement of the lemma, we have that S is (Π, \mathcal{F}) -factor-critical.

First, consider the case when S is an inclusion-wise minimal set in \mathcal{F} . Then, $\Pi_S(u) = \Pi(u)$, $\Psi_S(u) = \Psi(u)$ for every $u \in S$. By Claim 5.6, we have $\Delta \geq 0$. We are done if $\Delta = 0$. Otherwise, define the sets A^- and A^+ as in the claim using $\Delta_{\Pi, \Psi}(S)$.

Now consider an edge $uv \in E[S] \cap \text{supp}(x)$. By complementary slackness, we have $\Psi(u) + \Psi(v) = c(uv)$. By dual feasibility, we have $\Pi(u) + \Pi(v) \leq c(uv)$. Hence, if $u \in A^-$, then $v \in A^+$. Consequently, we have

$$\begin{aligned} |A^-| &= \sum_{u \in A^-} x(\delta(u)) = x(\delta(A^-, V \setminus S)) + x(\delta(A^-, A^+)) \\ &\leq 1 + \sum_{u \in A^+} x(\delta(u)) = 1 + |A^+| \leq |A^-|. \end{aligned}$$

Thus, we must have equality throughout, implying $x(\delta(A^-, V \setminus S)) = 1$. This precisely means that Ψ is consistent with Π inside S .

Next, let S be a non-minimal set. Let $T \in \mathcal{F}$ be a maximal set strictly contained in S . By Corollary 2.3, we know that T is also (Π, \mathcal{F}) -factor-critical. By Lemma 5.3, $x(\delta(T)) = 1$, therefore the inductive claim holds for T : Ψ is consistent with Π inside T , and $\Delta(T) = \Delta_{\Pi, \Psi}(T) \geq 0$.

We contract T w.r.t. Π and use Lemma 5.2. Let the image of the solutions x , Π , and Ψ be x' , Π' and Ψ' respectively and the resulting graph be G' with cost function c' . Then x' and Ψ'

are optimal solutions to $P_{\mathcal{F}'}(G', c')$ and to $D_{\mathcal{F}'}(G', c')$ respectively, and by Lemma 4.1(i), Π' is an \mathcal{F}' -critical dual. Let t be the image of T by the contraction. Now, consider the image S' of S in G' . Since G' is a smaller graph, it satisfies the induction hypothesis. Let $\Delta' = \Delta_{\Pi', \Psi'}(S')$ in G' . By induction hypothesis, $\Delta' \geq 0$. The next claim verifies consistency inside S and thus completes the proof. \square

Claim 5.7. *For every $u \in S$, $\Pi_S(u) - \Psi_S(u) \leq \Pi'_{S'}(u') - \Psi'_{S'}(u')$, and equality holds if there exists an edge $uv \in \delta(S) \cap \text{supp}(x)$. Consequently, $\Delta' = \Delta$.*

Proof. Let u' denote the image of u . If $u' \neq t$, then $\Pi'_{S'}(u') = \Pi_S(u)$, $\Psi'_{S'}(u') = \Psi_S(u)$ and therefore, $\Pi_S(u) - \Psi_S(u) = \Pi'_{S'}(u') - \Psi'_{S'}(u')$. Assume $u' = t$, that is, $u \in T$. Then $\Pi_S(u) = \Pi_T(u) + \Pi(T)$, $\Psi_S(u) = \Psi_T(u) + \Psi(T)$ by the maximal choice of T , and therefore

$$\begin{aligned} \Pi_S(u) - \Psi_S(u) &= \Pi_T(u) - \Psi_T(u) + \Pi(T) - \Psi(T) \\ &\leq \Delta(T) + \Pi(T) - \Psi(T) \\ &= \Pi'(t) - \Psi'(t) \quad (\text{Since } \Pi'(t) = \Pi(T), \Psi'(t) = \Psi(T) - \Delta(T)) \\ &= \Pi'_{S'}(t) - \Psi'_{S'}(t). \end{aligned} \tag{2}$$

Assume now that there exists a $uv \in \delta(S) \cap \text{supp}(x)$. If $u \in T$, then using the consistency inside T , we get $\Pi_T(u) - \Psi_T(u) = \Delta(T)$, and therefore (2) gives $\Pi_S(u) - \Psi_S(u) = \Pi'_{S'}(t) - \Psi'_{S'}(t) = \Delta'$. \square

Claim 5.6 can also be used to derive the following important property.

Lemma 5.8. *Given a laminar odd family $\mathcal{F} \subset \mathcal{O}$, let Λ and Γ be two dual feasible solutions to $D_{\mathcal{F}}(G, c)$. If a subset $S \in \mathcal{F}$ is both (Λ, \mathcal{F}) -factor-critical and (Γ, \mathcal{F}) -factor-critical, then Λ and Γ are identical inside S .*

Proof. Consider a graph $G = (V, E)$ with $|V|$ minimal, where the claim does not hold for some set S . Also, choose S to be the smallest counterexample in this graph. First, assume $S \in \mathcal{F}$ is a minimal set. Then consider Claim 5.6 for Λ and Γ and also by changing their roles, for Γ and Λ . If Λ and Γ are not identical inside S , then $\Delta = \max_{u \in S} |\Lambda_S(u) - \Gamma_S(u)| > 0$. The sets A^- and A^+ for Λ and Γ become A^+ and A^- for Γ and Λ . Then $|A^-| > |A^+| > |A^-|$, a contradiction.

Suppose now S contains $T \in \mathcal{F}$. It is straightforward by definition that T is also (Λ, \mathcal{F}) -factor-critical and (Γ, \mathcal{F}) -factor-critical. Thus, by the minimal choice of the counterexample S , we have that Λ and Γ are identical inside T . Now, contract the set T w.r.t. Λ , or equivalently, w.r.t. Γ . Let Λ' , Γ' denote the contracted solutions in G' , and let \mathcal{F}' be the contraction of \mathcal{F} . Then, by Lemma 4.1(i), these two solutions are feasible to $D_{\mathcal{F}'}(G', c')$, and S' is both Λ' -factor-critical and Γ' -factor-critical. Now, Λ' and Γ' are not identical inside S' , contradicting the minimal choice of G and S . \square

5.2 Finding a positively-critical dual optimal solution.

Let $\mathcal{F} \subseteq \mathcal{O}$ be a critical family with Π being an \mathcal{F} -critical dual. Let Ψ be a dual optimal solution to $D_{\mathcal{F}}(G, c)$. We present Algorithm 3 that modifies Ψ to an \mathcal{F} -positively-critical dual optimal solution. The correctness of the algorithm follows by showing that in every iteration, the modified solution $\bar{\Psi}$ is also dual optimal, and it is “closer” to Π .

Lemma 5.9. *Let $\mathcal{F} \subseteq \mathcal{O}$ be a critical family with Π being an \mathcal{F} -critical dual and let Ψ be a dual optimal solution to $D_{\mathcal{F}}(G, c)$. Suppose we consider a maximal set S such that Π and Ψ are not identical inside S , and $\Psi(S) > 0$. Define $\lambda = \min\{1, \Psi(S)/\Delta_{\Pi, \Psi}(S)\}$ if $\Delta_{\Pi, \Psi}(S) > 0$ and $\lambda = 1$ if*

Algorithm 3 Algorithm Positively-critical-dual-opt

Input: An optimal solution Ψ to $D_{\mathcal{F}}(G, c)$ and a \mathcal{F} -critical dual solution Π to $D_{\mathcal{F}}(G, c)$ Output: An \mathcal{F} -positively-critical dual optimal solution to $D_{\mathcal{F}}(G, c)$ 1. **Repeat** while Ψ is not \mathcal{F} -positively-critical dual.

- (a) Choose a maximal set $S \in \mathcal{F}$ with $\Psi(S) > 0$, such that Π and Ψ are not identical inside S .
- (b) Set $\Delta := \Delta_{\Pi, \Psi}(S)$.
- (c) Let $\lambda := \min\{1, \Psi(S)/\Delta\}$ if $\Delta > 0$ and $\lambda := 1$ if $\Delta = 0$.
- (d) Replace Ψ by the following $\bar{\Psi}$.

$$\bar{\Psi}(T) := \begin{cases} (1 - \lambda)\Psi(T) + \lambda\Pi(T) & \text{if } T \subsetneq S, \\ \Psi(S) - \Delta\lambda & \text{if } T = S, \\ \Psi(T) & \text{otherwise.} \end{cases} \quad (3)$$

2. **Return** $\bar{\Psi}$.

$\Delta_{\Pi, \Psi}(S) = 0$ and set $\bar{\Psi}$ as in (3). Then, $\bar{\Psi}$ is also a dual optimal solution to $D_{\mathcal{F}}(G, c)$, and either $\bar{\Psi}(S) = 0$ or Π and $\bar{\Psi}$ are identical inside S .

Proof. Let x be an optimal solution to $P_{\mathcal{F}}(G, c)$. Since $\Psi(S) > 0$, we have $x(\delta(S)) = 1$ and by Lemma 5.1, we have $\Delta = \Delta_{\Pi, \Psi}(S) \geq 0$. Now, the second conclusion is immediate from definition: if $\lambda = 1$, then we have that Π and $\bar{\Psi}$ are identical inside S ; if $\lambda < 1$, then we have $\bar{\Psi}(S) = 0$. For optimality, we show feasibility and verify the primal-dual slackness conditions.

The solution $\bar{\Psi}$ might have positive components on some sets $T \subsetneq S, T \in \mathcal{F}$ where $\Psi(T) = 0$ (but $\Pi(T) > 0$). However, $x(\delta(T)) = 1$ for all sets $T \subsetneq S, T \in \mathcal{F}$ by Lemma 5.3, since $x(\delta(S)) = 1$ by complementary slackness between x and Ψ . The choice of λ also guarantees $\bar{\Psi}(S) \geq 0$. We need to verify that all inequalities in $D_{\mathcal{F}}(G, c)$ are maintained and that all tight constraints in $D_{\mathcal{F}}(G, c)$ w.r.t. Ψ are maintained. This trivially holds if $uv \in E[V \setminus S]$. If $uv \in E[S] \setminus \text{supp}(x)$, the corresponding inequality is satisfied by both Π and Ψ and hence also by their linear combinations. If $uv \in E[S] \cap \text{supp}(x)$, then uv is tight for Ψ by the optimality of Ψ , and also for Π by Lemma 5.3.

