The Cutting Plane Method is Polynomial for Perfect Matchings*

Karthekeyan Chandrasekaran[1], László A. Végh[2], and Santosh S. Vempala[1]

[1]College of Computing, Georgia Institute of Technology
[2]Department of Management, London School of Economics
{*karthe, vempala*}*@cc.gatech.edu, l.vegh@lse.ac.uk*

Abstract

The cutting plane approach to optimal matchings has been discussed by several authors over the past decades [21, 14, 17, 25, 10], and its convergence has been an open question. We prove that the cutting plane approach using Edmonds' blossom inequalities converges in polynomial time for the minimum-cost perfect matching problem. Our main insight is an LP-based method to retain/drop candidate cutting planes. This cut retention procedure leads to a sequence of intermediate linear programs with a linear number of constraints whose optima are half-integral and supported by a disjoint union of odd cycles and edges. This structural property of the optima is instrumental in finding violated blossom inequalities (cuts) in linear time. Further, the number of cycles in the support of the half-integral optima acts as a potential function to show efficient convergence to an integral solution.

1 Introduction

Integer programming is a powerful and widely used approach to modeling and solving discrete optimization problems [19, 22]. Not surprisingly, it is NP-complete and the fastest known algorithms are exponential in the number of variables (roughly $n^{O(n)}$ [16]). 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 [7] and pioneered by Gomory [11, 12, 13]. 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.

*This work was supported in part by NSF award AF0915903. This work was done while the second author was affiliated with the College of Computing, Georgia Institute of Technology, supported by NSF Grant CCF-0914732.

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 the cutting plane method implemented using his cut-generation procedure converges to an integral solution [13]. There is a rich theory on the choice of cutting planes, both in general and for specific problems of interest. This theory includes interesting families of cutting planes with efficient cut-generation procedures [11, 1, 4, 2, 5, 20, 3, 18, 24], valid inequalities, closure properties and a classification of the strength of inequalities based on their *rank* with respect to cut-generating procedures [6] (e.g., the Chvátal-Gomory rank [4]), 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 [13]. It is possible that such a bound can be significantly improved for IPs with small Chvátal-Gomory rank [10]. 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 *(a)* has a polynomial-size IP-description (the LP-relaxation is polynomial-size), and *(b)* the convex-hull of integer solutions has a polynomial-time separation oracle. Such a problem admits a polynomial-time algorithm via the Ellipsoid method [15]. 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.*

A polyhedral characterization of the matching problem was discovered by Edmonds [8]: Basic solutions of the following linear program (extreme points of the polytope) correspond to perfect matchings of the graph.

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

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 [4]. 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 [21]. This suggests a natural cutting plane algorithm (Figure 1), proposed by Padberg and Rao [21] and discussed by Lovász and Plummer in their classic book on matching theory [17]. Experimental evidence suggesting that this method converges quickly was given by Grötschel and Holland [14], by Trick [25], and by Fischetti and Lodi [10]. 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.

The known polynomial-time algorithms for minimum-cost perfect matching are variants of Edmonds' weighted matching algorithm [8]. It is perhaps tempting to interpret the latter as a cutting plane algorithm, by adding cuts corresponding to the shrunk sets in the iterations of Edmonds' algorithm. However, there is no correspondence between the solution x of the LP given by non-negativity and degree constraints and a family $\mathcal{F}$ of blossom inequalities, and the partial matching M in the iteration of Edmonds' algorithm when $\mathcal{F}$ is the set of shrunk nodes. In particular, the next odd set S shrunk by Edmonds' algorithm might not even be a cut for x (i.e., $x(\delta(S)) \geq 1$). It is even possible, that the bipartite relaxation already has an integer optimal solution, whereas Edmonds' algorithm proceeds by shrinking and unshrinking a long sequence of odd sets.

1. Start with the bipartite relaxation. nodes.
2. While the current solution is fractional,
   (a) Find a violated blossom inequality and add it to the LP.
   (b) Solve the new LP.

Figure 1: Cutting plane method for matchings

The bipartite relaxation has the nice property that any basic solution is half-integral and its support is a disjoint union of edges and 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 [14, 25] 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 [21]. Thus, it is natural to wonder if there is a choice of cutting planes that maintains half-integrality of intermediate LP optimal solutions.

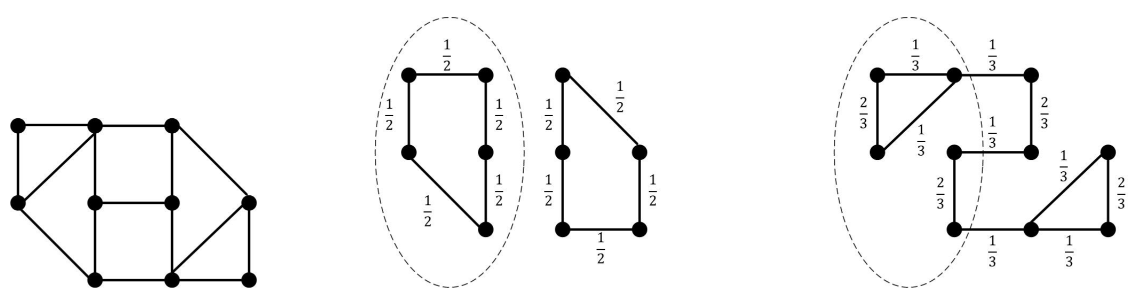

Graph G | The starting optimum x_0 and the cut to be imposed | Basic feasible solution obtained after imposing the cut

Figure 2: Counterexample to the half-integrality conjecture. All edge costs are one.

At first sight, maintaining half-integrality seems to be impossible. Figure 2 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 non half-integral basic feasible solution. Observe however, that in the example, the bipartite relaxation also has an integer optimal solution. The problem here is 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 the simple deterministic perturbation of the integer cost function, which increases the input size polynomially. We observe that this perturbation is only a first step towards maintaining half-integrality of intermediate LP optima. More careful cut retention and cut addition procedures are needed 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 $X, Y \in \mathcal{F}$, one of $X \cap Y = \emptyset$, $X \subseteq Y$, $Y \subseteq X$ holds.

Next we define a perturbation to the cost function that will help avoid some degeneracies. Given an integer cost function $c : E \to \mathbb{Z}$ 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 $c : E \to \mathbb{Z}$ 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 supported by a disjoint union of edges and odd cycles and

(iii) the total number of iterations to arrive at a minimum $\tilde{c}$-cost perfect matching is $O(n \log n)$.

Moreover, the collection of blossom inequalities used at each step can be identified by solving an LP of the same size as the current LP. 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. 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.

1.2 Cut selection via dual values

The main difficulty that we have to overcome is the following: ensuring unique optimal solutions does not suffice to maintain half-integrality of optimal solutions upon adding any sequence of blossom inequalities. In fact, even a laminar family of blossom inequalities is insufficient to guarantee the nice structural property on the intermediate LP optimum. Thus a careful choice is to be made while choosing new cuts and it is also crucial that we eliminate certain older ones that are no longer useful.

At any iteration, inequalities that are tight for the current optimal solution are natural candidates for retaining 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. 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 (later called *positively-critical dual* in this paper) that can be computed either by solving a single LP of the same complexity or combinatorially. Moreover, we also ensure that the set of cuts imposed in any LP are laminar and correspond to blossom inequalities.

Eliminating cutting planes that have zero dual values in any later iteration is common in most implementations of the cutting plane algorithm; although this is done mainly to keep the number of inequalities from blowing up, another justification is that a cut with zero dual value is not a facet contributing to the current LP optimum.

1.3 Algorithm C-P-Matching

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

$$\begin{aligned} \min \sum_{uv \in E} c(uv)x(uv) & \qquad (P_0(G,c)) \\ x(\delta(u)) = 1 \quad \forall u \in V & \\ x \geq 0 & \end{aligned}$$

$$\begin{aligned} \max \sum_{u \in V} \pi(u) & \qquad (D_0(G,c)) \\ \pi(u) + \pi(v) \leq c(uv) \quad \forall uv \in E & \end{aligned}$$

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 $\mathrm{supp}(x)$ is a disjoint union of edges and odd cycles. The bipartite relaxation of any graph has the following well-known property.

Proposition 1.2. *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.

$$\begin{aligned} \min \sum_{uv \in E} c(uv)x(uv) & \qquad (P_{\mathcal{F}}(G,c)) \\ x(\delta(u)) = 1 \quad \forall u \in \mathcal{V} & \\ x(\delta(S)) \geq 1 \quad \forall S \in \mathcal{F} & \\ x \geq 0 & \end{aligned}$$

$$\begin{aligned} \max \sum_{S \in \mathcal{V} \cup \mathcal{F}} \Pi(S) & \qquad (D_{\mathcal{F}}(G,c)) \\ \sum_{S \in \mathcal{V} \cup \mathcal{F} : uv \in \delta(S)} \Pi(S) \leq c(uv) \quad \forall uv \in E & \\ \Pi(S) \geq 0 \quad \forall S \in \mathcal{F} & \end{aligned}$$

For $\mathcal{F} = \emptyset$, $P_{\mathcal{F}}(G, c)$ is identical to $P_0(G, c)$, whereas for $\mathcal{F} = \mathcal{O}$, it 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).

The cutting plane implementation that we propose is shown in Figure 3. From the previous 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. In the first iteration, there is no need to solve the dual LP as $\mathcal{F}$ will be empty.

1. Let c denote the cost function on edges after perturbation (i.e., after ordering the edges arbitrarily and increasing the cost of edge i by $1/2^i$).
2. **Starting LP.** Let $\mathcal{F}$ be the empty set, so that the starting LP, $P_{\mathcal{F}}(G,c)$, is the bipartite relaxation and the starting dual Γ is identically zero.
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.

