Maximum Matchings and Popularity

Telikepalli Kavitha
Tata Institute of Fundamental Research, Mumbai, India

Abstract

Let G be a bipartite graph where every node has a strict ranking of its neighbors. For any node, its preferences over neighbors extend naturally to preferences over matchings. A maximum matching M in G is a *popular max-matching* if for any maximum matching N in G, the number of nodes that prefer M is at least the number that prefer N. Popular max-matchings always exist in G and they are relevant in applications where the size of the matching is of higher priority than node preferences. Here we assume there is also a cost function on the edge set. So what we seek is a min-cost popular max-matching. Our main result is that such a matching can be computed in polynomial time.

We show a compact extended formulation for the popular max-matching polytope and the algorithmic result follows from this. In contrast, it is known that the popular matching polytope has near-exponential extension complexity and finding a min-cost popular matching is NP-hard.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Bipartite graphs, Popular matchings, Stable matchings, Dual certificates

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.85

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: https://arxiv.org/abs/2011.03434

Funding *Telikepalli Kavitha*: Supported by the Department of Atomic Energy, Government of India, under project no. RTI4001.

Acknowledgements Thanks to Yuri Faenza and the reviewers for helpful comments and suggestions on the presentation.

1 Introduction

We consider a matching problem in a bipartite graph $G = (A \cup B, E)$ on n nodes and m edges where every node has a strict ranking of its neighbors. The bipartite graph G need not be complete. This is a very well-studied model in two-sided matching markets. This model has been used to match students to schools and colleges [1, 4] and doctors to residencies in hospitals [6, 28].

The goal is to find an optimal matching in G. The classical notion of optimality in such an instance is *stability*. A matching M is stable if there is no edge that *blocks* M; an edge (a, b) blocks M if a and b prefer each other to their respective assignments in M. Stable matchings always exist in G and one such matching can be computed in linear time by the Gale-Shapley algorithm [17].

In several applications, along with node preferences, the definition of optimality may involve attributes such as size, e.g., consider the problem of assigning doctors to hospitals during a pandemic. For instance, during the Covid-19 pandemic in Mumbai, public hospitals were overwhelmed with a rising number of patients and were severely short-staffed; so doctors associated with private clinics were asked to also work in public hospitals[1] [3, 22]. We want the maximum number of doctors to get assigned to hospitals, i.e., we do not wish to

[1] In the doctors-hospitals setting, one typically seeks a many-to-one matching since a hospital may need several doctors. Here we focus on one-to-one matchings.

compromise at all on the size of our matching. Thus the size of the matching is of higher priority here than node preferences. We refer to [5] for more such applications: these include school placement [1] and the assignment of sailors to billets [29, 34]. So the set of admissible solutions in all these applications is the set of maximum matchings and among all maximum matchings, we wish to find a *best* matching as per node preferences.

Note that a stable matching need not have maximum size. It is known that all stable matchings match the same set of nodes [18]. When stable matchings are not admissible, a natural alternative would be to seek a maximum matching with the least number of blocking edges. However this is an NP-hard problem [5] and as shown there, this is NP-hard to approximate within $n^{1-\varepsilon}$, for any $\varepsilon > 0$.

Observe that stability empowers every edge with the "veto power" to block a matching. Thus stability is a very strong notion and the notion of *popularity* is a meaningful relaxation of stability that captures collective welfare. We say node u prefers matching M to matching N if u prefers its assignment in M to its assignment in N; being unmatched is the worst choice for any node. We can compare any pair of matchings M and N by holding an election between them where nodes are voters.

Let $\phi(M, N)$ be the number of nodes that prefer M to N and similarly, let $\phi(N, M)$ be the number of nodes that prefer N to M. We say N is more popular than M if $\phi(N, M) > \phi(M, N)$.

▶ **Definition 1.** *A matching M is popular if $\Delta(M, N) \geq 0$ for all matchings N in G, where $\Delta(M, N) = \phi(M, N) - \phi(N, M)$.*

Thus a matching M is popular if there is no matching that is more popular than M. Every stable matching is popular[2] [19]. One of the main merits of popularity is that it allows larger matchings compared to stable matchings. Since a stable matching is a maximal matching, its size is at least $|M_{\max}|/2$ (where $M_{\max}$ is a maximum matching in G) and there are easy examples where this bound is tight. It is known that there is always a popular matching of size at least $2|M_{\max}|/3$ and there are simple instances with no larger popular matching [23].

Since a popular matching need not be maximum, a natural alternative is to ask for a maximum matching M that is popular *within* the set of maximum matchings, i.e., no maximum matching is more popular than M. Since it is only maximum matchings that are admissible and we are not willing to replace a maximum matching with a smaller one, elections that involve *non-maximum* or inadmissible matchings are not relevant here. Hence a natural candidate for a best maximum matching is a *popular max-matching* defined below.

▶ **Definition 2.** *Call a maximum matching M in $G = (A \cup B, E)$ a popular max-matching if $\Delta(M, N) \geq 0$ for all maximum matchings N in G.*

Thus what we seek is a *weak Condorcet winner* [7, 27] in the voting instance where maximum matchings are candidates and nodes are voters. The relation "more popular than" is not transitive and weak Condorcet winners need not exist in every voting instance. The question of whether every instance admits a popular max-matching was considered in [23] where it was shown that popular max-matchings always exist in $G = (A \cup B, E)$ and one such matching can be computed in $O(mn)$ time.

---

[2] In an election between a stable matching S and any matching M, if node u prefers M to S then the node $M(u)$ has to prefer S to M, otherwise $(u, M(u))$ blocks S, which is forbidden. Hence $\phi(M, S) \leq \phi(S, M)$.

In several applications there is a cost function $c : E \to \mathbb{R}$ and for any matching M, its cost $c(M) = \sum_{e \in M} c(e)$. So among all popular max-matchings in G, what we seek is a *min-cost* popular max-matching. There are no previous algorithmic/hardness results known for this problem. We know that it is NP-hard to find a min-cost popular matching [13].

Finding a min-cost popular max-matching is a natural and interesting problem in discrete optimization. Solving this problem efficiently implies efficient algorithms for a whole host of popular max-matching problems such as finding one with forced/forbidden edges or one with max-utility or one with min-regret. In general, a cost function allows us to "access" the entire set of popular max-matchings; note that G may have more than 2^n such matchings [32].

Let $\mathcal{M}_G$ denote the popular max-matching polytope of G, i.e., this polytope is the convex hull of the edge incidence vectors of popular max-matchings in G. A compact description of $\mathcal{M}_G$ (or some extension[3] of it) implies a polynomial time algorithm to compute a min-cost popular max-matching. Our main result is that the polytope $\mathcal{M}_G$ has a compact extended formulation. So unlike the min-cost popular matching problem, interestingly and quite surprisingly, the min-cost popular max-matching problem is tractable.

▶ **Theorem 3.** *Given $G = (A \cup B, E)$ where nodes have strict preferences and $c : E \to \mathbb{R}$, a min-cost popular max-matching can be computed in polynomial time.*

So Theorem 3 shows that a natural variant of the min-cost popular matching problem (which is NP-hard) admits a polynomial time algorithm. We also consider *Pareto-optimality* – this is a far more relaxed notion than popularity that any reasonable matching in this domain should satisfy. If M is a matching that is not Pareto-optimal then there is a matching N such that no node is worse-off in N than in M and at least one node is better-off.

The *unpopularity factor* of M is defined as follows [26]: $u(M) = \max_{N \neq M} \frac{\phi(N,M)}{\phi(M,N)}$. Observe that a popular matching M satisfies $u(M) \leq 1$. A matching M is Pareto-optimal if $u(M) < \infty$. A maximum matching M that satisfies $u(M) < \infty$ is a Pareto-optimal max-matching. We show the following hardness result here.

▶ **Theorem 4.** *Given $G = (A \cup B, E)$ with strict preferences and edge costs in $\{0, 1\}$, it is NP-hard to compute a min-cost Pareto-optimal matching/max-matching in G. Moreover, it is NP-hard to approximate this within any multiplicative factor.*

1.1 Background

The notion of popularity was introduced by Gärdenfors [19] in 1975 where he observed that every stable matching is popular. Many algorithmic questions in popular matchings have been investigated in the last 10-15 years and we refer to [8] for a survey. In the domain of popular matchings with two-sided preferences (every node has a preference order ranking its neighbors), other than a handful of positive results [10, 20, 23, 25], most optimization problems have turned out to be NP-hard [13].

Computing a min-cost *quasi-popular* matching M, i.e., $u(M) \leq 2$, is also NP-hard [12]. Compact extended formulations for the *dominant* matching[4] polytope [10, 12] and the popular fractional matching polytope [24] are known but the popular matching polytope has near-exponential extension complexity [12].

[3] A polytope Q that linearly projects to a polytope P is an *extension* of P and a linear description of Q is an *extended formulation* for P. The minimum size of an extension of P is the *extension complexity* of P.
[4] These are popular matchings that are more popular than all larger matchings.

Though an $O(mn)$ time algorithm to find a popular max-matching in G is known [23], there are no previous results on its *optimization* variant, i.e., to find a min-cost or max-utility popular max-matching. It is common in this domain to have an efficient algorithm to find a max-size matching in some class – say, popular matchings [20, 23] or Pareto-optimal matchings for one-sided preferences (only nodes in A have preferences) studied in [2], however finding a min-cost matching in these classes is NP-hard [2, 13]; Theorem 4 also shows such a hardness result. Thus an efficient algorithm to find some popular max-matching was no guarantee on the tractability of the min-cost popular max-matching problem.

There are several polynomial time algorithms to compute a min-cost stable matching and some variants of this problem [11, 14, 15, 16, 21, 30, 31, 33]. Moreover, the stable matching polytope of G has a linear-size description in $\mathbb{R}^m$ [30]. Thus in contrast to stable matchings, the landscape of popular matchings has only a few positive results.

1.2 Our Techniques

An algorithm called the "$|A|$-level Gale-Shapley algorithm" was given in [23] to find a popular max-matching in $G = (A \cup B, E)$. As we show in Section 2, this algorithm is equivalent to running the Gale-Shapley algorithm in an auxiliary instance G^* with $|A|$ copies of each node in A. The proof in [23] can be easily adapted to show that there is a map from the set of stable matchings in G^* to the set of popular max-matchings in G (see Theorem 6). Our novel contribution here is to show that this map is *surjective*, i.e., every popular max-matching in G is the image of a stable matching in G^*.