It remains to verify the constraint corresponding to edges uv with $u \in S, v \in V \setminus S$. The contribution of $\sum_{T \in \mathcal{F}: uv \in \delta(T), T \setminus S \neq \emptyset} \Psi(T)$ is unchanged. The following claim completes the proof of optimality. \square

Claim 5.10. $\bar{\Psi}_S(u) + \bar{\Psi}(S) \leq \Psi_S(u) + \Psi(S)$ with equality whenever $uv \in \text{supp}(x)$.

Proof. We have

$$\bar{\Psi}(T) - \Psi(T) = \begin{cases} \lambda(\Pi(T) - \Psi(T)) & \text{if } T \subsetneq S, \\ -\Delta\lambda & \text{if } T = S. \end{cases}$$

Thus,

$$\begin{aligned} \bar{\Psi}_S(u) + \bar{\Psi}(S) &= \lambda(\Pi_S(u) - \Psi_S(u)) + \bar{\Psi}(S) - \Psi(S) + \Psi_S(u) + \Psi(S) \\ &= \lambda(\Pi_S(u) - \Psi_S(u) - \Delta) + \Psi_S(u) + \Psi(S). \end{aligned}$$

Now, $\Pi_S(u) - \Psi_S(u) \leq \Delta$, and equality holds whenever $uv \in \text{supp}(x) \cap \delta(S)$ by the consistency of Ψ and Π inside S (Lemma 5.1). \square

Corollary 5.11. *Let \mathcal{F} be a critical family with Π being an \mathcal{F} -critical dual feasible solution. Algorithm Positively-critical-dual-opt in Algorithm 3 transforms an arbitrary dual optimal solution Ψ to an \mathcal{F} -positively-critical dual optimal solution in at most $|\mathcal{F}|$ iterations.*

Proof. The correctness of the algorithm follows by Lemma 5.9. We bound the running time by showing that no set $S \in \mathcal{F}$ is processed twice. After a set S is processed, by Lemma 5.9, either Π and Ψ will be identical inside S or $\Psi(S) = 0$. Once Π and Ψ become identical inside a set, it remains so during all later iterations.

The value $\Psi(S)$ could be changed later only if we process a set $S' \supsetneq S$ after processing S . Let S' be the first such set. At the iteration when S was processed, by the maximal choice it follows that $\Psi(S') = 0$. Hence $\Psi(S')$ could become positive only if the algorithm had processed a set $Z \supsetneq S'$, $Z \in \mathcal{F}$ between processing S and S' , a contradiction to the choice of S' . \square

5.3 Extremal dual solutions.

In this section, we prove Lemma 3.2. The end result of the iterative procedure of the previous section can also be achieved by optimizing over dual solutions. The key property of the objective function is that it puts less weight on larger laminar sets.

Assume $\mathcal{F} \subseteq \mathcal{O}$ is a critical family, with Π being an \mathcal{F} -critical dual. Let x be the unique optimal solution to $P_{\mathcal{F}}(G, c)$. Let $\mathcal{F}_x = \{S \in \mathcal{F} : x(\delta(S)) = 1\}$ the collection of tight sets for x . A Π -extremal dual can be found by solving the following LP.

$$\begin{aligned}
\min h(\Psi, \Pi) &= \sum_{S \in \mathcal{V} \cup \mathcal{F}_x} \frac{r(S)}{|S|} \\
-r(S) &\leq \Psi(S) - \Pi(S) \leq r(S) \quad \forall S \in \mathcal{V} \cup \mathcal{F}_x \\
\sum_{S \in \mathcal{V} \cup \mathcal{F}_x : uv \in \delta(S)} \Psi(S) &= c(uv) \quad \forall uv \in \text{supp}(x) \\
\sum_{S \in \mathcal{V} \cup \mathcal{F}_x : uv \in \delta(S)} \Psi(S) &\leq c(uv) \quad \forall uv \in E \setminus \text{supp}(x) \\
\Psi(S) &\geq 0 \quad \forall S \in \mathcal{F}_x
\end{aligned} \tag{D_{\mathcal{F}}^*}$$

The support of Ψ is restricted to sets in $\mathcal{V} \cup \mathcal{F}_x$. Primal-dual slackness implies that the feasible solutions to this program coincide with the optimal solutions of $D_{\mathcal{F}}(G, c)$, hence an optimal solution to $D_{\mathcal{F}}^*$ is also an optimal solution to $D_{\mathcal{F}}(G, c)$.

Lemma 5.12. *Let $\mathcal{F} \subset \mathcal{O}$ be a critical family with Π being an \mathcal{F} -critical dual. Then, a Π -extremal dual optimal solution is also an \mathcal{F} -positively-critical dual optimal solution.*

Proof. We will show that whenever $\Psi(S) > 0$, the solutions Ψ and Π are identical inside S . Assume for a contradiction that this is not true for some $S \in \mathcal{F}$. Let $\lambda = \min\{1, \Psi(S)/\Delta_{\Pi, \Psi}(S)\}$ if $\Delta_{\Pi, \Psi}(S) > 0$ and $\lambda = 1$ if $\Delta_{\Pi, \Psi}(S) = 0$. Define $\bar{\Psi}$ as in (3). By Lemma 5.9, $\bar{\Psi}$ is also optimal to $D_{\mathcal{F}}(G, c)$ and thus feasible to $D_{\mathcal{F}}^*$. We show $h(\bar{\Psi}, \Pi) < h(\Psi, \Pi)$, which is a contradiction.

For every $T \in \mathcal{V} \cup \mathcal{F}_x$, let $\tau(T) = |\Psi(T) - \Pi(T)| - |\bar{\Psi}(T) - \Pi(T)|$. With this notation,

$$h(\Psi, \Pi) - h(\bar{\Psi}, \Pi) = \sum_{T \in \mathcal{V} \cup \mathcal{F}_x} \frac{\tau(T)}{|T|}.$$

If $T \setminus S = \emptyset$, then $\bar{\Psi}(T) = \Psi(T)$ and thus $\tau(T) = 0$. If $T \subsetneq S$, $T \in \mathcal{V} \cup \mathcal{F}$, then $|\bar{\Psi}(T) - \Pi(T)| = (1 - \lambda)|\Psi(T) - \Pi(T)|$, and thus $\tau(T) = \lambda|\Psi(T) - \Pi(T)|$. Since $\bar{\Psi}(S) = \Psi(S) - \Delta\lambda$, we have $\tau(S) \geq -\Delta\lambda$.

Let us fix an arbitrary $u \in S$, and let $\gamma = \max_{T \subsetneq S: u \in T, T \in \mathcal{V} \cup \mathcal{F}_x} |T|$.

$$\begin{aligned} h(\Psi, \Pi) - h(\bar{\Psi}, \Pi) &= \sum_{T \in \mathcal{V} \cup \mathcal{F}_x} \frac{\tau(T)}{|T|} \\ &\geq \sum_{T \subsetneq S: u \in T, T \in \mathcal{V} \cup \mathcal{F}_x} \frac{\tau(T)}{|T|} + \frac{\tau(S)}{|S|} \\ &\geq \frac{\lambda}{\gamma} \sum_{T \subsetneq S: u \in T, T \in \mathcal{V} \cup \mathcal{F}_x} |\Psi(T) - \Pi(T)| - \frac{\Delta\lambda}{|S|} \\ &\geq \frac{\lambda}{\gamma} (\Pi_S(u) - \Psi_S(u)) - \frac{\Delta\lambda}{|S|}. \end{aligned}$$

Case 1: If $\Delta > 0$, then pick $u \in S$ satisfying $\Pi_S(u) - \Psi_S(u) = \Delta$. Then the above inequalities give

$$h(\Psi, \Pi) - h(\bar{\Psi}, \Pi) \geq \Delta\lambda \left(\frac{1}{\gamma} - \frac{1}{|S|} \right) > 0.$$

The last inequality follows since $|S| > \gamma$.

Case 2: If $\Delta = 0$, then $\lambda = 1$ and therefore,

$$h(\Psi, \Pi) - h(\bar{\Psi}, \Pi) \geq \frac{1}{\gamma} \sum_{T \subsetneq S: u \in T, T \in \mathcal{V} \cup \mathcal{F}_x} |\Psi(T) - \Pi(T)|$$

Now, if Π and Ψ are not identical inside S , then there exists a node $u \in S$ for which the RHS is strictly positive. Thus, in both cases, we get $h(\bar{\Psi}, \Pi) < h(\Psi, \Pi)$, a contradiction to the optimality of Ψ to $D_{\mathcal{F}}^*$. \square

Proof of Lemma 3.2. By Lemma 3.1, the unique optimal x to $P_{\mathcal{F}}(G, c)$ is proper-half-integral. Lemma 5.12 already shows that a Γ -extremal dual solution Π is also \mathcal{F} -positively-critical. We need to show that the next family of cuts is a critical family. Recall that the set of cuts for the next round is defined as $\mathcal{H}' \cup \mathcal{H}''$, where $\mathcal{H}' = \{T \in \mathcal{F} : \Pi(T) > 0\}$, and \mathcal{H}'' is defined based on some cycles in $\text{supp}(x)$. We need to show that every set of $\mathcal{H}' \cup \mathcal{H}''$ is Π -factor-critical. This is straightforward for sets of \mathcal{H}' by the definition of the \mathcal{F} -positively-critical property.

It remains to show that the sets of \mathcal{H}'' are also Π -factor-critical. These are defined for odd cycles $C \in \text{supp}(x)$. Now, $\hat{C} \in \mathcal{H}''$ is the union of $V(C)$ and the maximal sets S_1, \dots, S_ℓ of \mathcal{H}' intersecting $V(C)$. We have $\Pi(S_j) > 0$ for each $j = 1, \dots, \ell$ and hence $x(\delta(S_j)) = 1$.