Figure 3: Algorithm C-P-Matching

1.4 Overview of the analysis

Our analysis to show half-integrality is based on the intimate relationship that matchings have with factor-critical graphs (deleting any node leaves the graph with a perfect matching): for example, contracted sets are factor-critical both in the unweighted [9] and weighted [8] matching algorithms by Edmonds. We define a notion of factor-criticality for weighted graphs that also takes a laminar odd family $\mathcal{F}$ into account. This notion will be crucial to guarantee half-integrality of intermediate solutions.

We use the number of odd cycles in the support of an optimal half-integral solution as a potential function to show convergence. 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 odd(.) is the number of odd cycles in the support. 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. Since the maximum size of a laminar family of nontrivial odd sets is $n/2$, we get a bound of $O(n \log n)$ on the number of iterations.

The proof of the potential function behavior is quite intricate. It proceeds by designing a *half-integral* version of Edmonds primal-dual algorithm for minimum-cost perfect matching, and arguing that the optimal solution to the extremal dual LP must correspond to the one found by this primal-dual algorithm. We emphasize that this algorithm is used only in the analysis. Nevertheless, it is rather remarkable that even for analyzing the cutting plane approach, comparison with a new extension of Edmonds' classic algorithm provides the answer.

2 Factor-critical sets

In what follows, we formulate a notion of factor-critical sets and factor-critical duals, that play a central role in the analysis of our algorithm and are extensions of concepts central to the analysis of Edmonds' algorithm.

Let $H = (V, E)$ be a graph and $\mathcal{F}$ be a laminar family of 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 $S \subseteq V$, we call a set M of edges to be an *$(S, \mathcal{F})$-perfect-matching* if it is an $\mathcal{F}$-matching covering precisely the vertex set S.

A set $S \in \mathcal{F}$ is defined to be *$(H, \mathcal{F})$-factor-critical* or *$\mathcal{F}$-factor-critical in 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. For simplicity we will say that a set $S \in \mathcal{F}$ is $(\Pi, \mathcal{F})$-factor-critical if it is $(G_\Pi, \mathcal{F})$-factor critical, i.e., S is $\mathcal{F}$-factor-critical in G_Π. For a vertex $u \in S$, corresponding matching M_u is called the Π-*critical-matching* for u. If $\mathcal{F}$ is clear from the context, then we simply say S is Π-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 $(\Pi, \mathcal{F})$-factor-critical, and $\Pi(T) > 0$ for every non-maximal set T of $\mathcal{F}$. 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 crucial 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.* □

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 $(\Pi, \mathcal{F})$-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 $(\Pi, \mathcal{F})$-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 an even cycle C. Then the symmetric difference $M_u \Delta C$ is also a Π-critical matching for u.

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

(ii) Let $M' = M_u \Delta C$. First observe that since C is an even cycle, $u, v \notin V(C)$. Hence M' is a perfect matching on $S \setminus \{u\}$ using only tight edges w.r.t. Π. It remains to show that $|M' \cap \delta(T)| \le 1$ 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 $\gamma_u, \gamma_v \le 1$. 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 $|M' \cap \delta(T)| = |M_u \cap \delta(T)| - \gamma_u + \gamma_v$. □

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 $c : E \to \mathbb{R}$ satisfies:

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

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

Lemma 2.3. *Let $c : E \to \mathbb{Z}$ be an integer cost function, and $\tilde{c}$ be its perturbation. Then $\tilde{c}$ satisfies the uniqueness property (*).*

3 Analysis outline and proof of the main theorem

The proof of our main theorem is established in two parts. In the first part, we show that half-integrality of the intermediate primal optimum solutions is guaranteed by the existence of an $\mathcal{F}$-positively-critical dual optimal solution to $D_{\mathcal{F}}(G, c)$.

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 there exists an $\mathcal{F}$-positively-critical dual optimal solution, then x is proper-half-integral.*

Lemma 3.1 is shown using a basic contraction operation. Given Π, an $\mathcal{F}$-positively-critical dual optimal solution for the laminar odd family $\mathcal{F}$, contracting every set $S \in \mathcal{F}$ with $\Pi(S) > 0$ preserves primal and dual optimal solutions (similar to Edmonds' primal-dual algorithm). This is shown in Lemma 4.1. Moreover, if we had 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 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 in 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 dual solution. Then a Γ-extremal dual optimal solution Π is an $\mathcal{F}$-positively-critical dual optimal solution. 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 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, and let $\Delta = \max_{u \in S}(\Gamma_S(u) - \Pi_S(u))$. Then, every edge in $\text{supp}(x) \cap \delta(S)$ is incident to some node $u \in S$ such that $\Gamma_S(u) - \Pi_S(u) = \Delta$. Also (Lemma 5.8), if $S \in \mathcal{F}$ is both Γ- and Π-factor-critical, then Γ and Π are identical inside S.

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 (*). 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 cannot remain the same and 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 (*) 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* algorithm, a variation of Edmonds' primal-dual weighted matching algorithm that we design for this purpose.

We briefly describe this next, assuming familiarity with Edmonds' algorithm [8]. Our argument needs one phase of this algorithm and this is what we analyze in detail in Section 6.1. Although we introduce it for analysis, we note that the algorithm can be extended to a strongly-polynomial combinatorial algorithm for minimum-cost perfect matching. Unlike Edmonds' algorithm, which maintains an integral matching and extends the matching to cover all vertices, our algorithm maintains a proper half-integral solution.

The half-integral matching algorithm starts from a partial matching x in G, leaving a set W of nodes exposed, and a dual Π whose support is a laminar family $\mathcal{V} \cup \mathcal{F}$ with $\mathcal{F} \subseteq \mathcal{O}$; x and Π satisfy primal-dual slackness conditions. The algorithm transforms x to a proper-half-integral perfect matching and Π to a dual solution with support contained in $\mathcal{V} \cup \mathcal{K}$, satisfying complementary slackness. Before outlining the algorithm, we show how we apply it to prove Lemmas 3.3 and 3.4.

Let us consider two consecutive primal solutions x_i and x_{i+1} in the cutting plane algorithm, with duals Π_i and Π_{i+1}. We contract every set $S \in \mathcal{O}$ with $\Pi_{i+1}(S) > 0$; let $\hat{G}$ be the resulting graph. By Lemma 3.1 the image x'_{i+1} of x_{i+1} is the unique optimal solution to the bipartite relaxation in $\hat{G}$. The image x'_i of x_i is proper-half-integral in $\hat{G}$ with some exposed nodes W; let Π'_i be the image of Π_i. Every exposed node in W corresponds to a cycle in $\text{supp}(x_i)$. We start in $\hat{G}$ with the solutions x'_i and Π'_i, and we prove that it must terminate with the primal solution x'_{i+1}. The analysis of the half-integral matching algorithm reveals that the total number of exposed nodes and odd cycles does not increase; this will imply Lemma 3.3.

To prove Lemma 3.4, we show that if the number of cycles does not decrease between phases i and $i+1$, then the algorithm also terminates with the extremal dual optimal solution Π'_{i+1}. This enables us to couple the performance of Half-integral Matching between phases i and $i+1$ and between $i+1$ and $i+2$: the (alternating forest) structure built in the former iteration carries over to the latter one. As a consequence, all cuts added in iteration i will be imposed in all subsequent phases until the number of odd cycles decreases.

Let us now turn to the description of the Half-integral Matching algorithm. In every step of the algorithm, we maintain z to be a proper-half-integral partial matching with exposed set $T \subseteq W$, Λ to be a dual solution satisfying complementary slackness with z, and the support of Λ is $\mathcal{V} \cup \mathcal{L}$ for some $\mathcal{L} \subseteq \mathcal{F}$. We start with $z = x$, $T = W$, $\Lambda = \Pi$ and $\mathcal{L} = \mathcal{F}$. We work on the graph G^* resulting from the contraction of the maximal sets of $\mathcal{L}$.

By changing the dual Λ, we grow an alternating forest of tight edges in G^*, using only edges e with $z(e) = 0$ or $z(e) = 1$. The forest is rooted at the exposed set of nodes T. The solution z will be changed according to the following three scenarios. *(a)* If we find an alternating path between two exposed nodes, we change z by alternating along this path as in Edmonds' algorithm. *(b)* If we find an alternating path P from an exposed node to a 1/2-cycle C in $\text{supp}(z)$, we change z by alternating along P, and replacing it by a blossom (an alternating odd cycle) on C. *(c)* If we find an alternating path P from an exposed node to a blossom C, then we change z by alternating along P and replacing the blossom by a 1/2-cycle on C. If none of these cases apply, we change the dual Λ in order to extend the alternating forest. If $\Lambda(S)$ decreases to 0 for some $S \in \mathcal{L}$, then we remove S from $\mathcal{L}$ and unshrink it in G^*.

The modifications are illustrated on Figure 4. Note that in case *(c)*, Edmonds' algorithm would instead contract C. In contrast, we do not perform any contractions, but allow 1/2-cycles in the half-integral matching algorithm. For the special starting solution $x \equiv 0$, $\Pi \equiv 0$, our half-integral matching algorithm will return a proper-half-integral optimum to the bipartite relaxation.

The main theorem can now be proved using the above lemmas.