Loosely speaking, the instance G^* is made up of $|A|$ copies of G and is inspired by an instance from [10] that is made up of two copies of G. So our instance G^* is much larger than the instance used in [10]. In order to realize a popular max-matching in G as the image of a stable matching in G^*, rather than be guided by blocking edges (as done in [10]), we need a "global handle" over the given popular max-matching, i.e., we seek a function from $A \cup B$ to $\{0, \ldots, |A| - 1\}$ that guides us in how to "place" this matching in the instance G^*.

LP-duality gives us such a handle in terms of dual certificates. Dual certificates for popular matchings are well-understood: these are in $\{0, \pm 1\}^n$ [24]. To show a compact extended formulation for $\mathcal{M}_G$ (unlike the popular matching polytope), finding the right dual certificates is crucial. Our proof consists of two parts: when G admits a perfect matching, it is the easy case. We use total unimodularity, complementary slackness, and the fact that G has a perfect matching to show that the given popular max-matching M has a dual certificate $\vec{\alpha}$ where $\alpha_a \in \{0, -2, -4, \ldots\}$ for $a \in A$ and $\alpha_b \in \{0, 2, 4, \ldots\}$ for $b \in B$. Such an $\vec{\alpha}$ can be neatly used to realize M as the image of a stable matching in G^* (see Theorem 8).

The general case. When G does not admit a perfect matching, things are more complicated. The primal LP will not be as simple as (LP1) whose constraints describe the perfect matching polytope. So we reduce the general case to the case when M is a perfect matching, i.e., we use the dual solution $\vec{\alpha} \in \{0, \pm 2, \pm 4, \ldots\}^{2n_0}$ that certifies M's optimality in the subgraph $G' = G \setminus \{\text{nodes not matched in } M\}$ on $2n_0 = 2|M|$ nodes.

We need to update $\vec{\alpha}$ so that it certifies M's optimality in the entire graph G. Our main technical novelty is in how we update $\vec{\alpha}$ using certain rules (see Theorem 9). Let $A' \cup B'$ be the node set of G' and let U be the set of nodes not matched in M.

We use the fact that M is a maximum matching to prove that our update procedure terminates with a dual certificate $\vec{\alpha}$ for M in G' where $\alpha_a \in \{0, -2, -4, \ldots, -2(n_0 - 1)\}$ for $a \in A'$ and $\alpha_b \in \{0, 2, 4, \ldots, 2(n_0 - 1)\}$ for $b \in B'$ such that the neighbors of nodes in U take the highest possible α-values, i.e., (i) $\alpha_a = 0$ for $a \in A' \cap \mathsf{Nbr}(U)$ and (ii) $\alpha_b = 2(n_0 - 1)$ for $b \in B' \cap \mathsf{Nbr}(U)$. Roughly speaking, such an $\vec{\alpha}$ will *take care* of the edges of G missing in G' and will allow us to realize M as the image of a stable matching in G^* (see Theorem 10).

2 Popular Max-Matchings

In this section we first show a simple characterization of popular max-matchings. Then we show a method to construct matchings that satisfy this characterization. Let M be any matching in $G = (A \cup B, E)$. The following edge weight function wt_M will be useful here. For any $(a, b) \in E$:

$$\text{let } \mathsf{wt}_M(a,b) = \begin{cases} 2 & \text{if } (a,b) \text{ blocks } M; \\ -2 & \text{if } a \text{ and } b \text{ prefer their assignments in } M \text{ to each other;} \\ 0 & \text{otherwise.} \end{cases}$$

So $\mathsf{wt}_M(e) = 0$ for every $e \in M$. For any edge e, $\mathsf{wt}_M(e)$ is the sum of votes (each vote is in $\{0, \pm 1\}$) of the endpoints of e for each other versus their respective assignments in M.

For any cycle/path ρ in G, let $\mathsf{wt}_M(\rho) = \sum_{e \in \rho} \mathsf{wt}_M(e)$. Theorem 5 uses this edge weight function to characterize popular max-matchings. Recall that an alternating path (resp., cycle) with respect to matching M is a path (resp., cycle) whose alternate edges are in M.

▶ **Theorem 5.** *For any maximum matching M in G, M is a popular max-matching if and only if (1) there is no alternating cycle C wrt M such that $\mathsf{wt}_M(C) > 0$ and (2) there is no alternating path p with an unmatched node as an endpoint such that $\mathsf{wt}_M(p) > 0$.*

Proof. Let M be a popular max-matching. We need to show that conditions (1) and (2) given in the theorem statement hold. Suppose not. Then there exists either an alternating path with an unmatched node as an endpoint or an alternating cycle wrt M (call this path/cycle ρ) such that $\mathsf{wt}_M(\rho) > 0$. Since $\mathsf{wt}_M(e) \in \{0, \pm 2\}$, $\mathsf{wt}_M(\rho) \geq 2$.

Consider $N = M \oplus \rho$. This is a maximum matching in G and observe that $\Delta(N, M) \geq \mathsf{wt}_M(\rho) - 1$. We are subtracting 1 here to count for that endpoint of ρ (when ρ is a path) that is matched in M but will become unmatched in N. Since $\mathsf{wt}_M(\rho) \geq 2$, $\Delta(N, M) \geq 1$. So N is more popular than M; this is a contradiction to M's popularity within the set of maximum matchings. Thus conditions (1) and (2) have to hold.

To show the converse, suppose M is a maximum matching that obeys conditions (1) and (2). Consider the symmetric difference $M \oplus N$, where N is any maximum matching in G and let C be any alternating cycle here. We know from (1) that $\mathsf{wt}_M(C) \leq 0$. Let p be any alternating path in $M \oplus N$. Since M and N are maximum matchings, p is an alternating path with exactly one node not matched in M as an endpoint. We know from (2) that $\mathsf{wt}_M(p) \leq 0$. So we have $\Delta(N, M) \leq \sum_{\rho \in M \oplus N} \mathsf{wt}_M(\rho) \leq 0$. Thus no maximum matching is more popular than M. ◀

A new instance. We will now construct a new instance $G^* = (A^* \cup B^*, E^*)$ such that every stable matching in G^* maps to a maximum matching in G that satisfies properties (1) and (2) given in Theorem 5. As mentioned earlier, the structure of the instance G^* is inspired by an instance from [10] whose stable matchings map to dominant matchings in G.

We first describe the node sets A^* and B^*. Let $n_0 = |A|$. For every $a \in A$, the set A^* has n_0 copies of a: call them $a_0, \ldots, a_{n_0-1}$. So $A^* = \cup_{a \in A}\{a_0, \ldots, a_{n_0-1}\}$.

Let $B^* = \cup_{a \in A}\{d_1(a), \ldots, d_{n_0-1}(a)\} \cup \{\tilde{b} : b \in B\}$, where $\tilde{B} = \{\tilde{b} : b \in B\}$ is a copy of the set B. So along with nodes in $\tilde{B}$, the set B^* also contains $n_0 - 1$ nodes $d_1(a), \ldots, d_{n_0-1}(a)$ for each $a \in A$. These will be called *dummy* nodes. The purpose of $d_1(a), \ldots, d_{n_0-1}(a)$ is to ensure that in any stable matching in G^*, at most one node among $a_0, \ldots, a_{n_0-1}$ is matched to a neighbor in $\tilde{B}$.

The edge set. For each $(a, b) \in E$, the edge set E^* contains n_0 edges $(a_i, \tilde{b})$ for $0 \le i \le n_0-1$. For each $a \in A$ and $i \in \{1, \ldots, n_0 - 1\}$, E^* also has $(a_{i-1}, d_i(a))$ and $(a_i, d_i(a))$.

Preference orders. Let a's preference order in G be $b_1 \succ \cdots \succ b_k$. Then a_0's preference order in G^* is $\tilde{b}_1 \succ \cdots \succ \tilde{b}_k \succ d_1(a)$, i.e., it is analogous to a's preference order in G with $d_1(a)$ added as a_0's last choice.

- For $i \in \{1, \ldots, n_0 - 2\}$, the preference order of a_i in G^* is as follows: $d_i(a) \succ \tilde{b}_1 \succ \cdots \succ \tilde{b}_k \succ d_{i+1}(a)$. So a_i's top choice is $d_i(a)$ and last choice is $d_{i+1}(a)$.
- The preference order of a_{n_0-1} is $d_{n_0-1}(a) \succ \tilde{b}_1 \succ \cdots \succ \tilde{b}_k$.

For each $i \in \{1, \ldots, n_0 - 1\}$, the preference order of $d_i(a)$ is $a_{i-1} \succ a_i$. Since each of $a_0, \ldots, a_{n_0-2}$ and $d_1(a), \ldots, d_{n_0-1}(a)$ is a top choice neighbor for some node, every stable matching in G^* has to match all these nodes. So the only node among $a_0, \ldots, a_{n_0-1}, d_1(a), \ldots, d_{n_0-1}(a)$ that can possibly be left unmatched in a stable matching in G^* is a_{n_0-1}.

Consider any $b \in B$ and let its preference order in G be $a \succ \cdots \succ z$. Then the preference order of $\tilde{b}$ in G^* is

$$\underbrace{a_{n_0-1} \succ \cdots \succ z_{n_0-1}}_{\text{all subscript } n_0-1 \text{ neighbors}} \succ \underbrace{a_{n_0-2} \succ \cdots \succ z_{n_0-2}}_{\text{all subscript } n_0-2 \text{ neighbors}} \succ \cdots \succ \underbrace{a_0 \succ \cdots \succ z_0}_{\text{all subscript 0 neighbors}}$$

That is, $\tilde{b}$'s preference order in G^* is all its subscript $n_0 - 1$ neighbors, followed by all its subscript $n_0 - 2$ neighbors, so on, and finally, all its subscript 0 neighbors. For each $i \in \{0, \ldots, n_0 - 1\}$: within all subscript i neighbors, the order of preference for $\tilde{b}$ in G^* is the same as b's order of preference in G.

The set S'. For any stable matching S in G^*, define $S' \subseteq E$ to be the set of edges obtained by deleting edges in S that are incident to dummy nodes and replacing any edge $(a_i, \tilde{b}) \in S$ with the original edge $(a, b) \in E$. Since S matches at most one node among $a_0, \ldots, a_{n_0-1}$ to a neighbor in $\tilde{B}$, the set S' is a matching in G.