Let $u \in \hat{C}$ be an arbitrary node; we will construct the Π -critical matching \hat{M}_u in \hat{C} . Let us contract all sets S_1, \dots, S_ℓ to nodes s_1, \dots, s_ℓ w.r.t. Π . We know by Lemma 4.1(iii) that the image x' of x is proper-half-integral and that the odd cycle C projects to an odd cycle C' in $\text{supp}(x')$. Further, notice that $\{s_1, \dots, s_\ell\} \subseteq V(C')$, and therefore $V(C')$ is the image of the entire set \hat{C} . Let u' be the image of u ; since C' is an odd cycle, there is a perfect matching $M'_{u'} \subseteq C'$ of the set $V(C') \setminus \{u'\}$.

Assume first $u \in S_j$ for some $1 \leq j \leq \ell$. Then $u' = s_j$. The pre-image \hat{M} of $M'_{u'}$ in the original graph contains exactly one edge entering each S_k for $k \neq j$ and no edges entering S_j . Also, $\hat{M} \subseteq C$ and thus \hat{M} consists of tight edges w.r.t. Π . Consider the Π -critical matching M_u for u in S_j . For

$k \neq j$, if $a_k b_k \in \hat{M} \cap \delta(S_k)$, $a_k \in S_k$, then, let M_{a_k} be the Π -critical matching for a_k in S_k . The union of \hat{M} , M_u and the M_{a_k} 's give a Π -critical matching for u inside \hat{C} .

If $u \in \hat{C} \setminus (\cup_{j=1}^{\ell} S_j)$, then similarly there is a Π -critical matching M_{a_k} inside every S_k . The union of \hat{M} and the M_{a_k} 's give the Π -critical matching for u inside \hat{C} . We also have $\Pi(S) > 0$ for all non-maximal sets $S \in \mathcal{H}' \cup \mathcal{H}''$ since the only sets with $\Pi(S) = 0$ are those in \mathcal{H}'' , and they are all maximal ones. \square

6 Convergence.

The goal of this section is to prove Lemmas 3.3 and 3.4. Lemma 3.3 shows that the number of odd cycles in the support is nonincreasing. Lemma 3.4 shows that in a sequence of iterations where the number of cycles does not decrease, all the new cuts added continue to be included in subsequent iterations (till the number of cycles decreases). In order to establish Lemma 3.4, it is sufficient to show that the extremal dual solution has non-zero values on cuts that were added after the last decrease in the number of odd cycles.

These structural properties are established as follows. First we develop a primal-dual procedure that transforms a half-integral matching to satisfy a chosen subset of odd-set inequalities. Next, we apply this procedure starting with an appropriate primal/dual solution to obtain the optimal primal solution of the LP occurring in the cutting plane algorithm. The analysis of the procedure shows that the number of odd cycles is nonincreasing. The key ingredient in the proof of Lemma 3.4 is showing that whenever the number of odd cycles remains the same, then the extremal dual solution occurring in the cutting plane algorithm must be the same as the dual solution found by this procedure. As a consequence, properties of the dual solution found by this procedure also carry over to the extremal dual solution found by the algorithm.

6.1 The half-integral matching procedure.

We use the terminology of Edmonds' weighted matching algorithm [10] as described by Schrijver [26, Vol A, Chapter 26]. For a laminar family $\mathcal{L} \cup \mathcal{K}$, consider the following pair of primal and dual linear programs; note that the primal differs from $P_{\mathcal{L} \cup \mathcal{K}}(G, c)$ by requiring that the degree of every set in \mathcal{K} is precisely one, similar to the node constraints.

$$\begin{array}{ll}
\min \sum_{uv \in E} c(uv)z(uv) & (P_{\mathcal{L}}^{\mathcal{K}}(G, c)) \\
z(\delta(u)) = 1 \quad \forall u \in V & \\
z(\delta(S)) = 1 \quad \forall S \in \mathcal{K} & \\
z(\delta(S)) \geq 1 \quad \forall S \in \mathcal{L} & \\
z \geq 0 &
\end{array}
\qquad
\begin{array}{ll}
\max \sum_{S \in \mathcal{V} \cup \mathcal{L} \cup \mathcal{K}} \Lambda(S) & (D_{\mathcal{L}}^{\mathcal{K}}(G, c)) \\
\sum_{S \in \mathcal{V} \cup \mathcal{L} \cup \mathcal{K}: uv \in \delta(S)} \Lambda(S) \leq c(uv) \quad \forall uv \in E & \\
\Lambda(S) \geq 0 \quad \forall S \in \mathcal{L} &
\end{array}$$

We note that every feasible solution to $P_{\mathcal{L}}^{\mathcal{K}}(G, c)$ is also a feasible solution to $P_{\mathcal{L} \cup \mathcal{K}}(G, c)$, whereas a feasible solution to $D_{\mathcal{L}}^{\mathcal{K}}(G, c)$ is a feasible solution to $D_{\mathcal{L} \cup \mathcal{K}}(G, c)$ only if $\Lambda(S) \geq 0$ for all sets $S \in \mathcal{K}$.

The aim of the procedure is that for a given laminar family $\mathcal{F} \cup \mathcal{K}$ satisfying certain structural properties, we wish to transform a pair of primal and dual feasible solutions to $(P_{\mathcal{F}}^{\emptyset}(G, c), D_{\mathcal{F}}^{\emptyset}(G, c))$ to optimal solutions to a pair of primal and dual optimal solutions to $(P_{\mathcal{L}}^{\mathcal{K}}(G, c), D_{\mathcal{L}}^{\mathcal{K}}(G, c))$ for some $\mathcal{L} \subseteq \mathcal{F}$. The following notion of a *valid configuration* encapsulates these structural properties. We say that $(\mathcal{L}, \mathcal{K}, z, \Lambda)$ form a valid configuration, if the following hold:

- (A) $\mathcal{L} \cup \mathcal{K} \subset \mathcal{O}$ is a laminar family, and all sets in \mathcal{K} are disjoint from each other and all sets in \mathcal{L} . Λ is a feasible solution to $D_{\mathcal{L}}^{\mathcal{K}}(G, c)$ with $\Lambda(S) > 0$ for all $S \in \mathcal{L}$. Further, every set $S \in \mathcal{L} \cup \mathcal{K}$ is $(G_{\Lambda}, \mathcal{L} \cup \mathcal{K})$ -factor-critical, where G_{Λ} denotes the graph of tight edges wrt Λ .
- (B) z is proper-half-integral, satisfying all constraints of $P_{\mathcal{L}}^{\mathcal{K}}(G, c)$ except that $z(\delta(S)) = 0$ may hold for some $S \in \mathcal{K}$. The support of z is an odd cycle inside every such set S . Inside every other set $S \in \mathcal{K} \cup \mathcal{L}$, $\text{supp}(z)$ spans all vertices in S and is a disjoint union of edges and a (possibly empty) even path.
- (C) Every edge in $\text{supp}(z)$ is tight for Λ , and $z(\delta(S)) = 1$ for every $S \in \mathcal{L}$.

Algorithm 4 Half-integral Matching Procedure

Input. A graph G with edge costs c , and a valid configuration $(\mathcal{F}, \mathcal{K}, x, \Pi)$.

Output. A valid configuration $(\mathcal{L}, \mathcal{K}, z, \Lambda)$ with $\mathcal{L} \subseteq \mathcal{F}$, and z being a proper-half-integral optimal solution to $P_{\mathcal{L}}^{\mathcal{K}}(G, c)$.

1. Initialize $z = x$, $\Lambda = \Pi$, $\mathcal{L} = \mathcal{F}$. Let $G^* = (\mathcal{V}^*, E^*)$, where $E^* \subseteq E$ are edges that are tight w.r.t. Λ , and all maximal sets of $\mathcal{L} \cup \mathcal{K}$ are contracted w.r.t. Λ ; c^* and z^* are defined by the contraction. Let $T \subseteq \mathcal{V}^*$ denote the set of exposed nodes in z^* , and let $R(\supseteq T)$ be the set of isolated nodes and nodes incident to $\frac{1}{2}$ -edges in z^* .
2. While T is not empty,

Case I: There exists an alternating T - R -walk in G^ .* Let $P = v_0 v_1 \dots v_{2k+1}$ denote a shortest such walk.

 - (a) If P is an alternating path, and $v_{2k+1} \in T$, then change z by alternating along P .
 - (b) If P is an alternating path, and $v_{2k+1} \in R \setminus T$, then let C denote the odd cycle containing v_{2k+1} . Change z by alternating along P , and replacing z on C by a blossom with base v_{2k+1} .
 - (c) If P is not a path, then by Claim 6.2, it contains an even alternating path P_1 to a blossom C . Change z by alternating along P_1 , and setting $z^*(uv) = 1/2$ on every edge of C .

Case II: There exists no alternating T - R -walk in G^ .* Define

$$\begin{aligned} \mathcal{B}^+ &:= \{S \in \mathcal{V}^* : \exists \text{ an even alternating path from } T \text{ to } S\}, \\ \mathcal{B}^- &:= \{S \in \mathcal{V}^* : \exists \text{ an odd alternating path from } T \text{ to } S\}. \end{aligned}$$

For some $\varepsilon > 0$, reset

$$\Lambda(S) := \begin{cases} \Lambda(S) + \varepsilon & \text{if } S \in \mathcal{B}^+, \\ \Lambda(S) - \varepsilon & \text{if } S \in \mathcal{B}^-. \end{cases}$$

Choose ε to be the maximum value such that Λ remains feasible to $D_{\mathcal{L}}^{\mathcal{K}}(G, c)$.

- (a) If some new edge becomes tight, then E^* is extended.
 - (b) If $\Lambda(S) = 0$ for some $S \in \mathcal{L} \cap \mathcal{B}^-$ after the modification, then unshrink the node S . Set $\mathcal{L} := \mathcal{L} \setminus \{S\}$.
-

The input to the procedure (Algorithm 4) will be a graph G with costs c , and a valid configuration $(\mathcal{F}, \mathcal{K}, x, \Pi)$. The procedure is iterative. In each iteration, it maintains a valid configuration $(\mathcal{L}, \mathcal{K}, z, \Lambda)$, where $\mathcal{L} \subseteq \mathcal{F}$; the set \mathcal{K} never changes during the execution of the procedure. We terminate once z is feasible to $P_{\mathcal{L}}^{\mathcal{K}}(G, c)$. The complementary slackness conditions (C) imply that if z is feasible to $P_{\mathcal{L}}^{\mathcal{K}}(G, c)$, then (z, Λ) form an optimal primal-dual pair to $(P_{\mathcal{L}}^{\mathcal{K}}(G, c), D_{\mathcal{L}}^{\mathcal{K}}(G, c))$.