Proof of Theorem 1.1. We use Algorithm C-P-Matching given in Figure 3 for a perturbed cost function. By Lemma 2.3, this satisfies (*). Let i denote the index of the iteration. We prove by

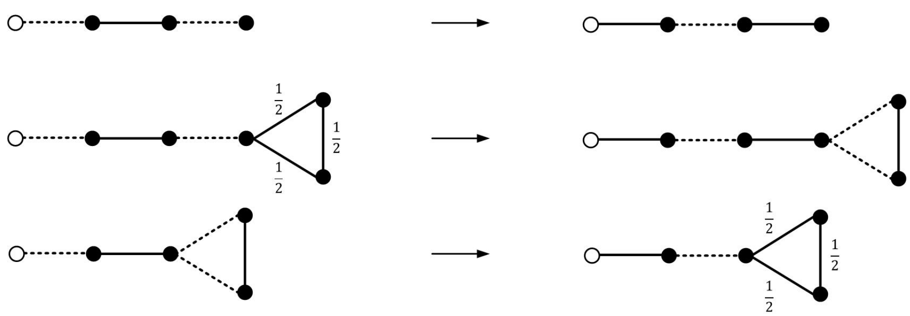

Figure 4: The possible modifications in the Half-integral Matching algorithm.

induction on i that every intermediate solution x_i is proper-half-integral and (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. Between iterations i and j, the set $\mathcal{H}''_k$ always contains ℓ cuts, 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 $j - i \le n/2\ell$. Consequently, the number of cycles must decrease from ℓ to $\ell - 1$ within $n/2\ell$ iterations. Since $\text{odd}(x_0) \le n/3$, 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 $c(M') \le c(M) - 1$. 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 $\tilde{c}(M') < c(M') + 1 \le c(M) < \tilde{c}(M)$, 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, Π be a feasible solution to $D_{\mathcal{F}}(G, c)$, and let $S \in \mathcal{F}$ be a $(\Pi, \mathcal{F})$-factor-critical set. Let us define

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

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. 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 $u \in S$, $v \in V - S$ add an edge sv. Let us define the image Π' of Π to be $\Pi'(T') = \Pi(T)$ for every $T' \in \mathcal{V}' \cup \mathcal{F}'$ and the image x' of x to be

$x'(u'v') = x(uv)$. Define the new edge costs

$$c'(u'v') = \begin{cases} c(uv) & \text{if } uv \in E[V \setminus S], \\ c(uv) - \Pi_S(u) & \text{if } u \in S,\, v \in V \setminus S. \end{cases}$$

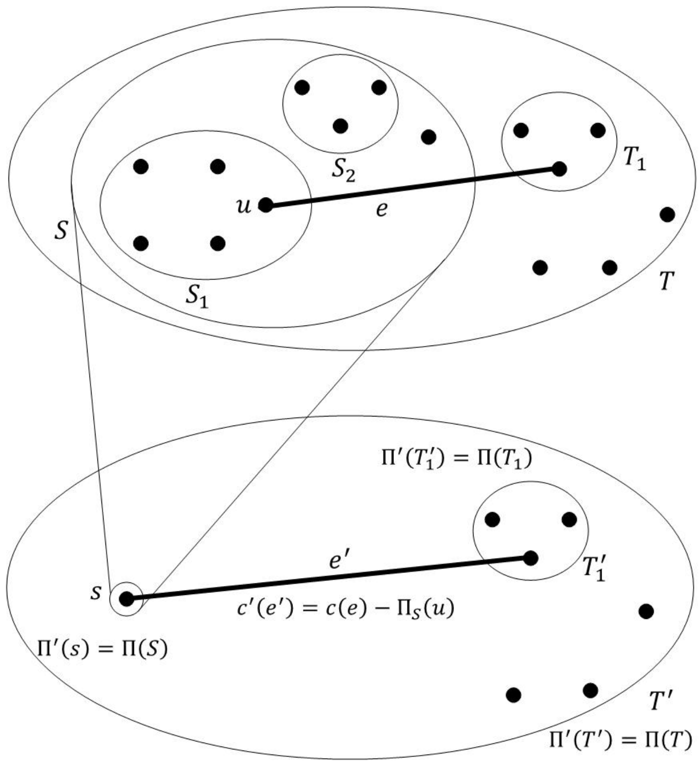

Figure 5: Contraction operation: image of dual and new cost function

Lemma 4.1. *Let $\mathcal{F}$ be a laminar odd family, x be an optimal solution to $P_{\mathcal{F}}(G,c)$, Π be a feasible solution to $D_{\mathcal{F}}(G,c)$. Let $S \in \mathcal{F}$ be a $(\Pi,\mathcal{F})$-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 and 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 $(\Pi,\mathcal{F})$-factor-critical, then its image T' is $(\Pi',\mathcal{F}')$-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) *If x is the unique optimum to $P_{\mathcal{F}}(G,c)$, and Π is an optimal solution to $D_{\mathcal{F}}(G,c)$, then x' is the unique optimum to $P_{\mathcal{F}'}(G',c')$. Moreover, x' is proper-half-integral if and only if x is proper-half-integral. Further, assume C' is an odd cycle in supp(x') and let T be the pre-image of $V(C')$ in G. Then, supp(x) inside T consists of an odd cycle and matching edges.*

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

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

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

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

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 $(\Pi,\mathcal{F})$-factor-critical with $T \setminus S \neq \emptyset$. It is sufficient to verify that T' is $(\Pi',\mathcal{F}')$-factor-critical whenever T contains S. Let $u' \in T'$ be the image of $u \in V$, and consider the projection M' of the Π-critical-matching M_u. The support of this projection is contained in the tight edges with respect to Π'. Let $Z' \subsetneq T'$, $Z' \in \mathcal{F}'$ and let Z be the pre-image of Z'. If $u' \neq s$, then $|M' \cap \delta(Z')| = |M_u \cap \delta(Z)| \le 1$ 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, $|M_u \cap \delta(Z)| \le 1$ implies $|M' \cap \delta(Z')| \le 1$ 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 $x'(\delta(T')) \ge 1$ 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 y' to $P_{\mathcal{F}'}(G', c')$. Let M_u be the Π-critical matching for u in S. Define $\alpha_u = y'(\delta(u))$ for every $u \in S$, i.e., α_u is the total y' value on the images of edges $uv \in E$ with $v \in V - S$. Note that $\sum_{u\in S} \alpha_u = y'(\delta(S)) = 1$. Take $w = \sum_{u\in S} \alpha_u M_u$ and

$$y(uv) = \begin{cases} y'(u'v') & \text{if } uv \in E \setminus E[S], \\ w(uv) & \text{if } uv \in E[S]. \end{cases}$$

Then, y is a feasible solution to $P_{\mathcal{F}}(G,c)$ and y satisfies complementary slackness with Π. Hence, y is an optimum to $P_{\mathcal{F}}(G,c)$ and thus by uniqueness, we get $y = x$. Consequently, $y' = x'$.

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, $w = M_u$ is a perfect matching on $S - u$. 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 $w = \frac{1}{2}(M_{u_1} + M_{u_2})$. The uniqueness of x implies the uniqueness of both M_{u_1} and M_{u_2}. Then by Lemma 2.2(ii), the symmetric difference of M_{u_1} and M_{u_2} may not contain any even cycles. Hence, 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 supp(x') and let T be the pre-image of the set $V(C')$ in G. Then, 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 $x(e) = 1$ $\forall\, e \in M$.* □

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. Finally, if C' is an odd cycle in $\text{supp}(x')$, then Claim 4.2 provides the required structure for x inside T. then the same argument as above shows that the pre-image of C' is an even cycle C in the support of x, a contradiction. □

Corollary 4.3. *Assume x is the optimal solution to $P_{\mathcal{F}}(G, c)$ and there exists an $\mathcal{F}$-positively-critical dual optimum 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 optimum to $P_{\mathcal{F}}(G, c)$, then $\hat{x}$ is the unique optimum 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 optimum, 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 optimum 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. □

5 Structure of dual solutions

In this section, we show two properties about positively-critical dual optimal solutions – (1) an optimum Ψ to $D_{\mathcal{F}}(G, c)$ can be transformed into an $\mathcal{F}$-positively-critical dual optimum (Section 5.2) if $\mathcal{F}$ is a critical family 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 assume uniqueness 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 $\Psi_S(u) := \sum_{T \in \mathcal{V} \cup \mathcal{F}: T \subsetneq S, u \in T} \Psi(T)$ for $u \in S$. Let us now define

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

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 section is to prove the following lemma.

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

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 $\mathcal{F} \subseteq \mathcal{O}$ be a critical family, with Π being an $\mathcal{F}$-critical dual solution and let Ψ be an optimal solution to $D_{\mathcal{F}}(G,c)$. Let x be an optimal solution to $P_{\mathcal{F}}(G,c)$. Suppose Ψ is consistent with Π inside $S \in \mathcal{F}$ and $x(\delta(S)) = 1$. Let $G', c', \mathcal{F}'$ denote the graph, costs and laminar family obtained by contraction. Then the image x' of x is an optimum to $P_{\mathcal{F}'}(G',c')$, and Ψ' (as defined above) is an optimum 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 for every edge uv such that $u, v \neq s$ since Ψ is a feasible solution for $D_{\mathcal{F}}(G,c)$. Consider an edge $uv \in E(G)$, $u \in S$. Let $sv \in E(G')$ be the image of uv in G', and let $\Delta = \Delta_{\Pi,\Psi}(S)$.