The proof of Theorem 6 is based on the proof of correctness of the $|A|$-level Gale-Shapley algorithm (from [23]) in the original instance G.

▶ **Theorem 6.** *Let S be a stable matching in G^*. Then S' is a popular max-matching in G.*

Proof. Partition the set A into $A_0 \cup \cdots \cup A_{n_0-1}$ where for $0 \le i \le n_0 - 2$: $A_i = \{a \in A : (a_i, \tilde{b}) \in S$ for some $\tilde{b} \in \tilde{B}\}$, i.e., a_i is matched in S to a neighbor in $\tilde{B}$. The left-out nodes in A, i.e., those in $A \setminus (A_0 \cup \cdots \cup A_{n_0-2})$, form the set A_{n_0-1} (see Fig. 1).

Similarly, partition B into $B_0 \cup \cdots \cup B_{n_0-1}$ where for $1 \le i \le n_0 - 1$: $B_i = \{b : (a_i, \tilde{b}) \in S$ for some $a \in A_i\}$, i.e., $\tilde{b}$'s partner in S is a subscript i node. Let $B_0 = B \setminus (B_1 \cup \cdots \cup B_{n_0-1})$ be the set of left-out nodes in B.

The following properties hold: (these are proved below)

1. $S' \subseteq \cup_{i=0}^{n_0-1}(A_i \times B_i)$. Moreover, S' restricted to each set $A_i \cup B_i$ is stable.
2. For any i and edge (a, b) where $a \in A_{i+1}, b \in B_i$: we have $\mathsf{wt}_{S'}(a, b) = -2$.

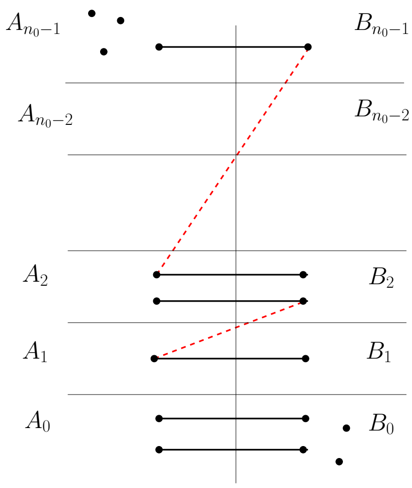

Figure 1 $A = A_0 \cup \cdots \cup A_{n_0-1}$ and $B = B_0 \cup \cdots \cup B_{n_0-1}$. The matching $S' \subseteq \cup_{i=0}^{n_0-1}(A_i \times B_i)$. All nodes unmatched in S' are in $A_{n_0-1} \cup B_0$. The dashed edges are blocking edges to S'.

3. G has no edge in $A_i \times B_j$ where $i \geq j+2$.
4. Any blocking edge to S' has to be in $A_i \times B_j$ where $i \leq j-1$.
5. All nodes that are unmatched in S' are in $A_{n_0-1} \cup B_0$.
6. S' is a maximum matching in G.

Properties 1–4 imply that for any alternating cycle C wrt S', $\mathsf{wt}_{S'}(C) \leq 0$. Similarly, properties 1–5 imply that for any alternating path p with one unmatched node as an endpoint, $\mathsf{wt}_{S'}(p) \leq 0$. We refer to [23, Theorem 2] for more details. Property 6 states that S' is a maximum matching in G. Hence S' is a popular max-matching in G (by Theorem 5). ◀

Properties 1-6. These six properties are proved below.

1. The inclusion $S' \subseteq \cup_{i=0}^{n_0-1}(A_i \times B_i)$ follows from the definition of the sets $B_0, \ldots, B_{n_0-1}$. Recall that for $1 \leq i \leq n_0 - 1$, B_i is the set of nodes b such that $(a_i, \tilde{b}) \in S$ for some $a \in A_i$. Also, B_0 contains all nodes b such that $(a_0, \tilde{b}) \in S$ for some $a \in A_0$. Thus $S' \subseteq \cup_{i=0}^{n_0-1}(A_i \times B_i)$.
   The stability of S' restricted to each set $A_i \cup B_i$ is by $\tilde{b}$'s preference order in G^*. Recall that within subscript i neighbors, the order of preference for $\tilde{b}$ in G^* is b's order of preference in G. Thus the stability of S in G^* implies the stability of S' restricted to $A_i \cup B_i$ for each i.
2. Let $a \in A_{i+1}$. Then $(a_{i+1}, d_{i+1}(a)) \notin S$. So it has to be the case that $(a_i, d_{i+1}(a)) \in S$. Recall that $d_{i+1}(a)$ is a_i's least preferred neighbor in G^*. So a_i prefers $\tilde{b}$ to its partner in S. Hence it follows from the stability of S in G^* that $\tilde{b}$ prefers its partner in S (this is a subscript i node z_i) to a_i, i.e., b prefers z to a.
   Since $\tilde{b}$ prefers subscript $i+1$ nodes to subscript i nodes, $\tilde{b}$ prefers a_{i+1} to its partner z_i in S. It follows from the stability of S in G^* that a_{i+1} has to prefer its partner $\tilde{w}$ in S to $\tilde{b}$, otherwise $(a_{i+1}, \tilde{b})$ would block S. Hence a prefers w to b. Thus $\mathsf{wt}_{S'}(a,b) = -2$.
3. Suppose $a \in A_i$ where $i \geq j+2$ and $b \in B_j$. So the edge $(a_{j+1}, d_{j+2}(a)) \in S$. Since $d_{j+2}(a)$ is a_{j+1}'s least preferred neighbor in G^*, the stability of S implies that $\tilde{b}$ prefers its partner in S to a_{j+1}. However $b \in B_j$ and so $\tilde{b}$'s partner in S is a subscript j node z_j. This contradicts $\tilde{b}$'s preference order that it prefers any subscript $j+1$ neighbor to a subscript j neighbor. Thus there is no edge (a,b) in G with $a \in A_i$ and $b \in B_j$ where $i \geq j+2$.

4. Property 4 follows from properties 1, 2, and 3 given above. Properties 2 and 3 tell us that there is no blocking edge in $A_i \times B_j$ where $i \geq j+1$. Property 1 tells us that there is no blocking edge in $A_i \times B_i$ for any i. So any blocking edge to S' has to be in $A_i \times B_j$ where $i \leq j-1$.
5. Property 5 follows from the definitions of the sets $A_0, \ldots, A_{n_0-2}$ and $B_1, \ldots, B_{n_0-1}$. For each $a \in A_i$ where $0 \leq i \leq n_0 - 2$: we have $(a_i, \tilde{b}) \in S$ for some $\tilde{b} \in \tilde{B}$ and thus $(a, b) \in S'$. Similarly, for each $b \in B_j$ where $1 \leq j \leq n_0 - 1$: we have $(a_j, \tilde{b}) \in S$ for some $a \in A_j$ and thus $(a, b) \in S'$. Hence all nodes unmatched in S' are in $A_{n_0-1} \cup B_0$.
6. Suppose S' is not a maximum matching in G. Then there is an augmenting path ρ with respect to S'. Let us refer to an edge e that satisfies $\mathsf{wt}_{S'}(e) = -2$ as a *negative* edge. The endpoints of a negative edge prefer their respective partners in S' to each other.
We know from property 5 above that all the nodes in A that are unmatched in S' are in A_{n_0-1} and all the nodes in B that are unmatched in S' are in B_0. We also know that $S' \subseteq \cup_i (A_i \times B_i)$ (by property 1 above). Moreover, all the edges e in $A_{j+1} \times B_j$ are negative edges (by property 2) and there is no edge in $A_i \times B_j$ where $i \geq j+2$ (by property 3).
Thus the path ρ starts in A_{n_0-1} at an unmatched node a and since there cannot be any negative edge incident to an unmatched node, all of a's neighbors have to be in B_{n_0-1}: this is because every edge e in $A_{n_0-1} \times B_{n_0-2}$ is a negative edge. The matched partners of a's neighbors are in A_{n_0-1}. Then the next node can be in B_{n_0-2} (this is by property 3) and its partner is in A_{n_0-2} and so on. Finally, there is no edge from A_1 to an unmatched node in B_0: this is because there is no negative edge incident to an unmatched node and we know all edges in $A_1 \times B_0$ are negative edges (by property 2).
So the *shortest* alternating path ρ from an unmatched $a \in A_{n_0-1}$ to an unmatched $b \in B_0$ moves across sets as follows: $A_{n_0-1} - B_{n_0-1} - A_{n_0-1} - B_{n_0-2} - A_{n_0-2} - B_{n_0-3} - \cdots - A_1 - B_0 - A_0 - B_0$. This implies there are at least $n_0 + 1$ nodes in A. However $|A| = n_0$. So there is no such alternating path, i.e., there is no augmenting path with respect to S'. In other words, S' is a maximum matching in G. This finishes our proof of these six properties.

Theorem 6 shows that every stable matching in G^* maps to a popular max-matching in G. In fact, as we show next, *every* popular max-matching in G has to be realized in this manner. This is the tough part of the proof and as mentioned earlier, we will use LP-duality here. We will see that appropriate dual certificates capture a very useful feature of popular max-matchings.

3 Proving Surjectivity in a Special Case

In this section we consider the case when G admits a perfect matching. Let M be a popular perfect matching in G. So no perfect matching in G is more popular than M.

Consider the following linear program (LP1) that computes a max-weight (wrt wt_M) perfect matching in G. For any node u, let $\delta(u)$ be the set of edges incident to u.

$$\text{maximize} \sum_{e \in E} \mathsf{wt}_M(e) \cdot x_e \qquad \text{(LP1)}$$

subject to

$$\sum_{e \in \delta(u)} x_e = 1 \quad \forall u \in A \cup B \qquad \text{and} \qquad x_e \geq 0 \quad \forall e \in E.$$

It follows from the definition of the function wt_M that the optimal value of (LP1) is $\max_N \Delta(N, M)$ where N is a perfect matching in G. So if M is a popular perfect matching then the optimal value of (LP1) is 0, which is $\Delta(M, M)$, i.e., the edge incidence vector of M is an optimal solution to (LP1). The linear program (LP2) is the dual of (LP1). Hence if M is a popular perfect matching then there exists a dual feasible $\vec{\alpha}$ such that $\sum_{u \in A \cup B} \alpha_u = 0$.

$$\text{minimize} \sum_{u \in A \cup B} y_u \qquad \text{(LP2)}$$

subject to

$$y_a + y_b \;\geq\; \mathsf{wt}_M(a, b) \quad \forall\, (a, b) \in E.$$

Let $\vec{\alpha}$ be an optimal solution to (LP2). Observe that there exists an integral optimal solution to (LP2) since the constraint matrix is totally unimodular. Thus we can assume that $\vec{\alpha} \in \mathbb{Z}^{2n_0}$, where $n_0 = |A| = |B|$.