The procedure works on the graph $G^* = (\mathcal{V}^*, E^*)$, obtained the following way from G : We first remove every edge in E that is not tight w.r.t. Λ , and then contract all maximal sets of $\mathcal{L} \cup \mathcal{K}$ w.r.t. Λ . The node set of \mathcal{V}^* is identified with the pre-images. Let c^* denote the contracted cost function and z^* denote the image of z . Since E^* consists only of tight edges, $\Lambda(u) + \Lambda(v) = c^*(uv)$ for every edge $uv \in E^*$. Let $T \subseteq \mathcal{V}^*$ denote the sets in \mathcal{K} for which $z(\delta(S)) = 0$; by property (B), T is the set of nodes in \mathcal{V}^* that have degree 0 in z^* , whereas all other nodes have fractional degree one.

Claim 6.1. *The vector z^* defined above is proper-half-integral. Assuming the uniqueness condition (\star) , the number of odd cycles in $\text{supp}(z^*)$ plus the number of exposed nodes in z^* equals $\text{odd}(z)$.*

Proof. It is clear that z^* is half-integral, and that the image of every odd cycle in $\text{supp}(z)$ is a cycle in $\text{supp}(z^*)$ or an exposed node. The last part of property (B) implies that all these cycles in $\text{supp}(z^*)$ must be odd. \square

In the execution of the procedure, we may decrease $\Lambda(S)$ to 0 for a set $S \in \mathcal{L}$. In this case, we remove S from \mathcal{L} . We modify G^* , c^* and z^* accordingly. This operation will be referred as ‘unshrinking’ S . New sets will never be added to \mathcal{L} , that is, no new sets will be shrunk after the initial contractions: $|\mathcal{V}^*|$ may only increase. In contrast, sets in \mathcal{K} are never unshrunk and the family \mathcal{K} does not change.

The procedure works by modifying the solution z^* and the dual solution Λ^* . An edge $uv \in E^*$ is called a 0-edge/ $\frac{1}{2}$ -edge/1-edge according to the value $z^*(uv)$. A modification of z^* in G^* naturally extends to a modification of z in G . Indeed, if $S \in \mathcal{L}$ is a shrunk node in \mathcal{V}^* , and z^* is modified so that there is an 1-edge incident to S in G^* , then let u_1v_1 be the pre-image of this edge in G , with $u_1 \in S$. Then modify z inside S to be identical with the Λ -critical-matching M_{u_1} inside S . If there are two half-edges incident to S in G^* , then let u_1v_1, u_2v_2 be the pre-image of these edges in G , with $u_1, u_2 \in S$. Then modify z inside S to be identical with the convex combination $(1/2)(\chi(M_{u_1}) + \chi(M_{u_2}))$ of the Λ -critical-matchings M_{u_1} and M_{u_2} inside S . Note that this modification preserves the second part of property (C).

A walk $P = v_0v_1v_2 \dots v_k$ in G^* is called an alternating walk, if every odd edge is a 0-edge and every even edge is a 1-edge. If every node occurs in P at most once, it is called an alternating path. By *alternating along the path P* , we mean modifying $z^*(v_i v_{i+1})$ to $1 - z^*(v_i v_{i+1})$ on every edge of P . If k is odd, $v_0 = v_k$ and no other node occurs twice, then P is called a *blossom* with base v_0 . We note that an alternating T - R -walk in the algorithm must be of odd length — the starting edge is a 0-edge and the ending edge also has to be a 0-edge since the last vertex is from R , has total (fractional) degree 1 and therefore has no 1-edge incident to it. The following claim is straightforward.

Claim 6.2 ([26, Thm 24.3]). *Let P be a shortest alternating T – R -walk. Either P is an alternating path, or it contains a blossom C and an even alternating path from v_0 to the base of the blossom.* \square

The procedure is described in Algorithm 4. Let us note that in [5] we extend it to a ‘complete’ algorithm to find a minimum-cost perfect matching where the intermediate solutions are half-integral and satisfy the degree constraints for all vertices.




Figure 5: The possible modifications in the Half-integral Matching Procedure.

The scenarios in *Case I* are illustrated in Figure 5. In Case II, we observe that $T \subseteq \mathcal{B}^+$ and further, $\mathcal{B}^+ \cap \mathcal{B}^- = \emptyset$ (otherwise, there existed a $T - T$ alternating walk and hence we have case I). The following claim is easy to verify. Note that $\mathcal{L} \cup \mathcal{K}$ will always be a critical family because of Claim 2.4.

Claim 6.3. *In every iteration of the procedure, $(\mathcal{L}, \mathcal{K}, z, \Lambda)$ is a valid configuration.*

The key to the proof of Lemma 3.3 is the following lemma, showing that $\text{odd}(z)$ is non-increasing during the execution of the procedure.

Lemma 6.4. *Let x be the input solution to the Half-Integral Matching procedure, z be the solution at the beginning of execution of an arbitrary iteration of the procedure, and let α be the number of odd cycles in $\text{supp}(x)$ disjoint from all members of \mathcal{K} , that are absent in $\text{supp}(z)$. Then $\text{odd}(x) \geq \text{odd}(z) + 2\alpha$. Further, if $\text{odd}(x) = \text{odd}(z)$, then cases I(a) and I(b) are never executed.*

Proof. We will investigate how the number of odd cycles in $\text{supp}(z^*)$ plus the number of exposed nodes change; by Claim 6.1, this equals $\text{odd}(z)$. In Case I(a), the number of exposed nodes decreases by two. In Case I(b), both the number of exposed nodes and the number of cycles decrease by one. In Case I(c), the number of exposed nodes decreases by one, but we obtain a new odd cycle, hence the total quantity remains unchanged. In Case II, the primal solution is not modified at all.

Further, the cycles in $\text{supp}(z^*)$ are in one-to-one correspondence with the cycles in $\text{supp}(z)$ that are not contained inside some member of \mathcal{K} . Such a cycle can be removed only by performing the operation in Case I(b). This must be executed α times, therefore $\text{odd}(z) \leq \text{odd}(x) - 2\alpha$. \square

Lemma 3.3 will be an immediate consequence of the next lemma. For the proof of this next lemma, we assume that the procedure terminates in finite number of iterations. In the next section, we will show that the procedure indeed terminates in strongly polynomial time.

Lemma 6.5. *Assume (\star) holds. Let \mathcal{F} be a critical family, and let x be an optimal solution to $P_{\mathcal{F}}(G, c)$, and Π an \mathcal{F} -positively-critical dual optimal solution to $D_{\mathcal{F}}(G, c)$. Define the sets \mathcal{H}' and \mathcal{H}'' as in steps 2(b) and (c) in Algorithm C-P matching (Algorithm 2), and let $\mathcal{H} = \mathcal{H}' \cup \mathcal{H}''$. Let y be an optimal solution to $P_{\mathcal{H}}(G, c)$ and let Ψ be an \mathcal{H} -positively-critical optimal solution to $P_{\mathcal{H}}(G, c)$. Then $\text{odd}(y) \leq \text{odd}(x)$, and if $\text{odd}(y) = \text{odd}(x)$ then $\Psi(S) > 0$ for every $S \in \mathcal{H}''$.*

Proof. We first note that Lemma 3.2 guarantees that \mathcal{H} is a critical family; hence the existence of Ψ is guaranteed. Further, (\star) guarantees the uniqueness of x and y . To prove the lemma by

contradiction, consider a counterexample (G, c, \mathcal{F}) with $|V|$ minimal. That is, either $\text{odd}(y) > \text{odd}(x)$, or $\text{odd}(y) = \text{odd}(x)$ but $\Psi(S) = 0$ for some $S \in \mathcal{H}''$.

Claim 6.6. $\Psi(S) = 0$ for every $S \in \mathcal{H}'$.

Proof. Consider a set $S \in \mathcal{H}'$ with $\Psi(S) > 0$. We first observe that Π and Ψ are feasible solutions to $D_{\mathcal{H}}(G, c)$. Further, if $S \in \mathcal{H}'$, then S is both (Π, \mathcal{H}) -factor-critical and (Ψ, \mathcal{H}) -factor-critical. Hence, by Lemma 5.8, Π and Ψ are identical inside S .

Let us contract the set S with respect to Π ; this is equivalent to contracting with respect to Ψ since Π and Ψ are identical inside S . Let $\hat{G}, \hat{c}, \hat{\mathcal{F}}, \hat{x}, \hat{\Pi}, \hat{\mathcal{H}}, \hat{y}, \hat{\Psi}$ denote the respective images to the contracted instance. Lemma 4.1 guarantees that in the contracted instances, $\hat{\Pi}$ and $\hat{\Psi}$ are critical families, \hat{x} and \hat{y} are optimal solutions to $P_{\hat{\mathcal{F}}}(\hat{G}, \hat{c})$ and $P_{\hat{\mathcal{H}}}(\hat{G}, \hat{c})$, respectively, and $\text{odd}(\hat{x}) = \text{odd}(x)$, $\text{odd}(\hat{y}) = \text{odd}(y)$. Furthermore, $\hat{\Pi}$ is an $\hat{\mathcal{F}}$ -positively-critical dual optimal solution to $D_{\hat{\mathcal{F}}}(\hat{G}, \hat{c})$. Now, by taking $\hat{\mathcal{H}} = \hat{\mathcal{H}}' \cup \hat{\mathcal{H}}''$, where the sets $\hat{\mathcal{H}}'$ and $\hat{\mathcal{H}}''$ are the ones obtained in steps 2(b) and (c) in Algorithm C-P matching applied for $(\hat{G}, \hat{c}, \hat{\mathcal{F}})$, we obtain a smaller counterexample. That is, if the counterexample (G, c, \mathcal{F}) is such that $\text{odd}(y) > \text{odd}(x)$, then $\text{odd}(\hat{y}) > \text{odd}(\hat{x})$ in the contracted instance; if the counterexample (G, c, \mathcal{F}) is such that $\text{odd}(y) = \text{odd}(x)$ but $\Psi(S) = 0$ for some $S \in \mathcal{H}''$, then $\text{odd}(\hat{y}) = \text{odd}(\hat{x})$, $\hat{\Psi}(\hat{S}) = 0$ and $\hat{S} \in \hat{\mathcal{H}}''$. \square

Let us define $\mathcal{K} := \{S \in \mathcal{H}'' : \Psi(S) > 0\}$, and apply the Half-integral Matching Procedure with input $(\mathcal{H}', \mathcal{K}, x, \Pi)$.

