

Discrete Optimization Problems in Popular Matchings and Scheduling

Vladlena Powers

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2020

c 2020
Vladlena Powers
All Rights Reserved

ABSTRACT

Discrete Optimization Problems in Popular Matchings and Scheduling

Vladlena Powers

This thesis focuses on two central classes of problems in discrete optimization: matching and scheduling.

Matching problems lie at the intersection of different areas of mathematics, computer science, and economics. In two-sided markets, Gale-Shapley's model has been widely used and generalized to assign, e.g., students to schools and interns to hospitals. The goal is to find a matching that respects a certain concept of fairness called *stability*. This model has been generalized in many ways. Relaxing the stability condition to *popularity* allows to overcome one of the main drawbacks of stable matchings: the fact that two individuals (a *blocking pair*) can prevent the matching from being much larger. The first part of this thesis is devoted to understanding the complexity of various problems around popular matchings. Chapter 2 and Chapter 3 deal with variations of the *maximum weighted popular matching problem*. In particular, we show various NP-hardness results, while on the other hand showing that a popular matching of maximum weight (if any) can be found in polynomial time if the input graph has bounded treewidth. Chapter 4 investigates algorithmic questions on the relationship between popular, stable, and Pareto-optimal matchings.

The second part of the thesis deals with a combinatorial scheduling problem arising in cyber-security. Moving target defense strategies allow to mitigate cyber attacks. We analyze a strategic game, PLADD, which is an abstract model for these strategies. We provide an IP formulation, show it has unbounded integrality gap, and write a software that solves the problem via enumeration and visualizes the optimal solution. This is investigated in Chapter 5.

Contents

List of Figures	iii
List of Tables	ix
Acknowledgements	x
Introduction	1
Chapter 1 Preliminaries	3
1.1 Notation	3
1.2 Sets, posets, lattices	4
1.3 Matching under preferences	4
1.4 SAT and hardness of approximation	15
1.5 Treewidth: definitions and some facts	17
Chapter 2 Popular matching problems with edge weights	20
2.1 Introduction	20
2.2 The reduction	24
2.3 Proof of Theorems 2.1.1–2.1.3	43
2.4 Proof of Theorems 2.1.4 and 2.1.5	48
Chapter 3 Popular matchings in graphs of bounded treewidth	56
3.1 Introduction	56
3.2 Leaders	62

3.3	The algorithm	63
3.4	Proof of Lemma 3.2.1	70
Chapter 4 The distance of stable and popular matchings from the Pareto-optimal frontier		79
4.1	Introduction	79
4.2	Further notation and basic facts	87
4.3	Phenomenology	91
4.4	Two easy cases	101
4.5	Proof of Theorem 4.1.3	104
4.6	Proofs of Theorem 4.1.5 and Theorem 4.1.6	116
Chapter 5 Combinatorial Scheduling for Adaptive Machine Learn- ing in Cybersecurity		124
5.1	Introduction	124
5.2	The PLADD game and a new scheduling problem	127
5.3	Integer Programming Formulation	129
5.4	Linear Programming Relaxation	132
5.5	Integrality Gap	137
5.6	Visualization Software	141
Chapter 6 Conclusions		143
Bibliography		146

List of Figures

2.1	The gadget associated to the positive (left) and negative (right) literals, and the corresponding canonical incomplete preference patterns.	25
2.2	A monotone 3-SAT instance $= (x_1 _ x_2 _ x_3) \wedge (: x_1 _ : x_2 _ : x_4) \wedge (x_2 _ x_4 _ x_5)$ and the associated graph $G_s()$ and preference system. Disk and circle labels of vertices provide a bipartition of $G_s()$	27
2.3	A monotone 3-SAT instance $= (x_1 _ x_2 _ x_3) \wedge (: x_1 _ : x_2 _ : x_4) \wedge (x_2 _ x_4 _ x_5)$ and the associated graph $G_g()$ and preference system. Disk and circle labels of vertices provide a bipartition of $G_g()$	28
2.4	A monotone 3-SAT instance $= (x_1 _ x_2 _ x_3) \wedge (: x_1 _ : x_2 _ : x_4) \wedge (x_2 _ x_4 _ x_5)$ and the associated graph $G_m()$ and preference system. Disk and circle labels of vertices provide a bipartition of $G_m()$	29
2.5	A monotone 3-SAT instance $= (x_1 _ x_2 _ x_3) \wedge (: x_1 _ : x_2 _ : x_4) \wedge (x_2 _ x_4 _ x_5)$ and the associated graph $G_h()$ and preference system. Disk and circle labels of vertices provide a bipartition of $G_h()$	30
2.6	An alternating path from s to w when all literals are False in positive (above) and negative (below) clauses.	35
2.7	A matching \mathcal{M} (bold edges) of G_m corresponding to a feasible assignment of $= (x_1 _ x_2 _ x_3) \wedge (: x_1 _ : x_2 _ : x_4) \wedge (x_2 _ x_4 _ x_5)$ from Fig. 2.4: $x_1; x_4; x_5$ are set to True and $x_2; x_3$ to False.	38
2.8	A dominant matching \mathcal{M}^θ (bold edges) corresponding to an example from Fig. 2.4.	54