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

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

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

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 the pre-image of sv for some $u \in S$. Then the tightness of the constraint follows since both the inequalities in (1) are tight – the first inequality is tight since uv is tight w.r.t. Ψ, and the second is tight since $\Pi_S(u) - \Psi_S(u) = \Delta(S)$ by the consistency property. Finally, if $\Psi'(T') > 0$ for some $T' \in \mathcal{F}'$, then $\Psi(T) > 0$ and hence $x(\delta(T)) = 1$ and therefore, $x'(\delta(T')) = 1$. □

Lemma 5.3. *Let $\mathcal{F}$ be a critical family with Π being an $\mathcal{F}$-critical dual, and let x be an optimal solution to $P_{\mathcal{F}}(G,c)$. If $x(\delta(S)) = 1$ for some $S \in \mathcal{F}$, then $x(\delta(T)) = 1$ for every $T \subsetneq S$, $T \in \mathcal{F}$, and all edges in $supp(x) \cap E[S]$ are tight w.r.t. Π.*

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.

$$\min \sum_{uv\in E[T]} c(uv)z(uv) \qquad (P_{\mathcal{F}}[S]) \qquad\qquad \max \sum_{T\subsetneq S, T\in \mathcal{V}\cup\mathcal{F}} (1-\alpha(T))\Gamma(T) \qquad (D_{\mathcal{F}}[S])$$

$$z(\delta(u)) = 1-\alpha_u \quad \forall u \in S \qquad\qquad \sum_{\substack{T\subsetneq S, T\in\mathcal{V}\cup\mathcal{F}\\ uv\in\delta(T)}} \Gamma(T) \le c(uv) \quad \forall uv \in E[S]$$

$$z(\delta(T)) \ge 1-\alpha(T) \quad \forall T \subsetneq S, T\in\mathcal{F}$$

$$z(uv) \ge 0 \ \forall \ uv \in E[S] \qquad\qquad \Gamma(Z) \ge 0 \quad \forall Z \subsetneq T, Z \in F$$

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 $T\in\mathcal{F}$, $T\setminus S \ne \emptyset$, then $x^z(\delta(T)) = x(T) \ge 1$. If $T\subseteq S$, then $x^z(\delta(T)) = z(\delta(T)) + x(\delta(T, V\setminus S)) \ge 1-\alpha(T)+\alpha(T) = 1$.

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

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 $D_{\mathcal{F}}[S]$. and $w(\delta(T)) = 1-\alpha(T)$ for every $T\subsetneq S$, $T\in\mathcal{F}$.*

Proof. Since Π is an $\mathcal{F}$-critical dual and $S\in F$, we have a Π-critical-matching M_u inside S for each $u\in S$. Let $w = \sum_{u\in S}\alpha_u M_u$. The claim follows by showing that w is feasible to $P_{\mathcal{F}}[S]$ and that w and $\bar\Pi$ satisfy complementary slackness. w and Γ are primal and dual feasible solutions satisfying complementary slackness, and hence both are optimal.

The degree constraint $w(\delta(u)) = 1-\alpha_u$ is straightforward and using Lemma 2.2, if $T\subsetneq S$, $T\in\mathcal{F}$, then $w(\delta(T)) = \sum_{u\in S\setminus T}\alpha_u = 1-\alpha(T)$. 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 $w(\delta(T)) = 1-\alpha(T)$ as verified above. Thus, all odd set constraints are tight in the primal. □

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 definition of the $\mathcal{F}$-critical-property, every non-maximal set T of $\mathcal{F}$ satisfies $\Pi(T) > 0$, and in particular, every set $T\subsetneq S$, $T\in\mathcal{F}$ satisfies $\bar\Pi(T) = \Pi(T) > 0$. Thus, complementary slackness gives $x(\delta(T)) = 1$.

□

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 $(\Lambda,\mathcal{F})$-factor-critical. Then $\Delta_{\Lambda,\Gamma}(S) = \max_{u\in S} |\Lambda_S(u) - \Gamma_S(u)|$. Further, if $\Delta_{\Lambda,\Gamma}(S) > 0$, define*

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

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

Proof. Let $\Delta = \max_{u\in S} |\Lambda_S(u) - \Gamma_S(u)|$, and define the sets A^- and A^+ with Δ instead of $\Delta_{\Lambda,\Gamma}(S)$. Since S is $(\Lambda,\mathcal{F})$-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 $\Gamma(u) + \Gamma(v) \le c(uv)$ 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, $|A^-| > |A^+|$. □

Proof of Lemma 5.1. We prove this by induction on $|V|$, and subject to this, on $|S|$. Let S be a $(\Pi,\mathcal{F})$-factor-critical set. First, consider the case when S is an inclusion-wise minimal set. Then, $\Pi_S(u) = \Pi(u)$, $\Psi_S(u) = \Psi(u)$ for every $u \in S$. By Claim 5.6, we have that $\Delta := \Delta_{\Pi,\Psi}(S) \ge 0$. 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 $\Pi(u) + \Pi(v) \le c(uv)$. Hence, if $u \in A^-$, then $v \in A^+$. Consequently, we have

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

Thus, we should have equality throughout. Hence, $x(\delta(A^-, V \setminus S)) = 1$. This precisely means that Ψ is consistent with Π inside S. $A^+ = \emptyset$ implies $|A^-| = 1$ and hence $|T| = 1$, a contradiction that T is a non-trivial odd set. This completes the proof of the base case.

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

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 optimum to $P_{\mathcal{F}'}(G',c')$ and $D_{\mathcal{F}'}(G',c')$ 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, $\Delta' \ge 0$. The following claim verifies consistency inside S and thus completes the proof. □

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

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

Assume now that there exists a $uv \in \delta(S) \cap \mathrm{supp}(x)$. 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.sets.

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 $(\Lambda, \mathcal{F})$-factor-critical and $(\Gamma, \mathcal{F})$-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 to see that T is also $(\Lambda, \mathcal{F})$-factor-critical and $(\Gamma, \mathcal{F})$-factor-critical by definition. 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 optimum solution to $D_{\mathcal{F}}(G, c)$. Our goal is to satisfy the property that for every $S \in F$, if $\Psi(S) > 0$, then Ψ and Π are identical inside S. By Lemma 5.8, it is equivalent to showing that Ψ is $\mathcal{F}$-positively-critical. We modify Ψ by the algorithm shown in Figure 6.

The correctness of the algorithm follows by showing that 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 optimum solution to $D_{\mathcal{F}}(G, c)$. Suppose we consider a maximal set S such that Π and Ψ are not identical inside S, take $\lambda = \min\{1, \Psi(S)/\Delta_{\Pi,\Psi}(S)\}$ if $\Delta_{\Pi,\Psi}(S) > 0$ and $\lambda = 1$ if $\Delta_{\Pi,\Psi}(S) = 0$ 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 Ψ 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:

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 $\bar{\Psi}$.

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

2. **Return** $\mathcal{F}$-positively-critical dual optimum Ψ.

Figure 6: Algorithm Positively-critical-dual-opt

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. $x(\delta(T)) = 1$.

The solution $\bar{\Psi}$ might have positive components on some sets $T \subsetneq S, T \in \mathcal{F}$ where $\Pi(T) > 0$. However, $x(\delta(T)) = 1$ for all such sets by Lemma 5.3. 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. □

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

Proof.

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

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). □

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 Figure 6 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. S' is processed after S and $\Psi(S)$ changed during the iteration when S' was processed. 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'. □

5.3 Extremal dual solutions

In this section, we prove Lemma 3.2. 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.

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

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 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 $\tau(S) \ge -\Delta\lambda$.

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

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

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 $D^*_{\mathcal{F}}$. □

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 $\mathrm{supp}(x)$. 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. Lemma 4.1(ii).

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

Let $u \in \hat{C}$ be an arbitrary node; we will construct the Π-critical matching $\hat{M}_u$. Let us contract all sets $S_1, \ldots, S_\ell$ to nodes $s_1, \ldots, s_\ell$ w.r.t. Π. We know by Lemma 4.1(iii) that the image x' of x is proper-half-integral and that the cycle C projects to a cycle in x'. The proper-half-integral property guarantees that C is contracted into an odd cycle C' in $\mathrm{supp}(x')$. Let u' be the image of u; the pre-image of every edge in C' is a tight edge w.r.t. Π in the original graph since Π is an optimum to the dual problem $D_{\mathcal{F}}(G, c)$. Since C' is an odd cycle, there is a perfect matching $M'_{u'}$ that covers every node in $V(C') \setminus \{u'\}$ using the edges in C'.

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. 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. □

6 Convergence

The goal of this section is to prove Lemmas 3.3 and 3.4. The first shows that the number of odd cycles in the support is nonincreasing, and the second 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). These structural properties are established as follows. First, we give an extension of Edmonds' primal-dual algorithm to half-integral matchings. Next, we argue that applying this algorithm to the current primal/dual solution leads to an optimal solution of the next LP. The analysis of the algorithm shows that the number of odd cycles is nonincreasing. Finally, we prove the extremal dual solution of the next LP must be the same as the one found by this combinatorial algorithm, and therefore the LP solution we use also satisfies the required properties.

As mentioned in the introduction, our extension of Edmonds' algorithm to the half-integral setting and its application are for the purpose of analysis. We describe this algorithm first, derive its properties, then use it in the main proof of convergence of the cutting plane algorithm.

6.1 The half-integral matching algorithm

The next algorithm will be applied in certain contractions of G, but here we present it for a general graph $G = (V, E)$ and cost c. We use the terminology of Edmonds' weighted matching algorithm [8] as described by Schrijver [23, Vol A, Chapter 26].