Claim 6.7. $(\mathcal{H}', \mathcal{K}, x, \Pi)$ is a valid configuration.

Proof. We verify only the nontrivial properties. We show that all sets in \mathcal{H}' are disjoint from all sets in \mathcal{K} . Consider a set $S \in \mathcal{K}$; thus $\Psi(S) > 0$. Hence S is both (Π, \mathcal{H}) -factor-critical and (Ψ, \mathcal{H}) -factor-critical, and therefore Lemma 5.8 is applicable. Consequently Ψ and Π must be identical inside S . For the sake of contradiction, assume there exists a set $T \subsetneq S$, $T \in \mathcal{H}'$. Then $\Psi(T) = \Pi(T) > 0$, contradicting Claim 6.6.

Thus, the sets in \mathcal{K} are disjoint from each other as well as the ones in \mathcal{H}' . This immediately implies that $\text{supp}(x)$ is an odd cycle inside every set in \mathcal{K} since $\mathcal{K} \subseteq \mathcal{H}''$. Finally, by Claim 4.1 (iii), we have that $\text{supp}(x)$ inside each set T is a disjoint union of edges and an even path. \square

Let $(\mathcal{L}, \mathcal{K}, z, \Lambda)$ denote the output of the Half-integral Matching Procedure applied to $(\mathcal{H}', \mathcal{K}, x, \Pi)$, where $\mathcal{L} \subseteq \mathcal{H}'$ and z is an optimal solution to $P_{\mathcal{L} \cup \mathcal{K}}^{\mathcal{K}}(G, c)$. Lemma 6.4 implies that $\text{odd}(z) \leq \text{odd}(x)$.

$$(\mathcal{H}', \mathcal{K}, x, \Pi) \xrightarrow[\text{Half-integral matching procedure}]{} (\mathcal{L}, \mathcal{K}, z, \Lambda)$$

Claim 6.8. $y = z$.

Proof. We know that y is an optimal solution to $P_{\mathcal{H}}(G, c)$. We observe that y is also an optimal solution to $P_{\mathcal{K}}(G, c)$. This follows since by Claim 6.6, we have $S \in \mathcal{K}$ whenever $\Psi(S) > 0$. Hence, the solution Ψ is feasible to $D_{\mathcal{K}}(G, c)$ and (y, Ψ) satisfy complementary slackness for primal-dual pair of programs $(P_{\mathcal{K}}(G, c), D_{\mathcal{K}}(G, c))$.

Further, y is a feasible solution to $P_{\mathcal{L}}^{\mathcal{K}}(G, c)$ by the choice of \mathcal{K} . We now note that the feasible region of $P_{\mathcal{L}}^{\mathcal{K}}(G, c)$ is a subset of the feasible region of $P_{\mathcal{L} \cup \mathcal{K}}(G, c)$, which in turn, is a subset of the feasible region of $P_{\mathcal{K}}(G, c)$. Therefore, y is also optimal to all three programs, and consequently, $c(y) = c(z)$. This in turn implies that z is also optimal to $P_{\mathcal{L} \cup \mathcal{K}}(G, c)$. Since $\mathcal{L} \cup \mathcal{K}$ is a critical family, the uniqueness assumption (\star) implies that y and z must be identical. \square

Hence we have $\text{odd}(y) \leq \text{odd}(x)$. Since (G, c, \mathcal{F}) was a counterexample, we must have $\text{odd}(y) = \text{odd}(x)$ and $\Psi(S) = 0$ for some $S \in \mathcal{H}''$. This is equivalent to \mathcal{K} being a strict subset of \mathcal{H}'' , that is, there exists a cycle C in $\text{supp}(x)$ that is disjoint from all sets in \mathcal{K} . Let $\hat{C} = V(C) \cup (\cup_{T \in \mathcal{H}': T \cap V(C) \neq \emptyset} T)$ denote the set included in \mathcal{H}'' corresponding to the cycle C in Algorithm C-P matching. The following claim implies that $y(\delta(\hat{C})) = 0$ contradicting the feasibility of y for $P_{\mathcal{H}}(G, c)$ and thus completes the proof. \square

Claim 6.9. *If $\text{odd}(x) = \text{odd}(z)$, then $z(\delta(\hat{C})) = 0$.*

Proof. By Lemma 6.4, $C \in \text{supp}(z)$ for otherwise, $\text{odd}(z) \leq \text{odd}(x) - 2$. Let C^* denote the cycle corresponding to C in the contracted graph at the beginning of the first iteration. Consider the case where all sets $T \in \mathcal{H}'$ that intersect $V(C)$ also belong to \mathcal{L} . Then all these sets are contracted in the current G^* . This implies that C^* is present in $\text{supp}(z^*)$ and hence $z^*(\delta(V(C^*))) = 0$. Therefore $z(\delta(\hat{C})) = 0$.

Consider the earliest iteration such that a set $T \in \mathcal{H}'$ intersecting $V(C)$ leaves \mathcal{L} . This means that the dual value on T decreased to zero in this iteration, that is $T^* \in \mathcal{B}^-$. Now, this is possible only if the cycle C^* is absent at the beginning of this iteration. A cycle in the support of the primal solution could have been broken only by the execution of Case I(b). This contradicts Lemma 6.4. \square

6.2 The non-decreasing scenario.

In this section we analyze further properties of the Half-integral Matching procedure to show termination, and also to establish structural properties needed for the proof of Lemma 3.4 in the subsequent section. If $\text{odd}(z)$ is unchanged during a certain number of iterations of the procedure, we say that these iterations form a *non-decreasing phase*. We say that the procedure itself is non-decreasing, if $\text{odd}(z)$ does not decrease anytime. We show that every non-decreasing phase may contain at most $|V| + |\mathcal{F}|$ iterations and therefore the procedure terminates in strongly polynomial time.

Let us now analyze the first non-decreasing phase \mathcal{P} of the procedure, starting from the input x . These results will also be valid for later non-decreasing phases. Consider an intermediate iteration with a valid configuration $(\mathcal{L}, \mathcal{K}, z, \Lambda)$ at the beginning of the iteration. Let $T \subseteq \mathcal{V}^*$ denote the set of exposed nodes in z^* , and $R \supseteq T$ is the set of exposed nodes and the node sets of the 1/2-cycles. Let us define the set of outer/inner nodes of G^* as those having even/odd length alternating walk from R in G^* . Let \mathcal{N}_o and \mathcal{N}_i denote their sets, respectively. Clearly, $\mathcal{B}^+ \subseteq \mathcal{N}_o$, $\mathcal{B}^- \subseteq \mathcal{N}_i$ in Case II of the algorithm.

Lemma 6.10. *If \mathcal{P} is a non-decreasing phase, then if a node in \mathcal{V}^* is outer in any iteration of phase \mathcal{P} , it remains a node in \mathcal{V}^* and an outer node in every later iteration of \mathcal{P} . If a node is inner in any iteration of \mathcal{P} , then in any later iteration of \mathcal{P} , it is either an inner node, or it has been unshrunk in an intermediate iteration.*

Proof. Since \mathcal{P} is a non-decreasing phase, Cases I(a) and (b) can never be executed. We show that the claimed properties are maintained during an iteration.

In Case I(c), a new odd cycle C is created, and thus the vertex set of C is added to R . Let $P_1 = v_0 \dots v_{2\ell}$ denote the even alternating path with $v_0 \in T$, $v_{2\ell} \in C$. If a node $u \in \mathcal{V}^*$ had an even/odd alternating walk from v_0 before changing the solution, it will have an even/odd walk alternating from $v_{2\ell} \in R$ after changing the solution.

In Case II, the alternating paths from T to the nodes in \mathcal{B}^- and \mathcal{B}^+ are maintained when the duals are changed. The only nontrivial case is when a set S is unshrunk; then all inner and outer nodes maintain their inner and outer property for the following reason: if u_1v_1 is a 1-edge and u_2v_2 is a 0-edge entering S after unshrinking, with $u_1, u_2 \in S$, we claim that there exists an even alternating path inside S from u_1 to u_2 using only tight edges wrt Λ . Indeed, during the unshrinking, we modify z to M_{u_1} inside S . Also, by the Λ -factor-critical property, all edges of M_{u_2} are tight w.r.t. Λ . Hence the symmetric difference of M_{u_1} and M_{u_2} contains an even alternating path from u_1 to u_2 .

We have to check that vertices in $\mathcal{N}_o \setminus \mathcal{B}^+$ and $\mathcal{N}_i \setminus \mathcal{B}^-$ also maintain their outer and inner property. These are the nodes having even/odd alternating paths from an odd cycle, but not from exposed nodes. The nodes in these paths are disjoint from $\mathcal{B}^- \cup \mathcal{B}^+$ and are thus maintained. Indeed, if $(\mathcal{B}^- \cap \mathcal{N}_o) \setminus \mathcal{B}^+ \neq \emptyset$ or $(\mathcal{B}^+ \cap \mathcal{N}_i) \setminus \mathcal{B}^- \neq \emptyset$, then we would get an alternating walk from T to an odd cycle, giving the forbidden Case I(b). \square

The termination of the algorithm is guaranteed by the following simple corollary.