3.1 An example of popular vs locally-popular. In each graph, let M be the matching given by bold edges. *Top graph.* M is not popular, since there is an M -alternating path in G_M with a $(+; +)$ edge and an unmatched node. But if we take $X = \overline{fv_1; v_2}g$ and $S = \overline{fv_3; v_4}g$, matching $M_{X|S}$ is $(S; X)$ -locally popular. *Bottom graph.* M is popular, hence it is $(X; S)$ -locally popular for every choice of S and X . But if we restrict M to the subgraph G^θ of G induced by $v_1; v_2; v_3; v_4; v_5$, the matching we obtain is not popular in G^θ , since there is a $(+; +)$ edge incident to the now unmatched node v_5 . 57

4.1 On the top: a marriage instance $(G; <)$. In the middle: the only stable matching M of the instance. It contains exactly one coalition cycle C . One can easily check (or resort to Lemma 4.2.2) that the only matching in $PO(M)$ is $M \cup C$. Hence, $\Lambda(M) = k + 2$. On the bottom: the matching realizing $\Delta(M) = 1$ 82

4.2 The Hasse diagram of the poset of matchings of a marriage instance with 4 students and schools. Stable matchings are represented by unfilled circles, while Pareto-optimal matchings correspond to elements of the poset without predecessors. There is exactly one Pareto-optimal matching M_0^θ (unfilled square) that dominates the student-optimal stable matching M_0 . Hence $\Gamma(M_0) = \Gamma(M_0; M_0^\theta)$. There is a second Pareto-optimal matching M_1^θ (filled large circle) that dominates the other stable matching M_1 . Hence $\Gamma(M_1) = \min\{\Gamma(M_1; M_0^\theta); \Gamma(M_1; M_1^\theta)\}g$ and $\Gamma(G; <) = \min\{\Gamma(M_0; M_0^\theta); \Gamma(M_1; M_0^\theta); \Gamma(M_1; M_1^\theta)\}g$. On the other hand, in order to compute $\Delta(G; <)$, we must also consider the other Pareto-optimal matchings (filled squares). 83

4.3 One block i consists of five men $m_i^0; m_i^1; m_i^2; m_i^3; m_i^4$ and four women $w_i^0; w_i^1; w_i^2; w_i^3$ with their preferences labeled on the edges. 95

4.4	One block i consists of three men $m_i^0; m_i^1; m_i^2$ and two women $w_i^0; w_i^1$ with their preferences labeled on the edges.	99
4.5	One gadget.	105
4.6	Corresponding graph for $\varphi = (x_1 _ x_2 _ x_3) \wedge (: x_1 _ : x_4 _ : x_3)$	108
4.7	Corresponding stable matching (bold edges) for $\varphi = (x_1 _ x_2 _ x_3) \wedge (: x_1 _ : x_4 _ : x_3)$, where $x_1; x_3 = T; x_2; x_4 = F$	109
4.8	One block i consists of three men $m_i^0; m_i^1; m_i^2$ and and three women $w_i^0; w_i^1; w_i^2$ with their preferences labeled on the edges.	115
4.9	One gadget.	117
4.10	A monotone 3-SAT instance $\varphi = (x_1 _ x_2 _ x_3) \wedge (: x_1 _ : x_2 _ : x_4) \wedge (x_2 _ x_4 _ x_5)$ with the corresponding preference system $(G(\varphi); <)$	119
5.1	An example of FlipIt game. Defender controls the resource in the beginning, then the attacker makes a move and the defender loses the control. Flags indicate the moves of the players. Red and blue colors show time when the defender and attacker control the resource respectively.	125

5.2 This diagram illustrates an example of PLADD game. Small flags indicate take moves and large flags indicate morph moves. The defender starts the game with a morph move. At the same time the attacker starts the attack that does not succeed immediately but have time-to-success wait time drawn from f_{base} distribution. During this wait time the defender makes two take moves that are wasted and do not have any influence on the system. After that the attacker gains the control of the resource until the defender places a take and regains the control of the system. The attacker knows that he lost control and immediately starts a new attack. This time the wait-to-success time is drawn from $f_{learned}$ distribution, since the attacker already acquired some knowledge about the system from the previous attack. The defender places a take that does not influence the game. The next defender's move, morph, cancels the attack and reset the distribution of wait-to-success time to base. The game resets. The attacker lost acquired knowledge and starts a new attack from f_{base} distribution. . 127

5.3	This diagram illustrates an example of PLADD game. There are three servers. Different jobs on one machine are colored differently for visualization purpose. On all servers j th jobs are the same color. For example, first and second jobs on all servers have red and blue colors respectively. All job lengths are measured in units of time that is needed to process them. The first server has 9 jobs of lengths 1, 4, 1, 1, 4, 1, 1, 4, 1. The second server has 9 jobs of lengths 2, 3, 1, 2, 3, 1, 2, 3, 1. The third server has 6 jobs, each of length 3. There are 9 takes (vertical black lines) and $T = 18$ (vertical red line). The first take t_0 is scheduled at time $t = 0$. The rest of the takes are scheduled at times 1, 2, 5, 6, 8, 12, 13, 15. The idle time is represented by white color between jobs on one server. Jobs after time T shown in dark color to illustrate that they were not processed. The diagram shows the optimal schedule for this instance with the total idle time of 6 units.	129
5.4	Optimal schedule for example 5.4.1. At time $t = 0$ all jobs are scheduled, the next take is placed at time $t = 3$, etc.	136
5.5	Optimal schedule for example 5.4.1. At time $t = 0$ all jobs are scheduled, the next take is placed at time $t = 1$, etc.	136
5.6	Optimal solution for LP for example 5.4.1. Takes are fractional.	136
5.7	For each server blocks of the same color represent the same job. Each server is divided into two parts for visualization of fractional takes. Vertical black lines represent takes, and vertical red line represents the global time T	138
5.8	Any schedule can be represented as a union of the blocks above. Red vertical line represent global time T	139
5.9	The top and bottom schedules are optimal for the problems with $n = 2$ and $n = 4$	141