Let $W \subseteq V$, and let $\mathcal{F} \subset \mathcal{O}$ be a laminar family of odd sets that are disjoint from W. Let $\mathcal{V}^W$ denote the set of one-element subsets of $V \setminus W$. The following primal $P_{\mathcal{F}}^W(G, c)$ and dual $D_{\mathcal{F}}^W(G, c)$ programs describe fractional matchings that leave the set of nodes in W exposed (unmatched) while satisfying the blossom inequalities corresponding to a laminar family $\mathcal{F}$. The primal program is identical to $P_{\mathcal{F}}(G \setminus W, c)$ while optimal solutions to $D_{\mathcal{F}}(G \setminus W, c)$ that are feasible to $D_{\mathcal{F}}^W(G, c)$ are also optimal solutions to $D_{\mathcal{F}}^W(G, c)$.

$$\begin{aligned} \min \sum_{uv \in E} c(uv)x(uv) & \qquad (P_{\mathcal{F}}^W(G, c)) \\ x(\delta(u)) = 1 \quad & \forall u \in V - W \\ x(\delta(u)) = 0 \quad & \forall u \in W \\ x(\delta(S)) \geq 1 \quad & \forall S \in \mathcal{F} \\ x \geq 0 & \end{aligned}$$

$$\begin{aligned} \max \sum_{S \in \mathcal{V}^W \cup \mathcal{F}} \Pi(S) & \qquad (D_{\mathcal{F}}^W(G, c)) \\ \sum_{S \in \mathcal{V}^W \cup \mathcal{F} : uv \in \delta(S)} \Pi(S) \leq c(uv) \quad & \forall uv \in E \\ \Pi(S) \geq 0 \quad & \forall S \in \mathcal{F} \end{aligned}$$

The algorithm is iterative. In each iteration, it maintains a set $T \subseteq W$, a subset $\mathcal{L} \subseteq \mathcal{F}$ of cuts, a proper-half-integral optimal solution z to $P_{\mathcal{L}}^T(G, c)$, and an $\mathcal{L}$-critical dual optimal solution Λ to

$D^T_{\mathcal{L}}(G,c)$ such that $\Lambda(S) > 0$ for every $S \in \mathcal{L}$. In the beginning $T = W$, $\mathcal{L} = \mathcal{F}$ and the algorithm terminates when $T = \emptyset$.

We work 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 $\mathcal{L}$ 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^* the image of z. Since E^* consists only of tight edges, $\Lambda(u) + \Lambda(v) = c^*(uv)$ for every edge $uv \in E^*$. Since $\mathcal{F}$ is disjoint from W, the nodes in $\mathcal{L}$ will always have degree 1 in z^*.

In the course of the algorithm, we may decrease $\Lambda(S)$ to 0 for a maximal set S of $\mathcal{L}$. In this case, we remove S from $\mathcal{L}$ and modify G^*, c^* and z^* accordingly. This operation will be referred as 'unshrinking' S. New sets will never be added to $\mathcal{L}$.

The algorithm 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, 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 $(1/2)(M_{u_1} + M_{u_2})$ of the Λ-critical-matchings M_{u_1} and M_{u_2} inside S.

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 $z^*(v_iv_{i+1})$ to $1 - z^*(v_iv_{i+1})$ 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. The following claim is straightforward.

Claim 6.1 ([23, Thm 24.3])**.** *Let $P = v_0v_1 \dots v_{2k+1}$ be an alternating 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.* □

The algorithm is described in Figure 7. The scenarios in *Case I* are illustrated in Figure 4. In Case II, we observe that $T \in \mathcal{B}^+$ and further, $\mathcal{B}^+ \cap \mathcal{B}^- = \emptyset$ (otherwise, there exists a T–T alternating walk and hence we should be in case I). The correctness of the output follows immediately due to complementary slackness. We show the termination of the algorithm along very similar lines as the proof of termination of Edmonds' algorithm.

Half-integral Matching

Input. A subset $W \subseteq V$, a critical family $\mathcal{F} \subset \mathcal{O}$ with all sets in $\mathcal{F}$ disjoint from W, a proper-half-integral optimal solution w to $P^W_{\mathcal{F}}(G,c)$, and an $\mathcal{F}$-critical dual optimal solution Γ to $D^W_{\mathcal{F}}(G,c)$.
Output. A proper-half-integral optimal solution z to $P_{\mathcal{L}}(G,c)$ and an $\mathcal{L}$-critical dual optimal solution Λ to $P_{\mathcal{L}}(G,c)$ for some $\mathcal{L} \subseteq \mathcal{F}$.

1. Initialize $z = w$, $\mathcal{L} = \mathcal{F}$, $\Lambda = \Gamma$, and $T = W$. Let $G^* = (\mathcal{V}^*, E^*)$, where $E^* \subseteq E$ are edges that are tight w.r.t. Λ, and all maximal sets of $\mathcal{L}$ w.r.t. Λ are contracted; c^* and z^* are defined by the contraction. Let $R \supseteq T$ be the set of exposed 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 $P = v_0 \dots v_{2k+1}$ 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 $v_{2k+1} \in R - T$, 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 6.1, 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

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

   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.

   (a) If some new edge becomes tight, then E^* is extended. $\mathcal{B}^-$ may be removed from E^*.

   (b) If $\Lambda(S) = 0$ for some $S \in \mathcal{L} \cap \mathcal{B}^-$ after the modification, then unshrink the node S. Set $\mathcal{L} := \mathcal{L} \setminus S$.

Figure 7: The Half-integral Matching Algorithm

Before we proceed to analyze the above algorithm, we briefly digress to note that it can be extended 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. We next describe this combinatorial half-integral algorithm and then return to our main analysis.

In the combinatorial half-integral algorithm, we make one modication to the Half-integral Matching subroutine. Namely, in Case II(b) we keep track of the dual values on vertices in W and if one of them decreases to zero, then the corresponding vertex is uncontracted and the vertex is removed from W.

The overall algorithm starts with the empty solution $w = 0, \Gamma = 0, W = V, \mathcal{F} = \emptyset$ and applies the following iteration repeatedly.

1. $[z, \Lambda, \mathcal{L}]$ = Half-integral Matching$(W, \mathcal{F}, w, \Gamma)$.
2. Contract cycles in z w.r.t. Λ; let the new W be the set of contracted vertices.
3. For each vertex of the original W that is incident to a single edge in z, uncontract the vertex, extend z and Λ using the uncontracted edges and earlier dual values, and recurse this way till there are no contracted vertices incident to matching edges in z. Set the new primal and dual solutions to be the ones at the end of this procedure. Include each uncontracted set in the new $\mathcal{F}$.

We can prove the following about this algorithm.

Theorem 6.2. *Combinatorial Half-Integral is a strongly polynomial algorithm for minimim-cost perfect matching.*

The proof builds on that of the main subroutine and is deferred to the end of this section.

Returning to the analysis of the Half-integral Matching algorithm, let $\beta(z)$ denote the number of exposed nodes plus the number of cycles in supp(z). We first note that $\beta(z) = \beta(z^*)$. This can be derived from Lemma 4.1(iii) (We apply this Lemma in $G \setminus T$, observing that $P^T_{\mathcal{L}}(G, c)$ is identical to $P_{\mathcal{L}}(G \setminus T, c)$). Our next lemma (Lemma 6.3) shows that $\beta(z)$ is non-increasing. If $\beta(z)$ is unchanged during a certain number of iterations of the algorithm, we say that these iterations form a *non-decreasing phase*. We say that the algorithm itself is non-decreasing, if $\beta(z)$ does not decrease anytime. In the next section, we investigate properties of non-decreasing phases. These results will also show that every non-decreasing phase may contain at most $|V| + |\mathcal{F}|$ iterations and therefore the algorithm terminates in strongly polynomial time.

Lemma 6.3. *Let z be an arbitrary solution during the algorithm, and let α be the number of odd cycles in supp(w) that are absent in supp(z). Then $|W| + odd(w) \geq \beta(z) + 2\alpha$. At termination, $|W| + odd(w) \geq odd(z) + 2\alpha$.*

Proof. Initially, $\beta(z) = |W| + \text{odd}(w)$. Let us check the possible changes in $\beta(z)$ during an iteration of the algorithm. In Case I(a), the number 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 $\beta(z)$ remains unchanged. In Case II, z is not modified.

The only way to remove a cycle from supp(z) is by performing the operation in Case I(b). This must be executed α times, therefore $\beta(z) \leq \beta(w) - 2\alpha$. Further, there are no exposed nodes at the end of . Thus, on termination $\beta(z) = \text{odd}(z)$, and the claim follows. ◻

6.1.1 The non-decreasing scenario

Let us now analyze the first non-decreasing phase $\mathcal{P}$ of the algorithm, starting from the input w. These results will also be valid for later non-decreasing phases as well. In this case, we say that the algorithm is non-decreasing. Consider an intermediate iteration with z, Λ being the solutions, $\mathcal{L}$ being the laminar family and T being the exposed nodes. Recall that $R \supseteq T$ 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 6.4. *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 performed. We show that the claimed properties are maintained during an iteration.

In Case I(c), a new odd cycle C is created, and thus 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 by the following: 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 alternating path from u_1 to u_2.