▶ **Lemma 7.** *If M is a popular perfect matching in G then there exists an optimal solution $\vec{\alpha}$ to* (LP2) *such that $\alpha_a \in \{0, -2, -4, \ldots, -2(n_0-1)\}$ for all $a \in A$ and $\alpha_b \in \{0, 2, 4, \ldots, 2(n_0-1)\}$ for all $b \in B$.*

Proof. The dual feasibility constraints are $\alpha_a + \alpha_b \geq \mathsf{wt}_M(a, b)$ for all $(a, b) \in E$. For each edge $(a, b) \in M$: $\alpha_a + \alpha_b = \mathsf{wt}_M(a, b) = 0$ by complementary slackness. Since $\alpha_b = -\alpha_a$ for $(a, b) \in M$ and because $\mathsf{wt}_M(e) \in \{0, \pm 2\}$ for each edge e, we can assume that in the sorted order of distinct α-values taken by nodes in A, for any two consecutive values $\alpha_{a'}, \alpha_{a''}$, where $\alpha_{a'} > \alpha_{a''}$, we have $\alpha_{a'} - \alpha_{a''} = 2$.

Thus we can assume that $\alpha_a \in \{k, k-2, k-4, \ldots, k-2(n_0-1)\}$ for all $a \in A$ and $\alpha_b \in \{-k, -k+2, -k+4, \ldots, -k+2(n_0-1)\}$ for all $b \in B$, for some $k \in \mathbb{Z}$. Observe that k has no impact on the objective function $\sum_{u \in A \cup B} \alpha_u$. This is because $|A| = |B|$ and so k's and $-k$'s cancel each other out.

Let us update $\vec{\alpha}$ as follows: $\alpha_a = \alpha_a - k$ for every $a \in A$ and $\alpha_b = \alpha_b + k$ for every $b \in B$. The updated vector $\vec{\alpha}$ continues to be dual feasible since $\alpha_a + \alpha_b$, for any edge (a, b), is unchanged by this update. Thus there is an optimal solution $\vec{\alpha}$ to (LP2) such that $\alpha_a \in \{0, -2, \ldots, -2(n_0-1)\}$ for all $a \in A$ and $\alpha_b \in \{0, 2, \ldots, 2(n_0-1)\}$ for all $b \in B$. ◀

Let M be a popular perfect matching in G. In order to define a stable matching S in G^* such that $M = S'$ (the set S' is defined above Theorem 6), we will use the vector $\vec{\alpha}$ described in Lemma 7. Since M is perfect, we know that for any $a \in A$, there is an edge $(a, b) \in M$ for some neighbor b of a. Recall that $\alpha_a + \alpha_b = \mathsf{wt}_M(a, b) = 0$ by complementary slackness. We will include the edge $(a_i, \tilde{b})$ in S where $\alpha_a = -2i$ and $\alpha_b = 2i$. Thus we define S as follows:

$$S = \cup_{i=0}^{n_0-1}\{(a_i, \tilde{b}) : (a, b) \in M \text{ and } \alpha_a = -2i, \alpha_b = 2i\} \cup \{\text{necessary edges incident to dummy nodes in } G^*\}.$$

In more detail, the edges incident to dummy nodes that are present in S are as follows: for each $a \in A$, these edges are $(a_j, d_{j+1}(a))$ for $0 \leq j \leq i-1$ and $(a_j, d_j(a))$ for $i+1 \leq j \leq n_0-1$, where $\alpha_a = -2i$.

Since $(a_i, \tilde{b}) \in S$, all the n_0 nodes $a_0, \ldots, a_{n_0-1}$ and the dummy nodes $d_1(a), \ldots, d_{n_0-1}(a)$ corresponding to a in G^* are matched in S. This holds for every $a \in A$. Also every $\tilde{b} \in \tilde{B}$ is matched in S since M is a perfect matching in G. Thus S is a perfect matching in G^*. It is easy to check that $S' = M$. What we need to prove is the stability of S in G^*.

▶ **Theorem 8.** *The matching S is stable in G^*.*

Proof. We need to show there is no edge in G^* that blocks S. There is no blocking edge incident to a dummy node: this is because a dummy node $d_i(a)$ has only two neighbors and when $d_i(a)$ is matched in S to its second choice neighbor a_i, its top choice neighbor a_{i-1} prefers its partner in S to $d_i(a)$.

Let us now show that no node in $a_0, \ldots, a_{n_0-1}$ has a blocking edge incident to it, for any $a \in A$. Let $(a_i, \tilde{b}) \in S$ where $(a, b) \in M$. All of $a_{i+1}, \ldots, a_{n_0-1}$ are matched to their respective top choice neighbors $d_{i+1}(a), \ldots, d_{n_0-1}(a)$. So there is no blocking edge incident to any of $a_{i+1}, \ldots, a_{n_0-1}$.

All of $a_0, \ldots, a_{i-1}$ are matched to their last choice neighbors – these are the dummy nodes $d_1(a), \ldots, d_i(a)$, respectively. Consider any neighbor $w \in B$ of a. We need to show that $\tilde{w} \in \tilde{B}$ is matched in S to a neighbor preferred to all of $a_0, \ldots, a_{i-1}$. We have $\alpha_a + \alpha_w \geq \mathsf{wt}_M(a, w)$. Since $\alpha_a = -2i$ and $\mathsf{wt}_M(e) \geq -2$ for every edge e, it follows that $\alpha_w \geq 2i - 2$.

So $(z, w) \in M$ for some neighbor z of w such that $\alpha_z = -\alpha_w \leq -(2i-2)$. Equivalently, $(z_j, \tilde{w}) \in S$ where $j \geq i-1$. Thus there is no blocking edge between $\tilde{w}$ and any of $a_0, \ldots, a_{i-2}$ by $\tilde{w}$'s preference order in G^*. We will now show that $(a_{i-1}, \tilde{w})$ cannot be a blocking edge.

- If $j \geq i$ then by $\tilde{w}$'s preference order in G^*, $\tilde{w}$ prefers z_j to a_{i-1} and so $(a_{i-1}, \tilde{w})$ does not block S.
- If $j = i-1$ then $\mathsf{wt}_M(a, w) \leq \alpha_a + \alpha_w = -2i + 2i - 2 = -2$. So both a and w prefer their respective partners in M to each other. Thus $\tilde{w}$ prefers z_{i-1} to a_{i-1}. So $(a_{i-1}, \tilde{w})$ does not block S.

Finally, we need to show there is no blocking edge incident to a_i. By the above arguments, we only need to consider edges $(a_i, \tilde{w})$ where $(z_i, \tilde{w}) \in S$. So $\mathsf{wt}_M(a, w) \leq \alpha_a + \alpha_w = -2i + 2i = 0$. Hence either $(a, w) \in M$ or at least one of a, w prefers its partner in M to the other. So either $(a_i, \tilde{w}) \in S$ or at least one of $a_i, \tilde{w}$ prefers its partner in S to the other; thus the edge $(a_i, \tilde{w})$ does not block S. Hence S is a stable matching in G^*. ◀

4 The General Case

We showed that when G has a perfect matching, our mapping from the set of stable matchings in G^* to the set of popular max-matchings in G is surjective. Now we look at the general case, i.e., G need not have a perfect matching. Let M be a popular max-matching in G.

Let $U \subseteq A \cup B$ be the set of nodes left unmatched in M. Consider (LP3) that computes a max-weight perfect matching (with respect to wt_M) in the subgraph G' induced on $V = (A \cup B) \setminus U$. Let E' be the edge set of G'. For any $v \in V$, let $\delta'(v) = \delta(v) \cap E'$.

$$\text{maximize} \sum_{e \in E'} \mathsf{wt}_M(e) \cdot x_e \qquad \text{(LP3)}$$

$$\text{subject to}$$

$$\sum_{e \in \delta'(v)} x_e = 1 \quad \forall v \in V \qquad \text{and} \qquad x_e \geq 0 \quad \forall e \in E'.$$

The optimal value of (LP3) is $\max_N \Delta(N, M)$ where N is a perfect matching in G'. Any perfect matching in G' is a maximum matching in G and since M is a popular max-matching in G, $\Delta(N, M) \leq 0$ for any perfect matching N in G'. Since $\Delta(M, M) = 0$, the edge incidence vector of M is an optimal solution to (LP3). The linear program (LP4) is the dual of (LP3).

$$\text{minimize} \sum_{u \in V} y_u \qquad \text{(LP4)}$$

subject to

$$y_a + y_b \;\geq\; \mathsf{wt}_M(a,b) \quad \forall\, (a,b) \in E'.$$

Since M is a popular max-matching, the optimal value of (LP3) is 0. So there exists an optimal solution $\vec{\alpha}$ to (LP4) such that $\sum_{u \in V} \alpha_u = 0$. Moreover, we can assume the following (see Lemma 7) where $A' = A \setminus U$ and $B' = B \setminus U$. Here $|A'| = |B'| = n_0$.

1. $\alpha_a \in \{0, -2, -4, \ldots, -2(n_0 - 1)\}$ for all $a \in A'$
2. $\alpha_b \in \{0, 2, 4, \ldots, 2(n_0 - 1)\}$ for all $b \in B'$.

For any $T \subseteq A \cup B$, let $\mathsf{Nbr}(T)$ be the set of neighbors in G of nodes in T. Theorem 9 is our main technical result here. Let $U_A = U \cap A$ and $U_B = U \cap B$.

▶ **Theorem 9.** *Let M be a popular max-matching in G and let U be the set of nodes left unmatched in M. There exists an optimal solution $\vec{\alpha}$ to* (LP4) *such that*
- $\alpha_a \in \{0, -2, \ldots, -2(n_0 - 1)\}$ *for* $a \in A'$ *and* $\alpha_b \in \{0, 2, \ldots, 2(n_0 - 1)\}$ *for* $b \in B'$
- *(i)* $\alpha_a = 0$ *for* $a \in \mathsf{Nbr}(U_B)$ *and (ii)* $\alpha_b = 2(n_0 - 1)$ *for* $b \in \mathsf{Nbr}(U_A)$.