Corollary 6.11. *The non-decreasing phase \mathcal{P} may consist of at most $|V| + |\mathcal{F}|$ iterations.*

Proof. Case I may occur at most $|\mathcal{K}|$ times as it decreases the number of exposed nodes. In Case II, either \mathcal{N}_i is extended, or a set is unshrunk. By Lemma 6.10, the first scenario may occur at most $|V| - |\mathcal{K}|$ times (since exposed nodes are outer nodes, and outer nodes remain outer nodes) and the second at most $|\mathcal{F}|$ times (the maximum number of shrunk nodes). \square

In the rest of the section, we focus on the case when the entire procedure is non-decreasing.

Lemma 6.12. *Assume the Half-integral Matching procedure is non-decreasing with input $(\mathcal{F}, \mathcal{K}, x, \Pi)$ and output $(\mathcal{L}, \mathcal{K}, z, \Lambda)$. Let \mathcal{N}_o and \mathcal{N}_i denote the final sets of outer and inner nodes in G^* .*

- *If $\Lambda(S) > \Pi(S)$ then S is an outer node in \mathcal{V}^* .*
- *If $\Lambda(S) < \Pi(S)$, then either $S \in \mathcal{F} \setminus \mathcal{L}$, (that is, S was unshrunk during the algorithm and $\Lambda(S) = 0$) or S is an inner node in \mathcal{V}^* , or S is a node in \mathcal{V}^* incident to an odd cycle in $\text{supp}(z)$.*

Proof. If $\Lambda(S) > \Pi(S)$, then $S \in \mathcal{B}^+$ in some iteration of the algorithm. By Lemma 6.10, this remains an outer node in all later iterations. The conclusion follows similarly for $\Lambda(S) < \Pi(S)$. \square

Lemma 6.13. *Assume the Half-integral Matching procedure is non-decreasing. Let $(\mathcal{L}, \mathcal{K}, z, \Lambda)$ be the terminating valid configuration, and let G^* be the corresponding contracted graph, \mathcal{N}_o and \mathcal{N}_i be the sets of outer and inner nodes. Let $\Theta : \mathcal{V}^* \rightarrow \mathbb{R}$ be an arbitrary optimal solution to the dual $D_0(G^*, c^*)$ of the bipartite relaxation. If $S \in \mathcal{V}^*$ is incident to an odd cycle in $\text{supp}(z)$, then $\Lambda(S) = \Theta(S)$. Further $S \in \mathcal{N}_o$ implies $\Lambda(S) \leq \Theta(S)$, and $S \in \mathcal{N}_i$ implies $\Lambda(S) \geq \Theta(S)$.*

Proof. For $S \in \mathcal{N}_o \cup \mathcal{N}_i$, let $\ell(S)$ be the length of the shortest alternating path. The proof is by induction on $\ell(S)$. Recall that there are no exposed nodes in z , hence $\ell(S) = 0$ means that S is contained in an odd cycle C . Then $\Theta(S) = \Lambda(S)$ is a consequence of Lemma 5.8: both Θ and Λ are optimal dual solutions in G^* , and an odd cycle in the support of the primal optimum z is both Λ -factor-critical and Θ -factor-critical.

For the induction step, assume the claim for $\ell(S) \leq i$. Consider a node $U \in \mathcal{V}^*$ with $\ell(U) = i+1$. There must be an edge f in E^* between S and U for some S with $\ell(S) = i$. This is a 0-edge if i is even and a 1-edge if i is odd.

Assume first i is even. By induction, $\Lambda(S) \leq \Theta(S)$. The edge f is tight for Λ , and $\Theta(S) + \Theta(U) \leq c^*(f)$. Consequently, $\Lambda(U) \geq \Theta(U)$ follows. Next, assume i is odd. Then $\Lambda(S) \geq \Theta(S)$ by induction. Then, $\Lambda(U) \leq \Theta(U)$ follows as f is tight for both Λ and Θ . \square

6.3 Proof of Lemma 3.4.

We follow a similar strategy to prove Lemma 3.4 as for Lemma 3.3: we formulate an analogue of Lemma 6.5. For that, we need the notion of outer nodes with respect to intermediate solutions in Algorithm C-P-matching. Let \mathcal{F} be a critical family, and let x be an optimal solution to $P_{\mathcal{F}}(G, c)$, and Π an \mathcal{F} -positively-critical dual optimal solution to $D_{\mathcal{F}}(G, c)$. Let us define the sets \mathcal{H}' and \mathcal{H}'' as in steps 2(b) and (c) in Algorithm C-P-matching (Algorithm 2), and let $\mathcal{H} = \mathcal{H}' \cup \mathcal{H}''$.

Let us consider the graph $G^* = (\mathcal{V}^*, E^*)$ obtained by first deleting all edges that are non-tight wrt Π and contracting all maximal sets of $\mathcal{H}' \cup \mathcal{H}''$ wrt Π . We observe that the resulting instance is identical to the contracted graph obtained in the execution of the Half-integral Matching procedure using the configuration $(\mathcal{H}', \mathcal{H}'', x, \Pi)$. Even though the configuration may not be a valid one, the contraction operation and the outer property of the nodes in G^* are well-defined. Let $\text{Outer}(x, \Pi, \mathcal{H})$ denote the sets in \mathcal{H} contained in the outer nodes of G^* , that is,

$$\text{Outer}(x, \Pi, \mathcal{H}) := \{S \in \mathcal{H} : \exists T \in \mathcal{V}^*, T \text{ is an outer node in } G^*, S \subseteq T \text{ and } \Pi(S) > 0\}.$$

Observe that $\mathcal{H}'' \subseteq \text{Outer}(x, \Pi, \mathcal{H})$ since every set in \mathcal{H}'' is exposed in G^* .

Lemma 6.14. *Assume (\star) holds, and use the notation $x, \Pi, \mathcal{F}, \mathcal{H}$ as above. Let y be an optimal solution to $P_{\mathcal{H}}(G, c)$ and let Ψ be a Π -extremal optimal solution to $D_{\mathcal{H}}(G, c)$. Let $\mathcal{H}_{\text{next}}$ denote the next round of cuts to be imposed by Algorithm C-P-matching from (y, Ψ) . If $\text{odd}(y) = \text{odd}(x)$ then $\text{Outer}(x, \Pi, \mathcal{H}) \subseteq \text{Outer}(y, \Psi, \mathcal{H}_{\text{next}})$.*

Proof of Lemma 3.4. Using the notation of the lemma, let \mathcal{F}_k be the set of constraints in iteration k , and $\mathcal{H}_k'' = \mathcal{F}_{k+1} \setminus \mathcal{F}_k$. Note that by definition $\mathcal{H}_k'' \subseteq \text{Outer}(x_k, \Pi_k, \mathcal{F}_{k+1})$. Also, by Lemma 6.14, if a set enters $\text{Outer}(x_k, \Pi_k, \mathcal{F}_{k+1})$ then it remains in $\text{Outer}(x_\ell, \Pi_\ell, \mathcal{F}_{\ell+1})$ for all $\ell > k$ if all iterations between the k 'th and the ℓ 'th are non-decreasing. Since $\Pi_\ell(S) > 0$ for all sets $S \in \text{Outer}(x_\ell, \Pi_\ell, \mathcal{F}_{\ell+1})$, all cuts corresponding to sets in $\text{Outer}(x_\ell, \Pi_\ell, \mathcal{F}_{\ell+1})$ will be imposed in the ℓ 'th iteration, as required. \square

The rest of the section is dedicated to proving Lemma 6.14. For a contradiction, let us choose a counterexample (G, c, \mathcal{F}) with $|V|$ minimal. By Lemma 6.5, $\Psi(S) > 0$ for every $S \in \mathcal{H}''$. By taking $\mathcal{K} = \mathcal{H}''$, Claims 6.6 and 6.7, and their proofs hold verbatim. Hence the support of Ψ is identical to \mathcal{H}'' . Hence, $(\mathcal{H}', \mathcal{H}'', x, \Pi)$ is a valid configuration. Let us run the Half-Integral Matching Procedure with input $(\mathcal{H}', \mathcal{H}'', x, \Pi)$, terminating with $(\mathcal{L}, \mathcal{H}'', z, \Lambda)$. As in the proof of Lemma 6.5, we have $y = z$. The key part of the proof is showing that Λ also coincides with Ψ .

Lemma 6.15. *We have $\Lambda = \Psi$.*

Lemma 6.14 is a straightforward consequence of Lemmas 6.15 and 6.10.

Proof of Lemma 6.14. By Lemma 6.15, $\mathcal{L} = \emptyset$ since the support of Ψ is \mathcal{H}'' . By Lemma 6.10, the sets $S \in \text{Outer}(x, \Pi, \mathcal{H})$ are outer nodes in the graph G^* obtained after deleting all edges that are non-tight wrt Ψ and contracting all maximal sets of \mathcal{H}'' wrt Ψ . Also, $\Psi(S) = \Lambda(S) \geq \Pi(S) > 0$ for every $S \in \text{Outer}(x, \Pi, \mathcal{H})$ since such sets are outer nodes at the beginning of the execution of the Half-Integral Matching Procedure and the dual values on outer nodes are non-decreasing throughout

the execution of the procedure. Therefore, $\text{Outer}(x, \Pi, \mathcal{H}) \subseteq \text{Outer}(y, \Psi, \mathcal{H}'')$. Further, by the choice of cuts, $\mathcal{H}'' \subseteq \mathcal{H}_{\text{next}}$ since \mathcal{H}'' is the support of Ψ . Hence, $\text{Outer}(x, \Pi, \mathcal{H}) \subseteq \text{Outer}(y, \Psi, \mathcal{H}_{\text{next}})$. \square

As a first step in the proof of Lemma 6.15, we show that Λ is optimal to $D_{\mathcal{H}}(G, c)$; note that Ψ was defined as the extremal dual solution to this problem.

Claim 6.16. Λ is an optimal solution to $D_{\mathcal{H}}(G, c)$.