We have to check that vertices in $\mathcal{N}_o - \mathcal{B}^+$ and $\mathcal{N}_i - \mathcal{B}^-$ 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). □

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

Corollary 6.5. *The non-decreasing phase $\mathcal{P}$ may consist of at most $|V| + |\mathcal{F}|$ iterations.*

Proof. Case I may occur at most $|W|$ 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 6.4, the first scenario may occur at most $|V|$ times and the second at most $|\mathcal{F}|$ times. □

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

Lemma 6.6. *Assume the half-integral matching algorithm is non-decreasing. Let Γ be the initial dual and z, Λ be the terminating solution and $\mathcal{L}$ be the terminating laminar family. Let $\mathcal{N}_o$ and $\mathcal{N}_i$ denote the final sets of outer and inner nodes in G^*.*

- *If $\Lambda(S) > \Gamma(S)$ then S is an outer node in $\mathcal{V}^*$.*

- *If $\Lambda(S) < \Gamma(S)$, 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 supp(z).*

Proof. If $\Lambda(S) > \Gamma(S)$, then $S \in \mathcal{B}^+$ in some iteration of the algorithm. By Lemma 6.4, this remains an outer node in all later iterations. The conclusion follows similarly for $\Lambda(S) < \Gamma(S)$. □

Lemma 6.7. *Assume the half-integral matching algorithm is non-decreasing. Let z, Λ be the terminating solution, $\mathcal{L}$ be the terminating laminar family and G^* the corresponding contracted graph, $\mathcal{N}_o$ and $\mathcal{N}_i$ be the sets of outer and inner nodes. Let $\Theta : \mathcal{V}^* \to \mathbb{R}$ 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 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. □

6.2 Proof of convergence

Let us consider two consecutive solutions in Algorithm C-P-Matching. Let x be the unique proper-half-integral optimal solution to $P_{\mathcal{F}}(G, c)$ and Π be an $\mathcal{F}$-positively-critical dual optimal solution to $D_{\mathcal{F}}(G, c)$. We define $\mathcal{H}' = \{S : S \in \mathcal{F}, \Pi(S) > 0\}$ and $\mathcal{H}''$ based on odd cycles in x, and use the critical family $\mathcal{H} = \mathcal{H}' \cup \mathcal{H}''$ for the next iteration. Let y be the unique proper-half-integral optimal solution to $P_{\mathcal{H}}(G, c)$, and let Ψ be an $\mathcal{H}$-positively-critical dual optimal solution to $D_{\mathcal{H}}(G, c)$. We already know that Π is an $\mathcal{H}$-critical dual feasible solution to $D_{\mathcal{H}}(G, c)$ by Lemma 3.2.

Let us now contract all maximal sets $S \in \mathcal{H}$ with $\Psi(S) > 0$ w.r.t. Ψ to obtain the graph $\hat{G} = (\hat{V}, \hat{E})$ with cost $\hat{c}$. Note that by Lemma 5.8, Π and Ψ are identical inside S, hence this is the same as contracting w.r.t. Π. Let $\hat{x}$, $\hat{y}$, $\hat{\Pi}$, and $\hat{\Psi}$ be the images of x, y, Π, and Ψ, respectively.

Let $\bar{\mathcal{H}}'' = \{S : S \in \mathcal{H}'', \Psi(S) > 0\}$, and let $W = \cup \bar{\mathcal{H}}''$ denote the union of the members of $\bar{\mathcal{H}}''$. Let $\hat{W}$ denote the image of W. Then $\hat{W}$ is the set of exposed nodes for $\hat{x}$ in $\hat{G}$, whereas the image of every set in $\mathcal{H}'' \setminus \bar{\mathcal{H}}''$ is an odd cycle in $\hat{x}$. Let $\mathcal{N} = \{T \in \mathcal{H}' : T \cap W = \emptyset\}$, $\mathcal{K} = \{T \in \mathcal{N} : \Psi(T) = 0\}$ and $\hat{\mathcal{N}}$ and $\hat{\mathcal{K}}$ be their respective images. All member of $\mathcal{N} \setminus \mathcal{K}$ are contracted to single nodes in $\hat{G}$; observe that $\hat{\mathcal{K}}$ is precisely the set of all sets in $\hat{\mathcal{N}}$ of size at least 3. sets $\hat{T}$ in $\hat{G}$ disjoint from $\hat{W}$; let $\mathcal{K}$ denote the collection of these sets.

Using the notation above, we will couple the solutions of the half-integral matching algorithm and the solutions of our cutting plane algorithm. For this, we will start the Half-integral Matching algorithm in $\hat{G}$ with $\hat{W}$, from the initial primal and dual solutions $\hat{x}$ and $\hat{\Pi}$. Claim 6.8(ii) justifies the validity of this input choice for .

Claim 6.8. *(i) For every $\hat{\mathcal{L}} \subseteq \hat{\mathcal{K}}$, $\hat{y}$ is the unique optimal solution to $P_{\hat{\mathcal{L}}}(\hat{G}, \hat{c})$ and $\hat{\Psi}$ is an optimal solution to $D_{\hat{\mathcal{L}}}(\hat{G}, \hat{c})$.*

(ii) $\hat{x}$ is a proper-half-integral optimal solution to $P^{\hat{W}}_{\hat{\mathcal{K}}}(\hat{G}, \hat{c})$ and $\hat{\Pi}$ is a $\hat{\mathcal{K}}$-positively-critical dual optimal solution to $D^{\hat{W}}_{\hat{\mathcal{K}}}(\hat{G}, \hat{c})$.

Proof. (i) For $\hat{\mathcal{L}} = \emptyset$, both conclusions follow by Corollary 4.3 – y is a unique optimal solution to $P_{\mathcal{H}}(G, c)$ and Ψ is an $\mathcal{H}$-positively-critical dual optimum.

For an arbitrary $\hat{\mathcal{L}} \subseteq \hat{\mathcal{K}}$, since $y(\delta(S)) \geq 1$ for every $S \in \mathcal{H}'$ in the pre-image of $\hat{\mathcal{L}}$, $\hat{y}$ is a feasible solution to $P_{\hat{\mathcal{L}}}(\hat{G}, \hat{c})$. Since $\hat{y}$ is optimum to the bipartite relaxation, this implies optimality of $\hat{y}$ to $P_{\hat{\mathcal{L}}}(\hat{G}, \hat{c})$. Now, $\hat{\Psi}$ is optimum to $D_{\hat{\mathcal{L}}}(\hat{G}, \hat{c})$ since $\hat{\Psi}$ satisfies complementary slackness with $\hat{y}$. Uniqueness follows since a different optimal solution to $P_{\hat{\mathcal{L}}}(\hat{G}, \hat{c})$ would also be optimal to $P_0(\hat{G}, \hat{c})$, but $\hat{y}$ is known to be the unique optimum to $P_0(\hat{G}, \hat{c})$.

(ii) We will use Lemma 4.1 on an appropriate graph. We first setup the parameters that we will be using to apply this Lemma. First observe that x is the unique optimum to $P_{\mathcal{F}}(G, c)$ as well as to $P_{\mathcal{H}'}(G, c)$ and Π is an optimum to $D_{\mathcal{H}'}(G, c)$.

Let G^W denote the graph $G \setminus W$, c^W denote the cost function restricted to this graph. Let Π^W denote Π restricted to the set $\mathcal{N}$ and x^W denote the solution x restricted on G^W. Since $x(\delta(W, V \setminus W)) = 0$, x^W is the unique optimal solution to $P_{\mathcal{N}}(G^W, c^W)$. By complementary slackness, Π^W is optimum to $D_{\mathcal{N}}(G^W, c^W)$.

We now apply Lemma 4.1(iii) by considering the graph G^W with cost c^W, and the laminar odd family $\mathcal{N}$. As noted above, x^W and Π^W are primal and dual optimal solutions. Let $S \in \mathcal{N} \setminus \mathcal{K}$, that is, $\Psi(S) > 0$. Since $S \in \mathcal{N} \subseteq \mathcal{H}'$, we also have $\Pi(S) = \Pi^W(S) > 0$ and hence S is a Π^W-factor-critical set. Let us contract all such sets S w.r.t. Π^W. Then, by the conclusion of the Lemma, we have that $\hat{x}^W$ is the unique proper-half-integral optimum to $P_{\hat{\mathcal{N}}}(\hat{G}^{\hat{W}}, \hat{c}^{\hat{W}})$. Since $\hat{K}$ is the set of nonsingular elements of $\hat{N}$, it is equivalent to saying that $\hat{x}$ is the unique optimum to $P^{\hat{W}}_{\hat{\mathcal{K}}}(\hat{G}, \hat{c})$. By Lemma 4.1(ii), we also have that $\hat{\Pi}^W$ is optimum to $D_{\hat{\mathcal{K}}}(\hat{G}^{\hat{W}}, \hat{c}^{\hat{W}}) = D_{\hat{\mathcal{K}}}(\hat{G} \setminus \hat{W}, \hat{c})$. Recall that every optimal solution to $D_{\hat{\mathcal{K}}}(\hat{G} \setminus \hat{W}, \hat{c})$ that is feasible to $D^{\hat{W}}_{\hat{\mathcal{K}}}(\hat{G}, \hat{c})$ is also optimal to $D^{\hat{W}}_{\hat{\mathcal{K}}}(\hat{G}, \hat{c})$. Observe that $\hat{\Pi}$ is 0 on every node and subset of $\hat{W}$. Consequently, all nonzero values of $\hat{\Pi}$ and $\hat{\Pi}^W$ coincide. Since $\hat{\Pi}$ is feasible to $D^{\hat{W}}_{\hat{\mathcal{K}}}(\hat{G}, \hat{c})$, we get that $\hat{\Pi}$ is also optimum to $D^{\hat{W}}_{\hat{\mathcal{K}}}(\hat{G}, \hat{c})$. Further, $\hat{\Pi}$ is $\hat{\mathcal{K}}$-positively-critical by definition of $\hat{\mathcal{K}}$. □

Lemma 6.9. *Suppose we start the Half-integral Primal-Dual algorithm in $\hat{G}$, $\hat{c}$, $\hat{\mathcal{K}}$, $\hat{W}$, from the initial primal and dual solutions $\hat{x}$ and $\hat{\Pi}$. Then the output $\hat{z}$ of the algorithm is equal to $\hat{y}$.*

Proof. Since $\hat{z}$ is the output of the Half-integral Matching algorithm, it is an optimal solution to $P_{\hat{\mathcal{L}}}(\hat{G}, \hat{c})$ for some $\hat{\mathcal{L}} \subseteq \hat{\mathcal{K}}$. By Claim 6.8(i), $\hat{y}$ is the unique optimal solution to this program, and consequently, $\hat{z} = \hat{y}$. □

Proof of Lemma 3.3. Using the notation above, let us start the Half-integral Matching algorithm in $\hat{G}$, $\hat{c}$, $\hat{\mathcal{K}}$, $\hat{W}$, from the initial primal and dual solutions $\hat{x}$ and $\hat{\Pi}$. Let $\hat{z}$ be the output of the half-integral matching algorithm.