Let us first finish our proof of surjectivity by assuming Theorem 9. Then we will prove Theorem 9. Let M be any popular max-matching in G. Corresponding to M, there is a vector $\vec{\alpha}$ as given in Theorem 9. We will use this vector $\vec{\alpha}$ to construct $S = \cup_{i=0}^{n_0-1}\{(a_i, \tilde{b}) : (a,b) \in M$ and $\alpha_a = -2i, \alpha_b = 2i\} \cup \{$necessary edges incident to dummy nodes in $G^*\}$.

The edges in S incident to dummy nodes are: $(a_j, d_{j+1}(a))$ for $0 \leq j \leq n_0 - 2$ for $a \in U_A$. For $a \in A \setminus U_A$, these are $(a_j, d_{j+1}(a))$ for $0 \leq j \leq i-1$ and $(a_j, d_j(a))$ for $i+1 \leq j \leq n_0 - 1$, where $\alpha_a = -2i$. It is easy to see that S is a matching in G^* and $S' = M$.

▶ **Theorem 10.** *The matching S is stable in G^*.*

Proof. Let E' be the edge set of G', where G' is the subgraph of G induced on $(A \cup B) \setminus U$. Consider any edge $(a_j, \tilde{w})$ in G^* where $(a,w) \in E'$ and $0 \leq j \leq n_0 - 1$. The proof of Theorem 8 shows that $(a_j, \tilde{w})$ does not block S.

Consider any edge (a,w) in $E \setminus E'$. Such an edge has a node in U as an endpoint. A useful observation is that every node in $\mathsf{Nbr}(U)$ has to be matched in M to some neighbor that it prefers to all its neighbors in U. Otherwise M would not be a popular max-matching.

- Suppose $a \in U_A$. So $w \in \mathsf{Nbr}(U_A)$ and $\alpha_w = 2(n_0 - 1)$ by property (ii) in Theorem 9. So $(z_{n_0-1}, \tilde{w}) \in S$ for some neighbor z that w prefers to a. Thus $(a_{n_0-1}, \tilde{w})$ does not block S. For $i \in \{0, \ldots, n_0 - 2\}$, none of the edges $(a_i, \tilde{w})$ can block S (by $\tilde{w}$'s preference order).
- Suppose $w \in U_B$. So $a \in \mathsf{Nbr}(U_B)$ and $\alpha_a = 0$ by property (i) in Theorem 9. Thus $(a_0, \tilde{b}) \in S$ for some neighbor b that a prefers to w. Hence $(a_0, \tilde{w})$ does not block S. Moreover, none of the edges $(a_i, \tilde{w})$ for $i \in \{1, \ldots, n_0 - 1\}$ can block S since $a_1, \ldots, a_{n_0-1}$ are matched to their respective top choice neighbors $d_1(a), \ldots, d_{n_0-1}(a)$.

Finally, no edge incident to a dummy node blocks S (by the same argument as given in the proof of Theorem 8). Hence S is a stable matching in G^*. ◀

Thus Theorem 9 allows us to show that for any popular max-matching M in G, there is a stable matching S in G^* such that $M = S'$. We will now prove Theorem 9.

Proof of Theorem 9. We know there is an optimal solution $\vec{\alpha}$ to (LP4) where $\alpha_a \in \{0, -2, \ldots, -2(n_0-1)\}$ for $a \in A'$ and $\alpha_b \in \{0, 2, \ldots, 2(n_0-1)\}$ for $b \in B'$. We now update $\vec{\alpha}$ so that it remains an optimal solution to (LP4) in the above format and it also satisfies properties (i) and (ii) given in the theorem statement.

Property (ii). Suppose the vector $\vec{\alpha} \in \{0, \pm 2, \ldots, \pm 2(n_0-1)\}^{2n_0}$ does not satisfy property (ii). So we have to update $\vec{\alpha}$ so that property (ii) is satisfied. First, we increase the α-values of the nodes in $\mathsf{Nbr}(U_A)$ to $2(n_0-1)$ and decrease the α-values of their partners in M to $-2(n_0-1)$. Now $\vec{\alpha}$ may no longer be a feasible solution to (LP4).

We use the following three update rules for all $a \in A'$ to make $\vec{\alpha}$ feasible again. Let $\alpha_a = -2i$ where $i \in \{0, \ldots, n_0-1\}$. Suppose there is some $(a,b) \in E'$ with $\alpha_a + \alpha_b < \mathsf{wt}_M(a,b)$. Let $M(b)$ be b's partner in M.

- *Rule 1.* If $\mathsf{wt}_M(a,b) = 0$ then update $\alpha_b = 2i$ and $\alpha_{M(b)} = -2i$.
- *Rule 2.* If $\mathsf{wt}_M(a,b) = -2$ then update $\alpha_b = 2(i-1)$ and $\alpha_{M(b)} = -2(i-1)$.
- *Rule 3.* If $\mathsf{wt}_M(a,b) = 2$ then update $\alpha_b = 2(i+1)$ and $\alpha_{M(b)} = -2(i+1)$.

At the onset, $\vec{\alpha}$ was a feasible solution to (LP4), so $\alpha_a + \alpha_b \geq \mathsf{wt}_M(a,b)$ for $(a,b) \in E'$. Then we moved the nodes in $\mathsf{Nbr}(U_A)$ and their partners in M to sets B_{n_0-1} and A_{n_0-1}, respectively, where $A_i = \{a \in A' : \alpha_a = -2i\}$ and $B_i = \{b \in B' : \alpha_b = 2i\}$ for all i. The subscript i will be called the *level* of nodes in $A_i \cup B_i$.

The nodes that moved to A_{n_0-1} have a lower α-value than earlier and it is these nodes that "pull" their neighbors upwards to higher levels as given by rules 1-3. Let a be a new node in level i and let b be a neighbor of a such that $\alpha_a + \alpha_b < \mathsf{wt}_M(a,b)$. Then b and $M(b)$ move to: (1) level i if $\mathsf{wt}_M(a,b) = 0$, (2) level $i-1$ if $\mathsf{wt}_M(a,b) = -2$, else (3) level $i+1$, i.e., if $\mathsf{wt}_M(a,b) = 2$.

In turn, the nodes in A' that have moved to these higher levels by rules 1-3 pull their neighbors and the partners of these neighbors upwards to higher levels by these rules. Thus we may get further new nodes in B_{n_0-1}, A_{n_0-1} and so on. While any of rules 1-3 is applicable, we apply that rule. So a rule may be applied many times to the same edge in E'.

Claim 11 (proved in Section 4.1) shows a useful property. We show in its proof that such a blocking edge creates a *forbidden* alternating cycle/path wrt M, as given in Theorem 5.

▷ Claim 11. By applying the above rules, suppose a node $v_0 \in A'$ moves to A_{n_0-1}. Then there is no blocking edge (e such that $\mathsf{wt}_M(e) = 2$) incident to v_0.

Applying rules 1-3 increases the α-values of some nodes in B' and it never decreases the α-value of any node in B'. The nodes in B' with increased α-values and their partners have moved to higher levels (see Fig. 1). This upwards movement of nodes has to terminate at level $n_0 - 1$. For the α-value of any $b \in B'$ to be increased beyond $2(n_0-1)$, we need a blocking edge (a,b) where $a \in A_{n_0-1}$ – this would cause rule 3 to be applied which would increase α_b to $2n_0$. However there is no such blocking edge (by Claim 11).

Since there are n_0 levels and because $|B'| = n_0$, there can be at most n_0^2 applications of these rules. When no rule is applicable, $\vec{\alpha}$ is a feasible solution to (LP4). Moreover, $\sum_{u \in V} \alpha_u$ is invariant under this update of α-values, since we maintain $\alpha_a + \alpha_b = 0$ for every $(a,b) \in M$. Hence $\vec{\alpha}$ is an optimal solution to (LP4). Thus for every popular max-matching M, there is an optimal solution $\vec{\alpha}$ to (LP4) in the desired format that satisfies property (ii).

Property (i). We now have an optimal solution $\vec{\alpha} \in \{0, \pm 2, \ldots, \pm 2(n_0 - 1)\}^{2n_0}$ to (LP4), where $\alpha_a \le 0$ for all $a \in A'$ and $\alpha_b \ge 0$ for all $b \in B'$, such that $\alpha_b = 2(n_0 - 1)$ for all $b \in \mathsf{Nbr}(U_A)$. Suppose property (i) is not satisfied.

Then we increase α-values of certain nodes in A' – this moves these nodes *downwards* with respect to their level (see Fig. 1) and ensures that property (i) holds. First, we increase the α-values of the nodes in $\mathsf{Nbr}(U_B)$ to 0 and their partners also have α-values updated to 0. Now $\vec{\alpha}$ may no longer be a feasible solution to (LP4).

So we will use the following three update rules for all $b \in B'$. Let $\alpha_b = 2i$ where $i \in \{0, \ldots, n_0 - 1\}$. Suppose there is an edge $(a, b) \in E'$ such that $\alpha_a + \alpha_b < \mathsf{wt}_M(a, b)$. Let $M(a)$ be a's partner in M.

- *Rule 4.* If $\mathsf{wt}_M(a, b) = 0$ then update $\alpha_a = -2i$ and $\alpha_{M(a)} = 2i$.
- *Rule 5.* If $\mathsf{wt}_M(a, b) = -2$ then update $\alpha_a = -2(i + 1)$ and $\alpha_{M(a)} = 2(i + 1)$.
- *Rule 6.* If $\mathsf{wt}_M(a, b) = 2$ then update $\alpha_a = -2(i - 1)$ and $\alpha_{M(a)} = 2(i - 1)$.

Applying rules 4-6 increases the α-values of some nodes in A' and it never decreases the α-value of any node in A'. The nodes in A' with increased α-values and their partners have moved to lower levels. Moreover, the movement of nodes downwards has to stop at level 0 since no blocking edge can be incident to any node that moves to B_0 (analogous to Claim 11). While any of the above three rules is applicable, we apply that rule.

When no rule is applicable, $\vec{\alpha}$ is a feasible solution to (LP4). Since $\sum_{u \in V} \alpha_u = 0$, $\vec{\alpha}$ is an optimal solution to (LP4). So there is an optimal solution $\vec{\alpha}$ to (LP4) in the desired format that satisfies property (i).