5.10 A screenshot of Schedule Simulator application 142

List of Tables

4.1 An example of the construction from Algorithm 6 (left) and Algorithm 7
(right) for $n = 3$. The unique stable matching is circled. 92

Acknowledgements

My Ph.D. journey has been a road full of excitement, challenges and discovery. It has been a road of search for continuous improvement and to prove to myself my capabilities and their usefulness to solve problems. First and foremost, I want to thank my advisor, Yuri, who has been my guide in this journey. I always looked up to him for his high standard of quality in everything and unlimited enthusiasm for advancing the research. His passion to research and teaching continuously spreads to other students, including me.

I want to thank my husband Nick. This thesis would not have been possible without him. Not only did he provide me emotional support but also his tremendous talent and genius in software development helped me to develop programs that I used continuously through my research to test ideas, find counterexamples and develop intuition to problems through experiments that are the foundation to my findings. On the third year of my Ph.D. program our daughter Scarlett was born. She is the biggest motivation and inspiration for me to be a better person for tomorrow both in research and personal development.

I would like to thank my mother Tetiana who prioritized my interests and became a mother to my daughter while I was working on my research.

I would like to give a special thank you to my friend Goutam for his continuous support and help.

To my daughter, husband and mother, for everything.

Introduction

This thesis deals with matching and scheduling problems.

The matching problems we will consider go under the generic name of *Matching under preferences*. Together with the input graph (often, but not always, bipartite), we are given, for each node, a strictly ordered list of its neighbors. The goal is to output a matching of the graph that satisfies additional properties. The most classical one, introduced by Gale and Shapley in [18], is *stability*. Recent literature has focused on properties other than stability. The most important for this thesis will be *popularity*, which generalizes stability, in the sense that every stable matching is also popular. While the complexity of many algorithmic problems on the class of stable matchings is well-known, much less was known on the classes of popular matchings. From a discrete optimization point of view, typical questions one would ask are: given a bipartite graph with strict preference lists and nonnegative profits, can one find efficiently a popular matching of maximum profit? What if we are given nonnegative costs, and one wants to find a popular matching of minimum cost? These questions has been asked multiple times in the literature [12, 13, 28, 33, 41]. As our first results, we settle them in Chapter 2, where we show tight hardness results for those and related problems. Interestingly, all the hardness results from Chapter 2 can be seen as a corollary of an easy-to-state problem on *stable matchings*, also introduced in this thesis. These techniques also allows us to give short proofs for other results on popular matchings that have recently appeared in the literature.

When a problem is hard, it is natural to add restrictions on the input, and investi-

gate whether the complexity of the problem changes. For instance, many notoriously hard problems become solvable in polynomial time when the input graph is assumed to have bounded *treewidth*. In Chapter 3, we show this is the case for (extensions of) the max profit / min cost popular matchings problems mentioned above. This also applies in the roommate case, i.e., when the input graph is not bipartite, for which in general graphs even deciding if a popular matching exists is NP-Complete. Results from Chapter 2 and Chapter 3 have appeared in [14, 16, 17].

Stability and popularity are only two among many concepts introduced in the literature. Another important concept is that of *Pareto-optimality*. While popularity generalizes stability, Pareto-optimality and stability are often seen as in opposition to one another. Of course, a matching that is both Pareto-optimal and stable (or popular) is highly sought after, because it satisfies desirable, seemingly contradicting, properties. So a natural question is: can we find a matching that is both stable (resp. popular) and Pareto-optimal, if any? And if there is none, can we find a matching that satisfies one property and is “not too far away” (in some well-defined sense) from the other? These and related questions are investigated in Chapter 4. This is a joint project with Yuri Faenza.

In Chapter 5 we investigate a new combinatorial scheduling problem arising from modeling moving target defense strategies in cyber-security. We analyze a strategic game, PLADD, which is an abstract model for these strategies and provide an IP formulation, show it has unbounded integrality gap, and write a software that solves the problem via enumeration and visualizes the optimal solution. This is a joint project with Ojas Parekh, Cynthia Phillips, Nouri Sakr and Cliff Stein.

We start this thesis by settling notation and introducing many useful concepts in Chapter 1.

Chapter 1

Preliminaries

1.1 Notation

As it is customary, we denote by $\mathbb{N} = \{1; 2; \dots; g\}$ the set of natural numbers, by \mathbb{Q} the set of rational numbers, and by \mathbb{R} the set of real numbers. For $S \subseteq \mathbb{R}$, S_+ denotes the intersection of S with the set of nonnegative reals, while we write $S_0 := S \cap [0; g]$. For $k \in \mathbb{N}$, we let $[k] := \{1; 2; \dots; k\}$. For a set T , $w : T \rightarrow \mathbb{R}$, and $S \subseteq T$, we write $w(S) := \sum_{e \in S} w(e)$. For a graph G , we denote by $V(G)$ its vertices and by $E(G)$ its edges.