By Lemma 6.9, $\hat{z} = \hat{y}$. We first observe that $\text{odd}(x) = |W| + \text{odd}(\hat{x})$. This easily follows by Lemma 4.1(iii), applied in $G \setminus W$. Let $\alpha = |\mathcal{H}'' \setminus \bar{\mathcal{H}}''|$. There is an odd cycle in $\text{supp}(x)$ corresponding to each set of $\mathcal{H}'' \setminus \bar{\mathcal{H}}''$. None of these cycles may be contained in $\text{supp}(\hat{z}) = \text{supp}(\hat{y})$ as otherwise the corresponding cut in $\mathcal{H}''$ would be violated by y. Thus Lemma 6.3 implies $\text{odd}(\hat{y}) = \text{odd}(\hat{z}) \leq |W| + \text{odd}(\hat{x}) - 2\alpha$ and Lemma 4.1(iii) implies $\text{odd}(y) = \text{odd}(\hat{y})$. Hence, $\text{odd}(y) \leq \text{odd}(x) - 2\alpha$. □

The next claim formulates a simple consequence of the above proof.

Claim 6.10. *If odd(y) = odd(x), then $\mathcal{H}'' = \bar{\mathcal{H}}''$. Further, the Half-integral Matching algorithm applied in $\hat{G}$, $\hat{c}$, $\hat{\mathcal{K}}$, $\hat{W}$, with starting solution $\hat{x}$, $\hat{\Pi}$ is non-decreasing.* □

This claim already implies Lemma 3.4 for $j = i + 1$. To analyze the scenario $\text{odd}(x) = \text{odd}(y)$, the results in Section 6.1.1 are applicable since the algorithm was non-decreasing. Let us start the half-integral matching algorithm in $\hat{G}$, $\hat{c}$, $\hat{\mathcal{K}}$, $\hat{W}$, from the initial primal and dual solutions $\hat{x}$ and $\hat{\Pi}$. Consider the final dual solution $\hat{\Lambda}$ with corresponding laminar family $\hat{\mathcal{L}}$ and define Λ in G as follows.

If $S \subsetneq T$ for some $T \in \mathcal{H}$, $\Psi(T) > 0$, then set $\Lambda(S) = \Psi(S)$ (this defines the dual solutions for sets and nodes inside T that were contracted to obtain $\hat{G}$). If $\hat{S} \in \hat{\mathcal{L}} \cup \hat{\mathcal{V}}$, then set $\Lambda(S) = \hat{\Lambda}(\hat{S})$ for its pre-image S (this defines the dual solutions for sets and nodes on or outside T that were contracted to obtain $\hat{G}$).

Claim 6.11. *Λ is an optimal solution to $D_{\mathcal{H}}(G, c)$.*

Proof. We show feasibility here. Optimality follows by complementary slackness with y and the definition of shrinking. The nontrivial part to show feasibility is to verify that $\Lambda(S) \geq 0$ if $S \in \mathcal{H}$. This is straightforward if $S \subseteq T$ for some $T \in \mathcal{H}$, $\Psi(T) > 0$, or if $\hat{S} \in \hat{\mathcal{L}}$. Let $S \in \mathcal{H}$ be a maximal set with $\Psi(S) > 0$, and let s denote its image in $\hat{G}$. Then during the Half-integral Matching algorithm, $\hat{\Lambda}(s)$ might have decreased below 0. We show that this cannot be the case.

By Claim 6.8(i), $\hat{\Psi}$ is an optimal solution to the bipartite relaxation $D_0(\hat{G}, \hat{c})$. If $\hat{\Lambda}(s) < 0$, then $\hat{\Lambda}(s) < 0 < \hat{\Pi}(s)$. We started the algorithm with the dual solution $\Gamma = \hat{\Pi}$, therefore by Lemma 6.6, either (1) s is an inner node or (2) $s \in \mathcal{V}^*$ and s is incident to an odd cycle in $\text{supp}(z)$. In both cases, by Lemma 6.7, we get that $\Lambda(S) = \hat{\Lambda}(s) \geq \hat{\Psi}(s) = \Psi(S) \geq 0$. □

Lemma 6.12. *Assume odd(x) = odd(y) for the consecutive solutions x and y. Then $\hat{\Lambda} = \hat{\Psi}$ and hence $\Lambda = \Psi$. solution. Then $\Lambda = \Pi_{i+1}$.*

Proof. By Lemma 5.8, if S is such that $\Psi(S) > 0$, then Ψ is identical to Π inside S, and so by definition, also to Λ. Therefore, it suffices to prove $\hat{\Lambda} = \hat{\Psi}$. By Claim 6.8(i), $\hat{\Psi}$ is an optimal solution to $D_0(\hat{G}, \hat{c})$. The following claim completes the proof by showing that $h(\Lambda, \Pi) \leq h(\Psi, \Pi)$, and equality can hold only if Λ and Ψ are identical. □

Claim 6.13. *For every $S \in \hat{\mathcal{V}} \cup \hat{\mathcal{K}}$, $|\hat{\Lambda}(S) - \hat{\Pi}(S)| \leq |\hat{\Psi}(S) - \hat{\Pi}(S)|$ and equality holds only if $\hat{\Lambda}(S) = \hat{\Psi}(S)$.*

Proof. The claim will follow by showing that for every $S \in \hat{\mathcal{V}} \cup \hat{\mathcal{K}}$, either $\hat{\Pi}(S) \leq \hat{\Lambda}(S) \leq \hat{\Psi}(S)$ or $\hat{\Pi}(S) \geq \hat{\Lambda}(S) \geq \hat{\Psi}(S)$. By Claim 6.10, $\text{odd}(x) = \text{odd}(y)$ implies that the half-integral matching algorithm was non-decreasing. Let $\hat{z}$ be the terminating primal solution, and let $G^* = (\mathcal{V}^*, E^*)$, c^* be the corresponding contraction of maximal sets of $\hat{\mathcal{L}}$ in $\hat{G}$, $\hat{c}$. Note that $\hat{\Psi}(S) = 0$ if $S \in \hat{\mathcal{L}}$. Let $\Theta : \mathcal{V}^* \to \mathbb{R}$ be defined as follows:

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

Claim 6.14. Θ *is an optimal solution to* $D_0(G^*, c^*)$*. Further,* $\Theta(S) \le 0$ *holds for every* $S \in \hat{\mathcal{L}}$*.*

Proof. If $S \in \hat{\mathcal{L}}$, then $\hat{\Lambda}(S) > 0$ and hence $\hat{y}(\delta(S)) = 1$. $\hat{\mathcal{L}}$ is a proper odd family in $\hat{G}$ with $\hat{\Lambda}$ being an $\hat{\mathcal{L}}$-critical dual (by definition of the algorithm). By Claim 6.8(i), $\hat{y}$ and $\hat{\Psi}$ are optimal solutions to $P_{\hat{\mathcal{L}}}(\hat{G}, \hat{c})$ and to $D_{\hat{\mathcal{L}}}(\hat{G}, \hat{c})$, respectively. Hence, by Lemma 5.1, $\hat{\Psi}$ is consistent with $\hat{\Lambda}$ inside every set $S \in \hat{\mathcal{L}}$ such that $\hat{y}(\delta(S)) = 1$. Further, $\Delta_{\hat{\Lambda},\hat{\Psi}}(S) \ge 0$ for every $S \in \hat{\mathcal{L}}$. Thus, by Lemma 5.2, Θ is an optimum dual solution to $D_0(G^*, c^*)$. The second part follows since $\hat{\Psi}(S) = 0$ and $\Delta_{\hat{\Lambda},\hat{\Psi}}(S) \ge 0$ if $S \in \mathcal{L}$. □

Now consider $S \subseteq \mathcal{V}$. If $\hat{\Lambda}(S) > \hat{\Pi}(S)$, then by Lemma 6.6, we have that $S \in \mathcal{V}^*$ and $S \in \mathcal{N}_o$. Consequently, by Lemma 6.7, $\Theta(S) \ge \hat{\Lambda}(S)$. If $S \in \hat{\mathcal{L}}$, then $0 \ge \Theta(S^*) \ge \hat{\Lambda}(S)$, a contradiction to $\hat{\Lambda}(S) > 0$. Thus, $S = \{s\}$ for some $s \in \hat{\mathcal{V}}$. Hence, $\hat{\Psi}(s) = \Theta(s) \ge \hat{\Lambda}(s) > \hat{\Pi}(s)$.

If $\hat{\Lambda}(S) < \hat{\Pi}(S)$, then by Lemma 6.6, we have that either *(1)* $S \in \hat{\mathcal{K}} \setminus \hat{\mathcal{L}}$, that is, $\hat{\Lambda}(S) = 0$ 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 supp(z). If *(1)*, then $\hat{\Psi}(S) = 0 = \hat{\Lambda}(S) < \Pi(S)$. In both cases *(2)* and *(3)*, Lemma 6.7 gives $\Theta(S) \le \hat{\Lambda}(S)$. If $S \in \hat{\mathcal{L}}$, then $\hat{\Psi}(S) = 0 \le \hat{\Lambda}(S) < \hat{\Pi}(S)$. If $S = \{s\}$ for some $s \in \hat{V}$, then $\hat{\Psi}(s) = \Theta(s) \le \hat{\Lambda}(s) \le \hat{\Pi}(s)$. □

Proof of Lemma 3.4. Let x_i be the solution in the i'th iteration (above, we used $x = x_i$ and $y = x_{i+1}$). Assume the number of odd cycles does not decrease between iterations i and j. By Claim 6.10, if we run the half-integral matching algorithm between x_k and x_{k+1}, for $i \le k < j$, it is always non-decreasing.