Properties (i) and (ii). Note that we cannot claim straightaway that the above $\vec{\alpha}$ satisfies both property (i) and property (ii). This is because applying rules 4-6 may have caused $\alpha_b < 2(n_0 - 1)$ for some $b \in \mathsf{Nbr}(U_A)$. Claim 12 shows this is not possible.

▷ Claim 12. The above $\vec{\alpha}$ satisfies property (ii), i.e., $\alpha_b = 2(n_0 - 1)$ for $b \in \mathsf{Nbr}(U_A)$.

Claim 12 is proved in Section 4.1. So we have an optimal solution $\vec{\alpha}$ to (LP4) in the desired format that satisfies both property (i) and property (ii). ◀

4.1 Proofs of Claim 11 and Claim 12

The proofs of Claim 11 and Claim 12 use the fact that certain alternating cycles/paths are forbidden for popular max-matchings. These include the ones given in Theorem 5 and also augmenting paths (since M is a maximum matching).

Proof of Claim 11. Let v_0 be the first node that moves to A_{n_0-1} with a blocking edge incident to it. Recall our update procedure – we initially added nodes in $\mathsf{Nbr}(U_A)$ and their partners in M to B_{n_0-1} and A_{n_0-1}, respectively. Then we applied rules 1-3 in some order and this resulted in the node v_0 moving to A_{n_0-1}. Corresponding to these rules, we will construct an alternating path $p = v_0 - M(v_0) - v_1 - M(v_1) - v_2 - \cdots - v_k - M(v_k) - u$ between v_0 and some node $u \in U_A$.

The path p can be partitioned into $k+1$ pairs of edges for some $k \ge 0$. For $0 \le i \le k-1$: the i-th pair consists of the matching edge $(v_i, M(v_i))$ of weight 0 and the non-matching edge $e_i = (M(v_i), v_{i+1})$ where v_{i+1} is the node that pulled $M(v_i)$ and v_i to their current level due to the application of one of the above three rules. Rule 1 implies $\mathsf{wt}_M(e_i) = 0$ while rule 2 implies $\mathsf{wt}_M(e_i) = -2$ and rule 3 implies $\mathsf{wt}_M(e_i) = 2$.

Observe that rule 1 places $M(v_i)$ in the same level as v_{i+1} while rule 2 places $M(v_i)$ one level lower than v_{i+1} and rule 3 places $M(v_i)$ one level higher than v_{i+1}. The last pair of edges in p are $(v_k, M(v_k)) \in A_{n_0-1} \times B_{n_0-1}$ and $(M(v_k), u)$, where the latter edge has weight 0. The node v_0 is in level $n_0 - 1$ and the node $M(v_k)$ is also in level $n_0 - 1$. So v_0 and v_k are at the same level, hence the number of edges in p of weight -2 is exactly the same as the number of edges of weight 2, thus $\mathsf{wt}_M(p) = 0$.

Suppose there is a blocking edge (v_0, w). If w belongs to p, then it is easy to see that the alternating cycle C obtained by joining the endpoints of the v_0-w subpath in p with the edge (v_0, w) satisfies $\mathsf{wt}_M(C) \geq 2$. This contradicts Theorem 5 since M is a popular max-matching. Hence w does not belong to path p. So let us add the 2-edge path $M(w) - w - v_0$ as a prefix to the v_0-u path p and call this alternating path q: we have $\mathsf{wt}_M(q) = \mathsf{wt}_M(p) + \mathsf{wt}_M(v_0, w) = 2$. Since $\mathsf{wt}_M(q) > 0$ and the unmatched node u is an endpoint of q, this again contradicts Theorem 5. Hence there is no blocking edge incident to v_0. ◁

Proof of Claim 12. Suppose there is a node $w_1 \in \mathsf{Nbr}(U_A)$ such that $\alpha_{w_1} < 2(n_0 - 1)$. So there is some $u \in U_A$ such that $(u, w_1) \in E$ and though $\alpha_{w_1} = 2(n_0 - 1)$ just before we started applying rules 4-6, the application of these rules caused α_{w_1} to become less than $2(n_0 - 1)$.

Initially we added nodes in $\mathsf{Nbr}(U_B)$ and their partners in M to A_0 and B_0, respectively. Then we repeatedly applied rules 4-6 and this resulted in w_1 moving to a level lower than $n_0 - 1$. Corresponding to what caused w_1 to be "pulled" downwards, we will construct an alternating path $p = w_1 - M(w_1) - w_2 - M(w_2) - \cdots - w_r - M(w_r) - u'$ between w_1 and a node $u' \in U_B$.

The path p will consist of r pairs of edges for some $r \geq 1$. For $1 \leq i \leq r - 1$: the i-th pair of edges is $(w_i, M(w_i))$ and $(M(w_i), w_{i+1})$ where w_{i+1} is the node that pulled $M(w_i)$ and w_i to their current level due to the application of one of rules 4-6. The last pair of edges in p is $(w_r, M(w_r))$ and $(M(w_r), u')$, where $w_r \in B_0, M(w_r) \in A_0$, and $u' \in U_B$. Thus we have an alternating path p between w_1 and $u' \in U_B$.

By adding the edge (u, w_1) as a prefix to the path p, we get an augmenting path $u - w_1 - \cdots - M(w_r) - u'$ with respect to M. However there cannot be any augmenting path wrt M since M is a maximum matching in G. Thus $\vec{\alpha}$ satisfies property (ii). ◁

This finishes the proof of Theorem 9. So the stable matching polytope of G^* yields a compact extended formulation for the popular max-matching polytope of G (by Theorem 6 and Theorem 10). This formulation is described in Section 4.2.

Linear programming on this formulation with $\min \sum_{e \in E} c(e) \cdot x_e$ as the objective function computes a min-cost popular max-matching in G in polynomial time. Equivalently, we can compute a min-cost stable matching S in G^* and return the corresponding matching S' in G. It follows from Theorem 6 and Theorem 10 that S' is a min-cost popular max-matching in G. This proves Theorem 3 stated in Section 1.

4.2 An Extended Formulation for the Popular Max-Matching Polytope

For any node u in G^*, let $\{v' \succ_u v\}$ be the set of all neighbors of u in G^* that it prefers to v. Let $\delta^*(u)$ denote the set of edges incident to u in G^*.

Let $T = \cup_{a \in A}(\{a_0, \ldots, a_{n_0-2}\} \cup \{d_1(a), \ldots, d_{n_0-1}(a)\})$. Every node in T is a top choice neighbor for some node, so every node in T has to be matched in all stable matchings in G^*. The constraints given below describe the stable matching polytope of G^*, as shown in [30].

$$\sum_{w \succ_{a_i} \tilde{b}} x_{(a_i,w)} + \sum_{z \succ_{\tilde{b}} a_i} x_{(z,\tilde{b})} + x_{(a_i,\tilde{b})} \geq 1 \qquad \forall (a_i,\tilde{b}) \in E^*$$

$$\sum_{e \in \delta^*(u)} x_e = 1 \qquad \forall\, u \in T$$

$$\sum_{e \in \delta^*(u)} x_e \leq 1 \quad \forall u \in A^* \cup B^* \qquad \text{and} \qquad x_e \geq 0 \qquad \forall e \in E^*.$$

1. The topmost constraint captures the *stability* constraint for edge $(a_i, \tilde{b}) \in E^*$ where $(a,b) \in E$ and $0 \leq i \leq n_0 - 1$.
2. The constraint in the second line for $u = a_{i-1}$ captures the stability constraint for the edge $(a_{i-1}, d_i(a))$ and for $u = d_i(a)$ captures the stability constraint for the edge $(a_i, d_i(a))$.
3. The constraints in the third line capture that $\vec{x}$ belongs to the matching polytope of G^*.

Consider the equations $x_{(a,b)} = \sum_{i=0}^{n_0-1} x_{(a_i,\tilde{b})}$ for all $(a,b) \in E$. It follows from Theorem 6 and Theorem 10 that these m equations along with the constraints of the stable matching polytope of G^* given above describe an extended formulation for the popular max-matching polytope $\mathcal{M}_G$. So the extension complexity of the polytope $\mathcal{M}_G$ is $O(mn)$.

5 A Hardness Result

In this section we show that it is NP-hard to compute a min-cost Pareto-optimal matching and a min-cost Pareto-optimal max-matching in an instance $G = (A \cup B, E)$ with edge costs in $\{0,1\}$.

Given a 3SAT formula ψ, we will build an instance G_ψ with edge costs in $\{0,1\}$ such that G_ψ admits a Pareto-optimal matching of cost 0 if and only if ψ is satisfiable. Any Pareto-optimal matching of cost 0 would have to be a perfect matching in G_ψ. Hence this will prove the NP-hardness of both the min-cost Pareto-optimal matching problem and the min-cost Pareto-optimal max-matching problem.

Our reduction resembles a hardness reduction from [9] that showed the NP-hardness of deciding if an instance G has a stable matching M that is also dominant. As done in this reduction, we will first transform ψ so that every clause contains either only positive literals or only negative literals; moreover, there will be a single occurrence of each negative literal in the transformed ψ. This is easy to achieve:

- let $X_1, \ldots, X_n$ be the starting variables. For $i \in [n]$: replace all occurrences of $\neg X_i$ with the same variable X_{n+i} (a new one) and add the two clauses $(X_i \vee X_{n+i}) \wedge (\neg X_i \vee \neg X_{n+i})$ to capture $\neg X_i \equiv X_{n+i}$. Thus there are $2n$ variables in the transformed ψ.

We build the graph G_ψ as follows. There are two types of gadgets: those that correspond to positive clauses and those that correspond to negative clauses. Fig. 2 (resp., Fig. 3) shows how a positive (resp., negative) clause gadget looks like.

We now describe the preference lists of nodes in a positive clause $C_\ell = x \vee y \vee z$ (see Fig. 2). The nodes a_x, a'_x, b_x, b'_x occur in x's gadget and a_y, a'_y, b_y, b'_y occur in y's gadget and a_z, a'_z, b_z, b'_z occur in z's gadget: these gadgets are in the ℓ-th clause gadget C_ℓ. Every occurrence of a literal has a separate gadget.