Given a graph $G(V; E)$, a *matching* is a collection M of edges from E such that each node of V has degree at most one in M . A node is *covered* by a matching M if its degree in M is exactly one. Given a matching M and $v \in V$, we denote by $V(M)$ the set of nodes covered by M and by $M(v)$ the unique node u such that $uv \in M$ if it exists, and \emptyset otherwise. Given a subgraph H of G and a matching M in G , we denote by M_H the matching induced by M in H , i.e., $M_H := M \cap E(H)$. We will often work with bipartite graphs $G(A \sqcup B; E)$. In this case, given a matching M of G , we write $A(M) := A \cap V(M)$ and similarly $B(M) := B \cap V(M)$.

All paths and cycles in this thesis are simple, hence they are identified by a sequence of nodes $P = v_1; v_2; \dots; v_k$. The *inverse* of P is the path $v_k; v_{k-1}; \dots; v_1$. For a node v of a graph G , we denote by $\Gamma(v)$ the neighborhood of v .

1.2 Sets, posets, lattices

All sets from this thesis are assumed to be finite. Given two sets $A; B$, we let $A \oplus B$ be their symmetric difference, i.e.,

$$A \oplus B := (A \setminus B) \cup (B \setminus A):$$

A *partial order* on a set S is a reflexive, anti-symmetric and transitive binary relation on elements of S . A set together with a partial order, $(S; \preceq)$, is called partially ordered set (or, *poset*), and denoted by $(S; \preceq)$. Given elements $a; b; c$ from a poset, we say that c is a *lower bound* (resp. upper bound) to $a; b$ if $c \preceq a; b$ (resp. $c \succeq a; b$). A *lattice* is a poset such that each pair of elements a and b has a unique *least* upper bound, denoted by $a \vee b$, and a unique *greatest* lower bound, denoted by $a \wedge b$. $a \vee b$ is called a least upper bound if there is no element c such that $a \preceq c, b \preceq c$ and $c \preceq a \vee b$. Similarly, $a \wedge b$ is called a greatest lower bound if there is no element c such that $c \preceq a, c \preceq b$ and $a \wedge b \preceq c$. A *sublattice* of a lattice $(S; \preceq)$ is given by a pair $(S^\emptyset; \preceq^\emptyset)$ with $S^\emptyset \subseteq S$ and \preceq^\emptyset the restriction of \preceq to S^\emptyset , with the property that, for each $a; b \in S$, we have that $a; b \in S^\emptyset$ implies $a \vee b; a \wedge b \in S^\emptyset$.

1.3 Matching under preferences

In the *marriage model* we are given a bipartite graph $G(A \cup B; E)$. We usually refer to A as the set of *men* and to B as the set of *women*. We are also given in input, for each node v of the graph (or *agent*), a strict *ranking* of its neighbors. We denote this ranking by $>_v$, with the convention that v prefers u to z if $u >_v z$. We use numerical labels for preferences, with the convention that between two numbers $i < j$, the smaller number i represents the more preferred choice. Hence, label 1 indicates the top choice. A *marriage instance* is therefore represented by the pair $(G; <)$ with

$\leq f <_v g_{V_2 A} [B]$. The input of the *roommate model* is still a pair $(G; <)$ as above, with the difference that the input graph is not required to be bipartite. We will also call an instance of the roommate model a *preference system*. Many problems in this thesis deals with finding matchings in instances of the marriage or roommate model with special properties. We next present basic concepts and facts that will be useful throughout the thesis, and unless otherwise noted, can be found in [22].

Stability

In the academic community, the marriage model came to prominence through the seminal work by Gale and Shapley [18], who introduced the concept of *stability*.

Consider a marriage instance $(G; <)$ with a matching M . An edge of graph G between nodes u and v will be denoted as $(u; v)$ or uv . We call an edge *blocking* or a *blocking pair* if the endpoints strictly prefer each other to their assignments in M . If M does not have blocking pairs, then M is *stable*. In [18] the authors showed that in the marriage model a stable matching always exists and can be found in linear time by the following algorithm:

Algorithm 1: Gale-Shapley
<pre> start from all men and women being unmatched while there is a man m (a node from set A) that is unmatched and has a neighboring woman (a node from set B) that he did not propose to do $w :=$ first woman in m's list to whom m has not proposed yet: if w unmatched then assign m and w to be matched to each other end else if w prefers m to its match m^l then assign m and w to be matched and m^l be unmatched end else w rejects m end end end end output formed pairs </pre>

We now know [46] that the problem was formulated and solved in practice in the US labor market for medical interns and residents well before the paper by Gale and Shapley. Nowadays, matching markets are used to model many different frameworks, from online dating [25], to house and school choice [2], to the assignment of users to servers in large networks [40].

We now present some basic facts about stable matchings. Some will apply to both roommate and marriage instances, while others only to the marriage case. We denote by $S(G; \prec)$ the set of stable matchings of $(G; \prec)$. We omit the dependency on $(G; \prec)$ when clear from the context. The first useful fact is that all stable matchings match the same vertices. In particular, they all have the same cardinality.