Proof. It is sufficient to show that Λ is feasible to $D_{\mathcal{H}}(G, c)$. Optimality then follows by the complementary slackness property (C) between $z = y$ and Λ . We have $\Lambda(S) = 0$ for all $S \in \mathcal{H} \setminus \mathcal{L}$, and $\Lambda(S) \geq 0$ for all $S \in \mathcal{L}$. Hence we only need to prove that $\Lambda(S) \geq 0$ for all $S \in \mathcal{H}''$. This follows by Lemma 6.10 since every such S is an outer node at the beginning of the algorithm. \square

Let us study the image of Ψ in the contracted graph (G^*, c^*) , where this graph is defined with respect to the terminating solution $(\mathcal{L}, \mathcal{H}'', z, \Lambda)$. Let us define $\Theta : \mathcal{V}^* \rightarrow \mathbb{R}$ as follows:

$$\Theta(S) = \begin{cases} \Psi(S) - \Delta_{\Lambda, \Psi}(S) & \text{if } S \in \mathcal{V}^* \cap \mathcal{L}, \\ \Psi(S) & \text{if } S \in \mathcal{V}^* \setminus \mathcal{L}. \end{cases}$$

Note that $\mathcal{V}^* \cap \mathcal{L}$ is the set of maximal sets in \mathcal{L} .

Claim 6.17. Θ is an optimal solution to $D_0(G^*, c^*)$. Further, $\Theta(S) \leq 0$ holds for every $S \in \mathcal{V}^* \cap \mathcal{L}$.

Proof. We note that G^* arises from G by contracting the maximal sets in $\mathcal{L} \cup \mathcal{H}''$ w.r.t. Λ . We claim that Ψ is a $\mathcal{L} \cup \mathcal{H}''$ -positively-critical dual optimal solution to $D_{\mathcal{L} \cup \mathcal{H}''}(G, c)$. The $\mathcal{L} \cup \mathcal{H}''$ -positively-critical property is a consequence of Lemma 3.2. For optimality, first note that Ψ is by definition an optimal solution to $D_{\mathcal{H}' \cup \mathcal{H}''}(G, c)$. Further, Ψ is feasible to $D_{\mathcal{L} \cup \mathcal{H}''}(G, c)$, and the feasible region of this program is contained inside that of $D_{\mathcal{H}' \cup \mathcal{H}''}(G, c)$. Consequently, Ψ is optimal to $D_{\mathcal{L} \cup \mathcal{H}''}(G, c)$ as well.

First consider a set $S \in \mathcal{H}''$. Inside S , Lemma 5.8 implies that Π and Ψ are identical. Further, Π and Λ are also identical inside such sets, because the sets in \mathcal{H}'' are never unshrunk during the procedure and hence the dual values inside do not change. Hence contracting such a set w.r.t. Λ is the same as contracting it w.r.t. Ψ .

Consider now a maximal set $S \in \mathcal{L}$. By property (A), S is $(G_{\Lambda}, \mathcal{L} \cup \mathcal{H}'')$ -factor-critical and $\Lambda(T) > 0$ for every $T \subseteq S, T \in \mathcal{L}$. By property (C), we have $z(\delta(S)) = 1$. By Claim 6.16, Λ is an optimal (in particular, feasible) solution to $D_{\mathcal{H}}(G, c)$. Thus Lemma 5.1 is applicable and shows that Λ and Ψ are consistent inside S with $\Delta_{\Lambda, \Psi}(S) \geq 0$. By Lemma 5.2, we may contract such a set S w.r.t. Λ , and we obtain an optimal dual solution from Ψ by subtracting $\Delta_{\Lambda, \Psi}(S)$ from $\Psi(S)$ and leaving the values on all other sets unchanged.

Applying the above arguments one-by-one for all maximal members of $\mathcal{L} \cup \mathcal{H}''$ we can conclude that Θ is an optimum dual solution to $D_0(G^*, c^*)$. The second part follows since $\Psi(S) = 0$ and $\Delta_{\Lambda, \Psi}(S) \geq 0$ if $S \in \mathcal{L}$. \square

Claim 6.18. For every $S \in \mathcal{V} \cup \mathcal{H}$, $|\Lambda(S) - \Pi(S)| \leq |\Psi(S) - \Pi(S)|$ and equality holds only if $\Lambda(S) = \Psi(S)$.

Proof. The claim will follow by showing that for every $S \in \mathcal{V} \cup \mathcal{H}$, either $\Pi(S) \leq \Lambda(S) \leq \Psi(S)$ or $\Pi(S) \geq \Lambda(S) \geq \Psi(S)$.

First, if $\Lambda(S) > \Pi(S)$, then by Lemma 6.12, we have that $S \in \mathcal{V}^*$ and $S \in \mathcal{N}_o$. Consequently, by Lemma 6.13, $\Theta(S) \geq \Lambda(S)$. If $S \in \mathcal{L}$, then $0 \geq \Theta(S) \geq \Lambda(S)$ using Claim 6.17. Otherwise, $\Psi(S) = \Theta(S) \geq \Lambda(S) > \Pi(S)$.

If $\Lambda(S) < \Pi(S)$, then by Lemma 6.12, we have that either (1) $S \in \mathcal{H}' \setminus \mathcal{L}$, that is, $\Lambda(S) = 0$ and S was unshrunk or (2) $S \in \mathcal{N}_i$ or (3) $S \in \mathcal{V}^*$ and S is incident to an odd cycle C in $\text{supp}(z)$. Note that $S \notin \mathcal{H}''$ since all sets in \mathcal{H}'' are in \mathcal{N}_o and $\Pi(S) > \Lambda(S)$. In case (1), we have $\Psi(S) = 0 = \Lambda(S) < \Pi(S)$. In both cases (2) and (3), Lemma 6.13 gives $\Theta(S) \leq \Lambda(S)$. If $S \in \mathcal{L}$, then $\Psi(S) = 0 \leq \Lambda(S) < \Pi(S)$. If $S \in \mathcal{V} \cup (\mathcal{H}' \setminus \mathcal{L})$, then $\Psi(S) = \Theta(S) \leq \Lambda(S) \leq \Pi(S)$. \square

Proof. Proof of Lemma 6.15. The proof is straightforward by the above two claims: By Claim 6.16, Λ is optimal to $D_{\mathcal{H}}(G, c)$, and Claim 6.18 shows that $h(\Lambda, \Pi) \leq h(\Psi, \Pi)$, and equality can hold only if Λ and Ψ are identical. \square

As remarked above, this completes the proof of Lemma 6.14, and hence of Lemma 3.4.

7 Uniqueness.

In Section 5.2, we proved that if \mathcal{F} is a critical family, then there always exists an \mathcal{F} -positively-critical optimal solution (Corollary 5.11). This argument did not use uniqueness. Indeed, it will also be used to derive Lemma 2.5, showing that a perturbation of the original integer cost function satisfies our uniqueness assumption (\star) . We will need the following simple claim.

Claim 7.1. *For a graph $G = (V, E)$, let $a, b : E \rightarrow \mathbb{R}_+$ be two vectors on the edges with $a(\delta(v)) = b(\delta(v))$ for every $v \in V$. If a and b are not identical, then there exists an even length closed walk C such that for every odd edge $e \in C$, $a(e) > 0$ and for every even edge $e \in C$, $b(e) > 0$.*

Proof. Due to the degree constraints, $z = a - b$ satisfies $z(\delta(v)) = 0$ for every $v \in V$, and since a and b are not identical, z has nonzero components. If there is an edge $uv \in E$ with $z(uv) > 0$ then there must be another edge $uw \in E$ with $z(uw) < 0$. This implies the existence of an alternating even closed walk C where for every odd edge $e \in C$, $0 < z(e) = a(e) - b(e)$, and for every even edge, $0 > z(e) = a(e) - b(e)$. This proves the claim. \square

Proof of Lemma 2.5. Let \tilde{c} denote the perturbation of the integer cost $c : E \rightarrow \mathbb{Z}$. Consider a graph $G = (V, E)$, perturbed cost \tilde{c} and critical family \mathcal{F} where (\star) does not hold. Choose a counterexample with $|\mathcal{F}|$ minimal. Let x and y be two different optimal solutions to $P_{\mathcal{F}}(G, \tilde{c})$. Since \mathcal{F} is a critical family, by Corollary 5.11, there exists an \mathcal{F} -positively-critical dual optimal solution, say Π .

First, assume $\mathcal{F} = \emptyset$. Then x and y are both optimal solutions to the bipartite relaxation $P_0(G, \tilde{c})$. As they are not identical, Claim 7.1 gives an even closed walk C such that $x(e) > 0$ on every even edge and $y(e) > 0$ on every odd edge. Let γ_1 and γ_2 be the sum of edge costs on even and on odd edges of C , respectively. Then for some $\varepsilon > 0$, we could modify x by decreasing $x(e)$ by ε on even edges and increasing on odd edges, and y the other way around. These give two other optimal matchings \bar{x} and \bar{y} , with $\tilde{c}^T \bar{x} = \tilde{c}^T x + (\gamma_2 - \gamma_1)\varepsilon$ and $\tilde{c}^T \bar{y} = \tilde{c}^T y + (\gamma_1 - \gamma_2)\varepsilon$. Since \bar{x} and \bar{y} are both optimal, this gives $\gamma_1 = \gamma_2$. However, the fractional parts of γ_1 and γ_2 must be different according to the definition of the perturbation, giving a contradiction.

The case $\mathcal{F} \neq \emptyset$ is slightly more complicated. First we will replace x, y by another pair of non-identical optimal solutions a, b . We will be able to identify an alternating closed walk C in $\text{supp}(a) \cup \text{supp}(b)$ with the additional property that if C intersects $\delta(S)$ for some set $S \in \mathcal{F}$, then

it does so in exactly one even and one odd edge. The modifications \bar{a} and \bar{b} defined as above would again be feasible, implying $\gamma_1 = \gamma_2$.

We first claim that $\Pi(S) > 0$ must hold for all $S \in \mathcal{F}$. Indeed, if $\Pi(S) = 0$ for some $S \in \mathcal{F}$, then x and y would be two different optimal solutions to $P_{\mathcal{F} \setminus \{S\}}(G, \bar{c})$, contradicting the minimal choice of \mathcal{F} . Let $T \in \mathcal{F} \cup \{V\}$ be a smallest set with the property that $x(\delta(u, V \setminus T)) = y(\delta(u, V \setminus T))$ for every $u \in T$, but x and y are not identical inside T (here $\delta(u, S)$ is the set of edges having u as one endpoint and a vertex from S as the other). Note that V trivially satisfies this property and hence such a set exists. Let \mathcal{S} denote the collection of maximal sets S in \mathcal{F} such that $S \subsetneq T$ (\mathcal{S} could possibly be empty).