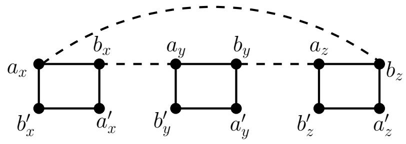

Figure 2 The clause gadget for a positive clause $C_\ell = x \vee y \vee z$. Every occurrence of a literal in ψ has a separate gadget. So we ought to use labels such as $a_{x,\ell}, b_{x,\ell}, \ldots$ here; for the sake of simplicity, we used the labels $a_x, b_x, \ldots$ here.

a_x	a'_x	a_y	a'_y	a_z	a'_z
b_z	b_x	b_x	b_y	b_y	b_z
b_x	b'_x	b_y	b'_y	b_z	b'_z
d'_x	–	d'_y	–	d'_z	–
b'_x	–	b'_y	–	b'_z	–

Here a_x's top choice is b_z, second choice b_x, third choice d'_x, fourth choice b'_x, and similarly for other nodes. For every occurrence of a positive literal x: there will be a pair of *consistency edges* – the pair (a_x, d'_x) and (b'_x, c_x) in Fig. 4 – between this gadget of x and the unique gadget of $\neg x$. In our preferences, the neighbors on consistency edges are marked in red.

b_x	b'_x	b_y	b'_y	b_z	b'_z
a_y	a'_x	a_z	a'_y	a_x	a'_z
a_x	c_x	a_y	c_y	a_z	c_z
a'_x	a_x	a'_y	a_y	a'_z	a_z

The preference lists of nodes that occur in a clause gadget with 2 positive literals will be totally analogous to the preference lists of nodes in a clause gadget with 3 positive literals.

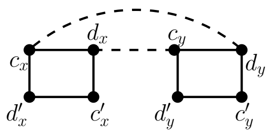

Figure 3 A clause gadget corresponding to a negative clause $D_k = \neg x \vee \neg y$; due to our transformation of ψ, every negative clause has only 2 literals.

We will now describe the preference lists of nodes in a negative clause k – the overall picture here is given in Fig. 3.

c_x	c'_x	c_y	c'_y
d_y	d_x	d_x	d_y
d_x	d'_x	d_y	d'_y
$b'_{x,i}$	–	$b'_{y,i'}$	–
$\cdots$	–	$\cdots$	–
$b'_{x,j}$	–	$b'_{y,j'}$	–
d'_x	–	d'_y	–

d_x	d'_x	d_y	d'_y
c_y	c'_x	c_x	c'_y
c_x	$a_{x,i}$	c_y	$a_{y,i'}$
c'_x	$\cdots$	c'_y	$\cdots$
–	$a_{x,j}$	–	$a_{y,j'}$
–	c_x	–	c_y

The nodes c_x, c'_x, d_x, d'_x and c_y, c'_y, d_y, d'_y occur in the gadgets of $\neg x$ and $\neg y$, respectively. The nodes $b'_{x,i}, \ldots, b'_{x,j}$ (resp., $b'_{y,i'}, \ldots, b'_{y,j'}$) in the preference lists above are the b'-nodes in the x-gadgets (resp., y-gadgets) in the various clauses that x (resp., y) occurs in. Similarly, $a_{x,i}, \ldots, a_{x,j}$ (resp., $a_{y,i'}, \ldots, a_{y,j'}$) are the a-nodes in the x-gadgets (resp., y-gadgets) in the various clauses that x (resp., y) occurs in. The preference order among the b'-nodes and among the a-nodes in these lists is not important. The consistency edges between a gadget of x and the gadget of $\neg x$ are shown in Fig. 4.

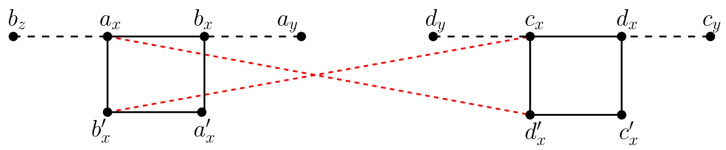

Figure 4 For the sake of simplicity, we use a_x, b_x, a'_x, b'_x to denote the 4 nodes in the gadget of x in the ℓ-th clause; c_x, d_x, c'_x, d'_x are the 4 nodes in the unique gadget of $\neg x$. The consistency edges are the red dashed edges.

Edge costs. For each edge e in G_ψ, we will set $\mathsf{cost}(e) \in \{0, 1\}$ as follows.

- For each variable $r \in \{X_1, \ldots, X_{2n}\}$: set $\mathsf{cost}(e) = 0$ where e is any of the 4 edges in any literal gadget $\langle a_r, b_r, a'_r, b'_r \rangle$ of r or any of the 4 edges in the gadget $\langle c_r, d_r, c'_r, d'_r \rangle$ of $\neg r$.
- For all other edges e, set $\mathsf{cost}(e) = 1$.

In particular, for any edge e in the consistency pair for any variable, we have $\mathsf{cost}(e) = 1$. In our figures, all dashed edges have cost 1 and all solid edges have cost 0.

Let M be a Pareto-optimal matching in G_ψ with $\mathsf{cost}(M) = 0$. So M has to use only cost 0 edges. Thus M is forbidden to use any edge other than the 4 edges in the gadget of any literal. Moreover, since M is Pareto-optimal, M cannot leave two adjacent nodes unmatched. Thus for $r \in \{X_1, \ldots, X_{2n}\}$:

1. From a gadget of r (say, on nodes a_r, b_r, a'_r, b'_r), either (i) $(a_r, b_r), (a'_r, b'_r)$ are in M or (ii) (a_r, b'_r), (a'_r, b_r) are in M.
2. From the gadget of $\neg r$ (the nodes are c_r, d_r, c'_r, d'_r), either (i) $(c_r, d_r), (c'_r, d'_r)$ are in M or (ii) (c_r, d'_r), (c'_r, d_r) are in M.

Thus any Pareto-optimal matching in G_ψ of cost 0 is a perfect matching. Lemma 13 will be useful to us.

▶ **Lemma 13.** *Let M be a Pareto-optimal matching in G_ψ. For any $r \in \{X_1, \ldots, X_{2n}\}$, both (a_r, b'_r) and (c_r, d'_r) cannot simultaneously be in M.*

Proof. The preferences of the nodes are set such that if both (a_r, b'_r) and (c_r, d'_r) are in M then both the non-matching edges (a_r, d'_r) and (b'_r, c_r) in the alternating cycle $\rho = a_r - (d'_r, c_r) - (b'_r, a_r) - d'_r$ are blocking edges to M. Consider $M \oplus \rho$ versus M. All the 4 nodes a_r, b'_r, c_r, d'_r prefer $M \oplus \rho$ to M while the other nodes are indifferent between $M \oplus \rho$ and M. Thus $\phi(M \oplus \rho, M) = 4$ and $\phi(M, M \oplus \rho) = 0$, so $u(M) = \infty$. This means M is not Pareto-optimal, a contradiction. Thus for any $r \in \{X_1, \ldots, X_{2n}\}$, we cannot have both (a_r, b'_r) and (c_r, d'_r) in M. ◀

Theorem 14 is our main result here.

▶ **Theorem 14.** *G_ψ has a Pareto-optimal matching M with $\mathsf{cost}(M) = 0$ if and only if ψ is satisfiable.*

Proof. Suppose G_ψ has a Pareto-optimal matching M with $\mathsf{cost}(M) = 0$. For any variable $r \in \{X_1, \ldots, X_{2n}\}$, consider the edges in $\neg r$'s gadget that are in M. If $(c_r, d'_r), (c'_r, d_r)$ are in M then set $r = \mathsf{false}$ else set $r = \mathsf{true}$.

Lemma 13 tells us that when we set r to false, the edges $(a_{r,i}, b_{r,i}), (a'_{r,i}, b'_{r,i})$ from r's gadget in the i-th clause have to be in M (where $a_{r,i}, b_{r,i}, a'_{r,i}, b'_{r,i}$ are the 4 nodes from r's gadget in the i-th clause).

▷ Claim 15. The above assignment satisfies ψ.

Claim 15 uses the Pareto-optimality of M to show that every clause has at least one literal set to true. Its proof is given after the proof of Theorem 14. Hence if G_ψ admits a Pareto-optimal matching M with $\mathsf{cost}(M) = 0$, then ψ is satisfiable.

The converse. We will now show that if ψ is satisfiable then there is a Pareto-optimal matching M in G_ψ such that $\mathsf{cost}(M) = 0$. There is a natural way of constructing the matching M – we will use the satisfying assignment for ψ to choose edges from each literal gadget. For any variable r, include the following edges in the matching M:

- if $r =$ true then take the edges $(c_r, d_r), (c'_r, d'_r)$ from $\neg r$'s gadget and the edges $(a_{r,i}, b'_{r,i})$, $(a'_{r,i}, b_{r,i})$ from r's gadget in clause i (for every clause i that r belongs to).
- if $r =$ false then take the edges $(c_r, d'_r), (c'_r, d_r)$ from $\neg r$'s gadget and the edges $(a_{r,i}, b_{r,i})$, $(a'_{r,i}, b'_{r,i})$ from r's gadget in clause i (for every clause i that r belongs to).

It is easy to see that $\mathsf{cost}(M) = 0$. Since M is a perfect matching, there is no alternating path ρ wrt M such that $\phi(M, M \oplus \rho) = 0$. This is because for every alternating path ρ wrt M, we have $|M \oplus \rho| < |M|$ and the nodes matched in M and unmatched in $M \oplus \rho$ prefer M to $M \oplus \rho$, so $\phi(M, M \oplus \rho) > 0$. Hence in order to prove M's Pareto-optimality, what we need to show is Claim 16.

▷ Claim 16. There is no alternating cycle ρ with respect to M such that $\phi(M \oplus \rho, M) > 0$ and $\phi(M, M \oplus \rho) = 0$.

The proof of Claim 16 is given below. This finishes the proof of Theorem 14. ◀

Proof of Claim 15. Suppose this assignment does not satisfy ψ. We have 3 cases here.