Theorem 1.3.1. *Let $(G; \prec)$ be a preference system and $M, M^0 \in S(G; \prec)$. Then $V(M) = V(M^0)$.*

Consider the following partial order over set of matchings:

$$M \preceq M^0 \iff M(m) \succeq_m M^0(m) \text{ for all } m \in A;$$

When $(G; \prec)$ is a marriage instance, $(S; \preceq)$ is a lattice. In particular, if $M, M^0 \in S$, then $M^A; M^O \in S$, where for each man $m \in A$, $M^A(m)$ (resp. M^O) is set to be the most (resp. least) favourite partner of m between $M(m)$ and $M^0(m)$. This implies that, among all stable matchings, there is one that all men (weakly) prefer to any other stable matching. It is called *man-optimal* matching. Similarly, there is a stable matching that all women prefer to any other stable matching: we call it *woman-optimal*.

Theorem 1.3.2. *Let $(G; \prec)$ be a marriage instance. Gale-Shapley algorithm, when men propose, results in the man-optimal stable matching. Similarly, Gale-Shapley algorithm, when women propose, results in the woman-optimal stable matching.*

The lattice structure can be exploited for developing fast algorithms for a number of problems related to stable matchings. In this thesis, we will be mostly interested in its consequences for edge-weighted stable matching problems.

Problem 1.3.1. *[weighted stable matching problem]* **Given:** A marriage instance $(G; <)$, and a function $w : E \rightarrow \mathbb{R}$. **Find:** a stable matching M that maximizes $w(M)$.

Theorem 1.3.3. *The weighted stable matching problem can be solved in polynomial time. Moreover, the convex hull $P(G; <)$ of the characteristic vectors of stable matchings in a marriage instance $(G; <)$ can be described via a polynomial number of constraints.*

Interestingly, the lattice structure is preserved if we restrict to the subset of stable matching containing and/or avoiding certain subsets of edges.

Lemma 1.3.1. *Let $(G; <)$ be an instance of the marriage problem, and let $E^{in}, E^{out} \subseteq E(G)$. Define*

$$S(E^{in}; E^{out}) := \{S \subseteq E : E^{in} \subseteq S \text{ and } E^{out} \cap S = \emptyset\}$$

Then $(S(E^{in}; E^{out}); \subseteq)$ is a sublattice of S . In particular, there exists a man-optimal stable matching in $(S(E^{in}; E^{out}); \subseteq)$.

Proof. Immediately from the fact that if $M, M' \in S(E^{in}; E^{out})$, then $M \wedge M', M \vee M' \in S(E^{in}; E^{out})$. □

The polysize description of $P(G; <)$ from Theorem 1.3.3 immediately implies the following.

Corollary 1.3.1. *Given a marriage instance $(G; <)$, a function $w : E \rightarrow \mathbb{R}$, sets $E^{in}, E^{out} \subseteq E(G)$, a matching of $(S(E^{in}, E^{out}); <)$ of maximum weight can be found in polynomial time.*

Popularity

In 1975 in [19], Gärdenfors introduced the notion of popularity and showed that it is a generalization of stability. Recall that every stable matching has the same size, see Theorem 1.3.1, and there could be just two agents who are blocking a matching. This is a very restricting property. In popular matchings, on the other hand, the input from all agents on a given matching is taken into consideration.

Consider a preference system $(G(V; E); <)$. We say a vertex $v \in V$ *prefers* matching M to matching M^θ if either v is matched in M and unmatched in M^θ , or v is matched in $M; M^\theta$ and strictly prefers its partner in M to its partner in M^θ . Let $\pi(M; M^\theta)$ be the number of vertices that prefer M to M^θ . A matching M is *more popular* than M^θ , if $\pi(M; M^\theta) > \pi(M^\theta; M)$. A matching M is *popular* if there is no matching M^θ that is more popular than M , meaning $\pi(M; M^\theta) \geq \pi(M^\theta; M)$ for all matchings M^θ in $(G; <)$. The following fact has been shown in [27].

Lemma 1.3.2. *Let S (resp. M) be a stable (resp. popular) matching of a preference system $(G; <)$. Then $V(S) = V(M)$. In particular, stable matchings are popular matchings of minimum size.*

Relaxing stability to popularity allows therefore to increase the size of matchings, while still preserving a notion of global stability. However, given a matching, checking if it is popular is not immediate, since we would have to compare it against all other matchings of the graph. However, this can be still performed in polynomial time, see [7, 27].

Theorem 1.3.4. *Given a preference system $(G; <)$ and a matching M of G , whether M is popular can be decided in polynomial time.*

One of the algorithms for efficiently recognizing whether a matching is popular is based on the following combinatorial characterization of popular matchings. Fix a matching M of $(G; <)$. An M -alternating path (resp. cycle) in G is a path (resp. cycle) whose edges are alternatively in M and in $E(G) \setminus M$. We associate labels to edges from $E \setminus M$ as follows: an edge uv is $(-;-)$ if both u and v prefer their respective partners in M to each other; $uv = (+;+)$ if u and v prefer each other to their partners in M ; $uv = (+;-)$ if u prefers v to its partner in M and v prefers its partner in M to u . To denote labels of edge uv , sometimes we write uv is a $(+;+)$ edge, and similarly for the other cases. If we want to stress that u prefers his current partner to the other endpoint of the edge, we write that e.g. uv is a $(-;+)$ edge. The graph G_M is defined as the subgraph of G obtained by deleting edges that are labeled $(-;-)$. Observe that M is also a matching of G_M , hence definitions of M -alternating path and cycles apply in G_M as well. These definitions can be used to obtain a characterization of popular matchings in terms of forbidden substructures of G_M [27]. For a path (resp. cycle) P and an edge e , we say that P contains e if $e = uv$ for two nodes $u;v$ that are consecutive in P (not necessarily in this order).