Inside each maximal set $S \in \mathcal{S}$, we modify x and y such that they are still both optimal and different from each other after the modification. Since $\Pi(S) > 0$, we have $x(\delta(S)) = 1$, $y(\delta(S)) = 1$ and S is Π -factor-critical. For $u \in S$, let M_u denote the Π -critical matching for u inside S . Let $\alpha_u := x(\delta(u, V \setminus S))$, $\alpha'_u := y(\delta(u, V \setminus S))$ and $w := \sum_{u \in S} \alpha_u \chi(M_u)$, $w' := \sum_{u \in S} \alpha'_u \chi(M_u)$. Consider the following modified solutions:

$$a(e) := \begin{cases} x(e) & \text{if } e \in \delta(S) \cup E[V \setminus S], \\ w(e) & \text{if } e \in E[S], \end{cases}$$

$$b(e) := \begin{cases} y(e) & \text{if } e \in \delta(S) \cup E[V \setminus S], \\ w'(e) & \text{if } e \in E[S], \end{cases}$$

We perform the above modifications inside every $S \in \mathcal{S}$, and denote by a and b the result of all these modifications. We claim that a and b are both optimal and non-identical. Optimality follows since both of them use only tight edges w.r.t. Π , and $a(\delta(Z)) = b(\delta(Z)) = 1$ still holds for every $Z \in \mathcal{F}$. Observe that, if for some S s.t. x and y are not identical inside S , and we have $\alpha_u = \alpha'_u$ for every $u \in S$, then S contradicts the minimality of T . This leaves two cases.

Case (1): There exists $S \in \mathcal{S}$ s.t. x and y are not identical inside S and $\alpha_u \neq \alpha'_u$ for some $u \in S$. Then we immediately have $a \neq b$ inside S .

Case (2): For all $S \in \mathcal{S}$, x is identical to y inside S . Since x is not identical to y inside T , we must have $a \neq b$ as only the edges inside the set S changed.

Claim 7.2. *Consider two edges $u_1v_1 \in \delta(S) \cap \text{supp}(a)$, $u_2v_2 \in \delta(S) \cap \text{supp}(b)$, $u_1, u_2 \in S$. Then there exists an even alternating path P_S inside S between u_1 and u_2 such that $a(e) > 0$ for every even edge and $b(e) > 0$ for every odd edge. Also, consider the extended path $P'_S = u_1v_1P_Su_2v_2$. If there exists a set $Z \subsetneq S, Z \in \mathcal{F}$ such that $V(P'_S) \cap Z \neq \emptyset$, then P'_S intersects $\delta(Z)$ in exactly one even and one odd edge.*

Proof. By the modification, $\text{supp}(a) \cap E[S]$ contains the Π -critical-matching M_{u_1} and $\text{supp}(b) \cap E[S]$ contains the Π -critical-matching M_{u_2} . Then the symmetric difference of M_{u_1} and M_{u_2} contains an $u_1 - u_2$ alternating path satisfying the requirements (by Lemma 2.2). \square

Let us now *contract* all sets in \mathcal{S} . We retain parallel edges in this contraction operation, as opposed to the contraction operation defined in Section 4. Let G' and \mathcal{F}' denote the resulting graph and laminar family respectively. Let a' and b' denote the projection of a and b on the contracted graph G' .

We claim that a' and b' are still not identical inside T' . Indeed, assume a' and b' are identical inside T' . We only modified x and y on the edge set $\cup_{S \in \mathcal{S}} E[S]$ to obtain a and b . Hence the contracted images must be identical: $x' = a'$ and $y' = b'$, and consequently, $x' = y'$. If Case (1)

holds, then x' cannot be identical to y' since the contraction operation retains parallel edges. Thus, we can only have Case (2). However, if Case (2) holds and $x' = y'$, then x and y are identical inside T contradicting the choice of T .

By the assumption on T , we also have $a'(\delta(u', V' \setminus T')) = b'(\delta(u', V' \setminus T'))$ for every $u' \in T'$. Consider the induced subgraph $G'[T']$ restricted to T' . The projection of a' and b' on this induced subgraph satisfy the hypothesis of Claim 7.1. This is because $a'(\delta(u')) = b'(\delta(u')) = 1$, and hence the fractional degrees inside T' must be the same. Hence, by Claim 7.1 we have an alternating closed walk C' inside T' with $a'(e) > 0$ on every even edge and $b'(e) > 0$ on every odd edge. Now, we can extend C' to an even alternating closed walk C in the original graph using the paths P_S as given by Claim 7.2. The resulting closed walk C will have the property that if there exists a set $Z \subsetneq S, Z \in \mathcal{F}$ such that $V(C) \cap Z \neq \emptyset$, then C intersects $\delta(Z)$ in exactly one even and one odd edge. \square

8 Open Questions.

Our initial motivation was to bound the number of iterations of the cutting plane method using the Padberg-Rao procedure. This question remains open and any analysis would have to deal with non-half-integral solutions.

Within our algorithm, Lemma 3.1 shows that it is sufficient to use positively-critical dual optimal solutions to maintain proper-half-integrality. Can we prove efficient convergence of our cutting plane algorithm using positively-critical dual optimal solutions (without using extremal dual solutions)? We believe that such a proof of convergence should depend on whether the following adversarial variation of Edmonds' algorithm for perfect matching is polynomial time. Suppose we run the Edmonds' perfect matching algorithm, but after every few primal/dual iterations, the adversary replaces the current dual solution with a different one, still satisfying complementary slackness with the (unchanged) primal solution, and having dual objective value at least as much as the previous one.

Given the encouraging results of this paper, it would be interesting to prove efficient convergence of the cutting plane method for other combinatorial polytopes. For example, one could try optimizing over the intersection of two matroid polytopes by a lazy-constraint approach—add violated rank constraints if the current LP optimum is not integral. Another direction could be to try this approach for optimizing over the subtour elimination polytope.

Acknowledgments.

We are grateful to William Cook and Richard Karp for their kind help and encouragement.

References

- [1] E. Balas. Intersection Cuts - A New Type of Cutting Planes for Integer Programming. *Operations Research*, 19(1):19–39, Jan 1971.
- [2] E. Balas. Disjunctive Programming. *Annals of Discrete Mathematics*, 5:3–51, 1979.
- [3] E. Balas, S. Ceria, and G. Cornuéjols. A lift-and-project cutting plane algorithm for mixed 0-1 programs. *Mathematical Programming*, 58(1):295–324, Jan 1993.

- [4] P. R. Bunch. *A simplex based primal-dual algorithm for the perfect b-matching problem - a study in combinatorial optimization algorithm in engineering*. PhD thesis, Purdue University, 1997.
- [5] K. Chandrasekaran, L. A. Végh, and S. S. Vempala. A half-integral algorithm for minimum weight matching. *Manuscript*, 2015.
- [6] V. Chvátal. Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete Mathematics*, 4(4):305–337, Apr 1973.
- [7] W. Cook, R. Kannan, and A. Schrijver. Chvátal closures for mixed integer programming problems. *Mathematical Programming*, 47(1):155–174, May 1990.
- [8] G. Cornuéjols and Y. Li. Elementary closures for integer programs. *Operations Research Letters*, 28(1):1–8, Feb 2001.
- [9] G. Dantzig, R. Fulkerson, and S. Johnson. Solution of a large-scale traveling-salesman problem. *Operations Research*, 2(4):393–410, Nov 1954.
- [10] J. Edmonds. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of Research of the National Bureau of Standards*, 69(1-2):125–130, 1965.
- [11] J. Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17(3):449–467, 1965.
- [12] M. Fischetti and A. Lodi. Optimizing over the first Chvátal closure. *Mathematical Programming*, 110(1):3–20, 2007.
- [13] R. Gomory. Outline of an algorithm for integer solutions to linear programs. *Bulletin of the American Mathematical Society*, 64(5):275–278, 1958.
- [14] R. Gomory. Solving linear programs in integers. In *Proceedings of Symposia in Applied Mathematics*, volume 10, pages 211–215, July 1960.
- [15] R. Gomory. An algorithm for integer solutions to linear programs. *Recent advances in mathematical programming*, page 269, 1963.
- [16] M. Grötschel and O. Holland. Solving matching problems with linear programming. *Mathematical Programming*, 33(3):243–259, 1985.
- [17] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer, 1988.
- [18] R. Kannan. Minkowski’s convex body theorem and integer programming. *Mathematics of Operations Research*, 12(3):415–440, 1987.
- [19] L. Lovász and M. Plummer. *Matching theory*. North Holland, 1986.
- [20] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal of Optimization*, 1:166–190, 1991.
- [21] D. L. Miller and J. F. Pekny. A staged primal-dual algorithm for perfect b-matching with edge capacities. *ORSA Journal on Computing*, 7(3):298–320, 1995.

- [22] G. Nemhauser and L. Wolsey. *Integer and Combinatorial Optimization*. Wiley-Interscience, 1999.
- [23] G. L. Nemhauser and L. A. Wolsey. A recursive procedure to generate all cuts for 0-1 mixed integer programs. *Mathematical Programming*, 46(1):379–390, Jan 1990.
- [24] M. Padberg and M. Rao. Odd minimum cut-sets and b -matchings. *Mathematics of Operations Research*, 7(1):67–80, 1982.
- [25] A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & sons, 1998.
- [26] A. Schrijver. *Combinatorial optimization*. Springer, 2003.
- [27] H. D. Sherali and W. P. Adams. A hierarchy of relaxations and convex hull characterizations for mixed-integer zero-one programming problems. *Discrete Applied Mathematics*, 52(1):83–106, Jul 1994.
- [28] M. Trick. *Networks with additional structured constraints*. PhD thesis, Georgia Institute of Technology, 1987.