1. Let $C_i = x \vee y \vee z$. Suppose all the three variables x, y, z are in false state. Consider the following alternating cycle ρ wrt M:

$$b_{z,i} - (a_{x,i}, b_{x,i}) - (a_{y,i}, b_{y,i}) - (a_{z,i}, b_{z,i}) - a_{x,i}.$$

All non-matching edges in this alternating cycle, i.e., the edges $(b_{z,i}, a_{x,i})$, $(b_{x,i}, a_{y,i})$, $(b_{y,i}, a_{z,i})$, are blocking edges with respect to M. In the $M \oplus \rho$ versus M comparison, these 6 nodes $a_{x,i}, b_{x,i}, a_{y,i}, b_{y,i}, a_{z,i}, b_{z,i}$ prefer $M \oplus \rho$ to M while all other nodes in G_ψ are indifferent between $M \oplus \rho$ and M. Thus we have $\phi(M \oplus \rho, M) = 6$ and $\phi(M, M \oplus \rho) = 0$. Hence $u(M) = \infty$, contradicting the Pareto-optimality of M.

2. Let $C_j = x \vee y$, i.e., this is a positive clause with 2 literals. Suppose both x and y are in false state. Consider the following alternating cycle ρ wrt M:

$$b_{y,j} - (a_{x,j}, b_{x,j}) - (a_{y,j}, b_{y,j}) - a_{x,j}.$$

In the $M \oplus \rho$ versus M comparison, the 4 nodes $a_{x,j}, b_{x,j}, a_{y,j}, b_{y,j}$ prefer $M \oplus \rho$ to M while all the other nodes in G_ψ are indifferent between $M \oplus \rho$ and M. Thus $\phi(M \oplus \rho, M) = 4$ and $\phi(M, M \oplus \rho) = 0$. Hence $u(M) = \infty$, contradicting the Pareto-optimality of M.

3. Let $D_k = \neg x \vee \neg y$. Suppose both $\neg x$ and $\neg y$ are in false state. Consider the following alternating cycle ρ wrt M:

$$d_y - (c_x, d_x) - (c_y, d_y) - c_x.$$

In the $M \oplus \rho$ versus M comparison, the 4 nodes c_x, d_x, c_y, d_y prefer $M \oplus \rho$ to M while all the other nodes in G_ψ are indifferent between $M \oplus \rho$ and M. So $\phi(M \oplus \rho, M) = 4$ and $\phi(M, M \oplus \rho) = 0$. Hence $u(M) = \infty$, contradicting the Pareto-optimality of M.

Thus every clause in ψ has at least one literal in true state. ◁

Proof of Claim 16. We need to show there is no alternating cycle ρ with respect to M such that $\phi(M \oplus \rho, M) > 0$ and $\phi(M, M \oplus \rho) = 0$. Every non-matching edge in such an alternating cycle ρ has to be a blocking edge wrt M.

First, we argue that every consistency edge is a *non-blocking* edge to M; say, this is a consistency edge corresponding to variable r in clause i. It follows from our construction of M that M contains either:

1. $(a_{r,i}, b'_{r,i}), (a'_{r,i}, b_{r,i})$ and $(c_r, d_r), (c'_r, d'_r)$ or
2. $(a_{r,i}, b_{r,i}), (a'_{r,i}, b'_{r,i})$ and $(c_r, d'_r), (c'_r, d_r)$.

- case 1: the node d'_r prefers $M(d'_r) = c'_r$ to $a_{r,i}$ and the node c_r prefers $M(c_r) = d_r$ to $b'_{r,i}$.
- case 2: the node $a_{r,i}$ prefers $M(a_{r,i}) = b_{r,i}$ to d'_r and the node $b'_{r,i}$ prefers $M(b'_{r,i}) = a'_{r,i}$ to c_r.

Thus in both cases, the consistency edges $(a_{r,i}, d'_r)$ and $(b'_{r,i}, c_r)$ are non-blocking edges to M. Let H be the subgraph of G_ψ obtained by preserving only the edges that are in M and also blocking edges wrt M. Thus no non-blocking edge (other than edges in M) is included in H – so no consistency edge belongs to H.

Since there are no consistency edges in H, any alternating cycle in H has to be contained within a single clause. We will now show there is no such cycle in H by using the fact that we constructed M using a satisfying assignment for ψ: thus every clause has at least one literal set to true.

Let $C = x \vee y \vee z$ and suppose $y = \mathsf{true}$ in ψ. Then (a_y, b'_y) and (a'_y, b_y) are in M, however the edge (a'_y, b'_y) is non-blocking wrt M and hence it is missing in H. Thus there is no alternating cycle in H that is contained within the clause C. Now consider a negative clause $D = \neg x \vee \neg y$ and suppose $x = \mathsf{false}$ in ψ. Then (c_x, d'_x) and (c'_x, d_x) are in M, however the edge (c'_x, d'_x) is non-blocking wrt M and it is missing in H. Thus there is no alternating cycle in H that is contained within the clause D.

Consider the 4 edges of any literal gadget (say, r) in G_ψ: if $(a_r, b_r) \in M$ then (a_r, b'_r) is a non-blocking edge wrt M and if $(a'_r, b_r) \in M$ then (a'_r, b'_r) is a non-blocking edge wrt M. Similarly, in the gadget of $\neg r$: if $(c_r, d_r) \in M$ then (c_r, d'_r) is a non-blocking edge wrt M and if $(c'_r, d_r) \in M$ then (c'_r, d'_r) is a non-blocking edge wrt M. Thus there is no alternating cycle wrt M in H. So there is no alternating cycle ρ in G_ψ such that $\phi(M \oplus \rho, M) > 0$ and $\phi(M, M \oplus \rho) = 0$. ◁

Since any Pareto-optimal matching in G_ψ of cost 0 is a perfect matching, Theorem 14 shows that the min-cost Pareto-optimal matching problem and the min-cost Pareto-optimal max-matching problem are NP-hard. Moreover, these problems are NP-hard to approximate to any multiplicative factor. Thus we have shown Theorem 4 stated in Section 1.

References

1 A. Abdulkadiroğlu and T. Sönmez. School choice: a mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.

2 D. J. Abraham, K. Cechlřová, D. F. Manlove, and K. Mehlhorn. Pareto optimality in house allocation problems. In *Proceedings of the 15th International Symposium on Algorithms and Computation (ISAAC)*, pages 3–15, 2004.

3 T. Barnagarwala. 1172 private doctors sign up to treat patients at govt hospitals in Mumbai. Indian Express, May 11, 2020.

4 S. Baswana, P. P. Chakrabarti, S. Chandran, Y. Kanoria, and U. Patange. Centralized admissions for engineering colleges in India. *INFORMS Journal on Applied Analytics*, 49(5):338–354, 2019.

5 P. Biro, D. F. Manlove, and S. Mittal. Size versus stability in the marriage problem. *Theoretical Computer Science*, 411:1828–1841, 2010.

6 Canadian Resident Matching Service. How the matching algorithm works. `http://carms.ca/algorithm.htm`.

7 M.-J.-A.-N. de C. (Marquis de) Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix.* L'Imprimerie Royale, 1785.

8 Á. Cseh. Popular matchings. Trends in Computational Social Choice, Ulle Endriss (ed.), 2017.

9 Á. Cseh, Y. Faenza, T. Kavitha, and V. Powers. Understanding popular matchings via stable matchings. `arXiv:1811.06897`.

10 Á. Cseh and T. Kavitha. Popular edges and dominant matchings. *Mathematical Programming*, 172(1):209–229, 2018.

11 P. Eirinakis, D. Magos, I. Mourtos, and P. Miliotis. Polyhedral aspects of stable marriage. *Mathematics of Operations Research*, 39(3):656–671, 2014.

12 Y. Faenza and T. Kavitha. Quasi-popular matchings, optimality, and extended formulations. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 325–344, 2020.

13 Y. Faenza, T. Kavitha, V. Powers, and X. Zhang. Popular matchings and limits to tractability. In *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2790–2809, 2019.

14 T. Feder. A new fixed point approach for stable networks and stable marriages. *Journal of Computer and System Sciences*, 45(2):233–284, 1992.

15 T. Feder. Network flow and 2-satisfiability. *Algorithmica*, 11(3):291–319, 1994.

16 T. Fleiner. A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research*, 28(1):103–126, 2003.

17 D. Gale and L.S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.

18 D. Gale and M. Sotomayor. Some remarks on the stable matching problem. *Discrete Applied Mathematics*, 11:223–232, 1985.

19 P. Gärdenfors. Match making: assignments based on bilateral preferences. *Behavioural Science*, 20:166–173, 1975.

20 C.-C. Huang and T. Kavitha. Popular matchings in the stable marriage problem. *Information and Computation*, 222:180–194, 2013.

21 R. W. Irving, P. Leather, and D. Gusfield. An efficient algorithm for the "optimal" stable marriage. *Journal of the ACM*, 34(3):532–543, 1987.

22 P. B. Jaiswal. Mumbai struggles to tackle increasing number of COVID-19 cases. The Week, May 6, 2020.

23 T. Kavitha. A size-popularity tradeoff in the stable marriage problem. *SIAM Journal on Computing*, 43(1):52–71, 2014.

24 T. Kavitha. Popular half-integral matchings. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 22:1–22:13, 2016.

25 T. Kavitha. Popular matchings with one-sided bias. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 70:1–70:18, 2020.

26 M. McCutchen. The least-unpopularity-factor and least-unpopularity-margin criteria for matching problems with one-sided preferences. In *Proceedings of the 8th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 593–604, 2008.

27 S. Merrill and B. Grofman. *A unified theory of voting: directional and proximity spatial models.* Cambridge University Press, 1999.

28 National Resident Matching Program. Why the Match? `http://www.nrmp.org/whythematch.pdf`.

29 P. A. Robards. Applying the two-sided matching processes to the United States Navy enlisted assignment process. Master's Thesis, Naval Postgraduate School, Monterey, Canada, 2001.

30 U. G. Rothblum. Characterization of stable matchings as extreme points of a polytope. *Mathematical Programming*, 54:57–67, 1992.

31 C.-P. Teo and J. Sethuraman. The geometry of fractional stable matchings and its applications. *Mathematics of Operations Research*, 23(4):874–891, 1998.

32 E. G. Thurber. Concerning the maximum number of stable matchings in the stable marriage problem. *Discrete Mathematics*, 248(1-3):195–219, 2002.

33 J. H. Vande Vate. Linear programming brings marital bliss. *Operations Research Letters*, 8(3):147–153, 1989.

34 W. Yang, J. A. Giampapa, and K. Sycara. Two-sided matching for the US Navy detailing process with market complication. Technical Report CMU-R1-TR-03-49, Robotics Institute, Carnegie Mellon University, 2003.