Theorem 1.3.5. *Let $(G; <)$ be a preference system, and M be a matching of G . Then M is popular if and only if G_M does not contain any of the following:*

1. *an M -alternating cycle that contains a $(+;+)$ edge.*
2. *an M -alternating path that contains two distinct $(+;+)$ edges.*
3. *an M -alternating path starting from an M -exposed node that contains a $(+;+)$ edge.*

If the matching M under consideration is clear, we omit mentioning that labels are wrt M , and paths, cycles etc. occur in the graph G_M : we simply say e.g. that uv

is a $(+; +)$ edge and P is an M -alternating path with a $(+; +)$ edge. In particular, unless stated otherwise, when talking about M -alternating paths or cycles, we always mean in the graph G_M (hence, after $(-; -)$ edges have been removed).

We say a matching M *defeats* a matching M^θ (and M^θ is defeated by M) if either of these two conditions holds:

- (i) M is more popular than M^θ , i.e., $(M; M^\theta) > (M^\theta; M)$;
- (ii) $(M; M^\theta) = (M^\theta; M)$ and $jMj > jM^\theta j$.

A *dominant* matching is defined to be a matching that is never defeated. A dominant matching is by definition a popular matching of maximum size.

A maximum size popular matching can be twice of the size of a stable matching in the same instance. In [32], Kavitha shows how to find a maximum size popular matching efficiently in marriage instances using a combination of the Gale-Shapley algorithm and promotion of nodes that were rejected by all adjacent nodes. In fact, her algorithm finds a dominant matching.

The following characterization of dominant matchings appeared in [13], where it is also shown that a dominant matching always exists, and all dominant matchings are of the same size (a proof of this follows by Lemma 1.3.3 below).

Theorem 1.3.6. *Let M be a popular matching in a preference system $(G; <)$. M is dominant if and only if there is no M -augmenting path in G_M .*

The following fact is attributed in [12] to [24]. We give a proof here for completeness.

Lemma 1.3.3. *Let $(G; <)$ be a preference system, and M (resp. S , D) be a popular matching (resp. a popular matching of minimum size, a popular matching of maximum size) in $(G; <)$. Then $V(S) = V(M) = V(D)$.*

Proof. The fact that stable matchings are popular matchings of minimum size and that $V(S) = V(M)$ whenever S is stable and M is popular follows from Lemma 1.3.2. Hence, for any popular matching S^θ of minimum size, we have $V(S^\theta) = V(S)$ and $jS^\theta j = jSj$, which implies $V(S^\theta) = V(S) = V(M)$ for any popular matching M , settling the first part. Now suppose that M again is popular, D is a dominant matching, and there exists $v \in V(M) \cap V(D)$. Then, $G(M \oplus D)$ contains an alternating path P starting at v with an edge of M . Since both matchings are popular, the number of nodes that prefer M to D in P are exactly the same number of nodes that prefer D to M in P . In particular, P has an even number of nodes and is therefore D -augmenting in G . Now label edges of $P \cap D$ with $+$; as discussed at the beginning of the section. Note that the number of $(+;+)$ edges and $(-;-)$ edges in P coincide, otherwise either D or M is not popular. We consider three cases: if there is no $(+;+)$ edge, hence no $(-;-)$, the path is D -augmenting in G_D , contradicting the fact that D is dominant. Now suppose there is a $(+;+)$ edge e . Since the number of $(+;+)$ and $(-;-)$ edges in P coincide, one of the following two facts happen: either we can assume wlog that e is traversed in P before any $(-;-)$ edge, or there is a subpath of P that contains two $(+;+)$ edges and no $(-;-)$ edge. In both circumstances, we contradict Theorem 1.3.5: condition (ii) in the second case, and condition (iii) in the first. Hence $V(M) = V(D)$. If M^θ is a popular matching of maximum size, we conclude $V(M^\theta) = V(D)$, as required. \square

Similarly to what was done for stable matching, in the bipartite case we can define (and solve efficiently) the problem of computing a dominant matching of maximum weight.

Problem 1.3.2. [*weighted dominant matching problem*] **Given:** A marriage instance $(G; <)$, and a function $w : E \rightarrow \mathbb{R}$. **Find:** a dominant matching M that maximizes $w(M)$.