We first run on the contracted graph $\hat{G} = \hat{G}_i$ starting from primal solution $\hat{x} = \hat{x}_i$ and dual solution $\hat{\Pi} = \hat{\Pi}_i$. Lemmas 6.9 and 6.12 show that it terminates with the primal optimal solution $\hat{y} = \hat{x}_{i+1}$ and dual optimal solution $\hat{\Lambda} = \hat{\Psi}$.

For $j = i + 1$, the statement follows by Claim 6.10 since $\bar{\mathcal{H}}'' = \mathcal{H}''$ means that all cuts added in iteration i have positive dual value in iteration $i + 1$. Further, all sets in $\mathcal{H}''$ were contracted to exposed nodes in $\hat{x}_i$. By Lemma 6.6, these will be outer nodes on termination of the half-integral matching algorithm as well. Let G^* be the contracted graph upon termination of the Half-Integral Primal-Dual algorithm.

Let $\mathcal{J} = \mathcal{J}' \cup \mathcal{J}''$ be the set of cuts imposed in the $(i+2)$'th round, with $\mathcal{J} = \{S \in \mathcal{H} : \Psi(Z) > 0\}$, and let $\mathcal{J}''$ be defined according to odd cycles in x_{i+1}. Let Φ be the extremal dual optimal solution to $D_{\mathcal{J}}(G, c)$.

Let us run the half-integral matching algorithm from x_{i+1} to x_{i+2}. We start the algorithm with the contracted graph $\hat{G}_{i+1}$, which results by contracting all sets with $\Phi(S) > 0$, $S \in \mathcal{J}$. Let $\hat{G}^*_{i+1}$ be the initial contraction of $\hat{G}_{i+1}$ used by the algorithm.

The key observation is that while the underlying graphs $\hat{G}_i$ and $\hat{G}_{i+1}$ are different, $\hat{G}^*_{i+1}$ can be obtained from G^* by contracting those odd cycles corresponding to the sets of $\mathcal{J}''$. Every other node that was inner or outer node in G^* will also be inner or outer node in $\hat{G}^*_{i+1}$, including the members of $\mathcal{H}''$. By Lemma 6.6, the members of $\mathcal{H}''$ will be outer nodes at termination, along with the new outer nodes $\mathcal{J}''$.

Iterating this argument one can show that every set that was imposed based on an odd cycle between iterations i and k will be outer nodes at the termination of the Half-integral Matching algorithm from x_k to x_{k+1}. □

We conclude this section with a proof of convergence of the combinatorial half-integral algorithm (Theorem 6.2).

Proof of Theorem 6.2. We need to verify two things. First that the inputs to the main subroutine in each iteration are valid, and second that the algorithm makes progress towards an integral solution.

To see the first, we simply use the contraction properties, namely any set that is uncontracted has a half-integral cycle, which can then be adjusted to maintain a half-integral solution after uncontraction. Moreover, the dual values for vertices inside will be the values at the point of contraction. An argument similar to the proof of Lemma 6.9 establishes the validity of this dual solution.

To show progress, we use odd(x) as in the analysis of the cutting plane algorithm. In the cutting plane analysis, in a non-decreasing phase, the cuts that are added in each iteration remain in all iterations until the end of the non-decreasing phase. Instead, what we argue here is that the number of contracted vertices plus the size of the laminar family at the end of a phase is at least the total number of half-integral cycles encountered during the course of the non-decreasing phase. To see this, we first note that each half-integral cycle at the end of a phase is contracted. Next, when a vertex is uncontracted, the corresponding set is added to the laminar family. Now we need to verify that when a set is added to the laminar family it remains in the family till the end of the phase. This is similar to the proof of Lemma 3.4. Therefore the length of a non-decreasing phase is bounded, and polynomiality of the algorithm follows. □

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 2.3, showing that a perturbation of the original integer cost function satisfies our uniqueness assumption (*). We will need the following simple claim.

Claim 7.1. *For a graph $G = (V, E)$, let $a, b : E \to \mathbb{R}_+$ 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. □

Proof of Lemma 2.3. Let $\tilde{c}$ denote the perturbation of the integer cost $c : E \to \mathbb{Z}$. Consider a graph $G = (V, E)$, perturbed cost $\tilde{c}$ and critical family $\mathcal{F}$ where (*) 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.

If $\mathcal{F} \neq \emptyset$, after modifying x, y to non-identical optimal solutions a, b respectively (if necessary), we will similarly 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 $P_{\mathcal{F}\setminus\{S\}}(G, \tilde{c})$, contradicting the minimal choice of $\mathcal{F}$. Let $T \in \mathcal{F} \cup \{V\}$ be a smallest set with the property that $x(\delta(u, V - T)) = y(\delta(u, V - T))$ for every $u \in T$, but x and y are not identical inside T. 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 $\alpha_u := x(\delta(u, V - S)), \alpha'_u = y(\delta(u, V - S))$ and $w := \sum_{u\in S} \alpha_u M_u, w' := \sum_{u\in S} \alpha'_u M_u$. If x and y are not identical inside S and $\alpha_u = \alpha'_u$ for every $u \in S$, then S contradicts the minimality of T. Hence, for each maximal set $S \in \mathcal{S}$ we have two cases – either *(1)* x and y are identical inside S, and thus for every $u \in S$, $\alpha_u = \alpha'_u$ or *(2)* $\alpha_u \neq \alpha'_u$ for some $u \in S$.

Fix $S \in \mathcal{S}$. The following modified solutions

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

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

are also optimal solutions since we used only tight edges w.r.t. Π for these modifications.

Claim 7.2. *Consider two edges $u_1v_1 \in \delta(S) \cap supp(a)$, $u_2v_2 \in \delta(S) \cap supp(b)$, $u_1, u_2 \in S$. 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). □

We perform the above modifications inside every $S \in \mathcal{S}$, and denote by a and b the result of all these modifications. Let us now contract all sets in $\mathcal{S}$ w.r.t. Π; let $G', \tilde{c}', \mathcal{F}'$ denote the resulting graph, costs and laminar family respectively. By Lemma 4.1(ii), the images a' and b' are both optimal solutions in $P'_{\mathcal{F}'}(G', \tilde{c}')$.

We claim that a' and b' are still not identical inside T'. Indeed, assume a' and b' are identical inside T'. Then we must have had case *(1)* for every $S \in \mathcal{S}$. Note that we only modified x and y on the edge set $\cup_{S\in\mathcal{S}} E[S]$, and hence the contracted images must be identical: $x' = a'$ and $y' = b'$, and consequently, $x' = y'$. As x and y are not identical inside T, they must differ inside at least one $S \in \mathcal{S}$. This contradicts the fact that case *(1)* applied for every $S \in \mathcal{S}$.

By the assumption on T, we also have $a'(\delta(u', V' - T')) = b'(\delta(u', V' - T'))$ for every $u' \in T'$. Hence Claim 7.1 is applicable, giving an alternating closed walk C' 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 in the above claim. 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. □

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 iterations, the adversary replaces the current dual solution with a different one, still satisfying complementary slackness with the (unchanged) primal solution.

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 a similar approach for finding an optimal solution for b-matchings. Another direction could be to try this approach for optimizing over the subtour elimination polytope.

Acknowledgment. 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*, volume 5, pages 3–51. Elsevier, 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] V. Chvátal. Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete Mathematics*, 4(4):305–337, Apr 1973.

[5] W. Cook, R. Kannan, and A. Schrijver. Chvátal closures for mixed integer programming problems. *Mathematical Programming*, 47(1):155–174, May 1990.

[6] G. Cornuéjols and Y. Li. Elementary closures for integer programs. *Operations Research Letters*, 28(1):1–8, Feb 2001.

[7] G. Dantzig, R. Fulkerson, and S. Johnson. Solution of a large-scale traveling-salesman problem. *Operations Research*, 2(4):393–410, Nov 1954.

[8] 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.

[9] J. Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17(3):449–467, 1965.

[10] M. Fischetti and A. Lodi. Optimizing over the first Chvátal closure. *Mathematical Programming*, 110(1):3–20, 2007.

[11] R. Gomory. Outline of an algorithm for integer solutions to linear programs. *Bulletin of the American Mathematical Society*, 64(5):275–278, 1958.

[12] R. Gomory. Solving linear programs in integers. In *Proceedings of Symposia in Applied Mathematics*, volume 10, pages 211–215, July 1960.

[13] R. Gomory. An algorithm for integer solutions to linear programs. *Recent advances in mathematical programming*, page 269, 1963.

[14] M. Grötschel and O. Holland. Solving matching problems with linear programming. *Mathematical Programming*, 33(3):243–259, 1985.

[15] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization.* Springer, 1988.

[16] R. Kannan. Minkowski's convex body theorem and integer programming. *Mathematics of Operations Research*, 12(3):415–440, 1987.

[17] L. Lovász and M. Plummer. *Matching theory.* North Holland, 1986.

[18] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal of Optimization*, 1:166–190, 1991.

[19] G. Nemhauser and L. Wolsey. *Integer and Combinatorial Optimization.* Wiley-Interscience, 1999.

[20] 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.

[21] M. Padberg and M. Rao. Odd minimum cut-sets and b-matchings. *Mathematics of Operations Research*, 7(1):67–80, 1982.

[22] A. Schrijver. *Theory of Linear and Integer Programming.* John Wiley & sons, 1998.

[23] A. Schrijver. *Combinatorial optimization.* Springer, 2003.

[24] 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.

[25] M. Trick. *Networks with additional structured constraints.* PhD thesis, Georgia Institute of Technology, 1987.