Theorem 1.3.7. [13] *Problem 1.3.2 can be solved in polynomial time for marriage instances. Moreover, the convex hull $D(G; <)$ of the characteristic vectors of dominant matchings in a marriage instance $(G; <)$ has a linear extension¹ that can be described via a polynomial number of constraints.*

Given $(G; <)$, we often represent $<$ via a *preference pattern*, i.e., a function L that maps ordered pairs $(u; v)$ with $uv \in E$ to the position of v in $<_u$. If $L(u; v) = 1$, we say that v is the *favorite partner* of u . Sometimes it makes sense to specify a partially described preference pattern for a subgraph of the input graph. More formally, an *incomplete preference pattern* I for a bipartite graph G is a map $I : E(G) \rightarrow \mathbb{N} \cup \{0\}$, such that $I(u; v) = I(u; x)$ for some $u \in U, x \in V, v \in V$ implies $I(u; v) = I(u; x) = 0$, and $I(u; v) \geq \mathbb{N}$ implies $I(u; v) = |\Gamma(u)|$. We say that a preference pattern L for a supergraph G^\flat of G *agrees* with an incomplete preference pattern I if, for all $uv \in E(G)$ with $I(u; v) \neq 0$, we have: $L(u; v) = I(u; v)$ whenever $I(u; v) \geq \mathbb{N}$ and $L(u; v) = |\Gamma_{G^\flat}(u)|$ whenever $I(u; v) = 1$ (here $\Gamma_{G^\flat}(u)$ is the neighborhood of u in G^\flat).

On possible applications for popular matchings

Stability has proven to be a powerful concept with a rich mathematical structure used in numerous real-world applications in matching markets, ranging from house allocation, to assignment of students to schools, users to servers, medical students to hospitals, refugees to cities. In [41] authors give an overview of the topic and mention many more applications such as kidney exchanges with incompatible patient-donor pairs [47, 48, 5], auction mechanisms for sponsored search [6], assigning children to daycare [36], placement of graduating rabbis [8], supply-chain networks [43], online dating in the US [26], and online matrimony in India [31], etc. However, it is known

¹A linear extension of a polytope P is a polytope Q that linearly projects to P .

that all stable matchings have the same size, see Theorem 1.3.1, and the set of matched vertices is the same. In certain cases it is strongly preferred to obtain a matching of larger size (e.g. more people are matched or assigned) with the sacrifice of having some blocking pairs but still with the global notion of stability.

Popular matchings are the potential substitute for more restrictive stable matching in many of the above applications. For example, one potential practical use of popular matching is online dating. In [26] authors studied data from an online dating service and made predictions using the Gale-Shapley algorithm. The resulted matchings are similar to matchings formed by the dating site. The preferences are estimated from user's information and interactions on the website. In a new study [45] sociologists found that the most popular way heterosexual couples meet is online. If more people become users of online dating websites every year, then it is crucial to develop algorithms that would result in matchings leading to happy married life for as many people as possible. Since stable matchings are proven to underlie matchings on one of the major dating websites [26], then potential application of popular matchings could lead even to more optimistic results, where the number of people matched is larger.

Pareto-optimality

Popular matching provide an opportunity to overcome drawbacks of stability. There is another class of matchings, *Pareto-optimal*, that also relax stability by ignoring the preferences of one side. They are proven to be useful in numerous applications, such as assignment students to housings. They are also well-studied in economics. Consider a marriage instance $(G(A [B); <)$. A matching M *Pareto beats* a matching M^0 of $(G; <)$ if there exists a man which prefers M to M^0 and no man prefers M^0 to M . A matching M is *Pareto-optimal* if there is no matching M^0 which Pareto beats M . Popularity provides global stability by ensuring overall satisfaction, while

Pareto-optimally maximizes the satisfaction of one side.

A matching M is maximal if there is no edge $mw \in E(G)$ such that m and w are both unmatched in M . M is *trade-in-free* if there is no edge mw , $m \in A$; $w \in B$ such that woman w is unmatched and man m prefers w to its partner $M(m)$. An M -alternating cycle C is called a *coalition cycle* of M , if for each edge $mw \in E(C) \cap M$, $m \in A$; $w \in B$, m prefers w to $M(m)$. If M does not contain coalition cycles, then it is called *coalition-free*.

The following Lemma provides a characterization of Pareto-optimal matchings [4].

Lemma 1.3.4. *Let $(G; <)$ be a preference system, and M be a matching of G . Then M is Pareto-optimal if and only if M is maximal, trade-in-free, and coalition-free.*

Theorem 1.3.8 provides a characterization of popular Pareto-optimal matchings.

Theorem 1.3.8. *Let $(G; <)$ be a preference system, and M be a popular matching of G . Then M is Pareto-optimal if and only if M is coalition-free.*

Proof. Let M be popular and Pareto-optimal. From Lemma 1.3.4 we conclude that M is coalition-free.

Now let M be popular and coalition-free. We prove that M is maximal and trade-in-free. Assume by contradiction that M is not maximal, and there is an edge mw with both m and w being unmatched. In this case mw is an M -alternating path in G_M from an unmatched node m to $mw = (+; +)$. This contradicts popularity (see Lemma 1.3.5). Assume by contradiction that M is not trade-in-free, then there is an unmatched woman w and m such that m prefers w to $m_M(w)$. Similarly, mw is an M -alternating path in G_M from an unmatched node w to $mw = (+; +)$. This is a contradiction to popularity, see again Lemma 1.3.5. □

In [49] the following matching mechanism *serial dictatorship* was introduced.

<p>Algorithm 2: Serial dictatorship</p> <pre> start from all men and women being unmatched mark all women "unmatched", and let A be the set of men while A is not empty do choose any man m from A $w :=$ top choice for man m among women with mark "unmatched" assign w to m remove m from A change w's mark from "unmatched" to "matched" end output formed pairs </pre>
--

In [1] it was shown the following characterization of Pareto-optimality.

Theorem 1.3.9. *Let $(G; <)$ be a preference system, and M be a matching of G . Then M is Pareto-optimal if and only if M is obtained from the serial dictatorship mechanism.*

1.4 SAT and hardness of approximation

A boolean formula in CNF (*conjunctive normal form*) is a conjunction of clauses such that each clause is a disjunction of literals. Each literal can be positive or affirmative (denoted by x_i), or negative or negated (denoted by \bar{x}_i). Given a CNF boolean formula, we define an *assignment* to be a mapping from literals to $\{0, 1\}$, and as usual we say that a literal that maps to 0 is set to *false*, while a literal that maps to 1 is set to *true*. We say an assignment is *feasible* if in each clause, there is at least one literal set to true. We say that an assignment is *consistent* if a literal corresponding to x_i is set to true if and only if all literals corresponding to x_i are set to true and all literals corresponding to \bar{x}_i are set to false. Note that consistent assignments are in bijection with maps from the set of *variables* to $\{0, 1\}$ (assuming wlog that each variable appears at least once). Hence, we often refer to a consistent assignment of literals as an *assignment of variables*. An assignment of literals that is both consistent and feasible is called *satisfying*.

A *Boolean satisfiability problem* is the following decision problem.

Problem 1.4.1. [SAT] *Given:* A boolean formula ϕ in CNF. *Decide:* If there exists a satisfying assignment to ϕ .

A SAT instance is *satisfiable* if it has a satisfying assignment. For $\alpha \in (0;1]$, we say that a SAT instance ϕ is α -satisfiable if there exists an assignment of variables that satisfies at least a α fraction of the clauses of ϕ . The following special case of SAT is well-studied.

Problem 1.4.2. [3-SAT] *Given:* A boolean formula ϕ in CNF, where each clause has at most 3 literals. *Decide:* If there exists a satisfying assignment to ϕ .

Indeed, the following result is known to hold.

Theorem 1.4.1. [23] *Let $\epsilon > 0$ be a constant. Unless $P=NP$, there is no polynomial-time algorithm that, given an instance ϕ of 3-SAT, accepts if ϕ is satisfiable, rejects if ϕ is not $(7/8 + \epsilon)$ -satisfiable, and is free to accept or reject if none of these cases apply.*

Monotone 3-SAT (called NM-E3SAT in [21]) is the version of 3-SAT where each clause consists of only positive or only negative literals.

Problem 1.4.3. [Monotone 3-SAT] *Given:* A boolean formula ϕ in CNF where each clause contains at least 2 and at most 3 variables, either all negated or all positive. *Decide:* If there exists a satisfying assignment to ϕ .

Monotone 3-SAT is also known to be NP-hard, since as observed in [11], a reduction from an instance of 3-SAT can be achieved using the following mapping: replace clause $C = (x_1 \vee x_2 \vee x_3)$ by $(x_1 \vee x_c) \wedge (x_c \vee x_2 \vee x_3)$ and replace clause $C = (x_1 \vee x_2 \vee x_3)$ by $(x_1 \vee x_2 \vee x_c) \wedge (x_c \vee x_3)$, where for each clause C , x_c is a new variable, and similarly for the clauses with two variables only. The other clauses remain unchanged. In [21], Theorem 1.4.1 is extended to the maximization version of Monotone 3-SAT.

Theorem 1.4.2. Let $\epsilon > 0$. Unless $P=NP$, there is no polynomial-time algorithm that, given an instance of Monotone 3-SAT, accepts if it is satisfiable, rejects if it is not $(7/8 + \epsilon)$ -satisfiable, and is free to accept or reject if none of these cases apply.

1.5 Treewidth: definitions and some facts

The treewidth of a graph G is a parameter that measures, roughly speaking, how far G is from being a tree. For many optimization problems, one can formulate dynamic programming algorithms that depend polynomially on the size of the graphs, and exponentially on the treewidth. Hence, when the treewidth is bounded, many NP-Hard problems become tractable, or admit better approximation algorithms. Examples of such problems include the Maximum Independent Set, Minimum Vertex Cover, and the Graph Coloring Problem.

We now formally introduce those concepts. Let $G = (V; E)$ be a graph. A *tree decomposition* of G is a pair $(T; B)$ where T is a tree and $B = \{B(i) : i \in V(T)\}$ is a family of subsets of $V(G)$, called *bags*, one for each vertex of T , satisfying the following:

1. For each $(u; v) \in E$, there is at least one bag $B \in B$ such that $(u; v) \subseteq B$.
2. If $i \neq j \neq k \in V(T)$ are such that k is on the unique path from i to j in T then $B(i) \cap B(j) \subseteq B(k)$.

Note that the second property implies that, for any vertex $u \in V(G)$, the bags which contain u form a subtree of T . We will sometimes abuse notation and denote by B both a vertex of $V(T)$ and the bag corresponding to it.

The *width* of a tree decomposition $(T; B)$ is $\max_{B \in B} |B| - 1$. The *treewidth* of G is the minimum integer t such that there is a tree decomposition of G of width t .

Let $G = (V; E)$ be a graph and $(T; B)$ be a tree decomposition of G of width t . Wlog we can assume that, for each pair of bags $B; B'$ adjacent in T , $B \cap B'$ is a vertex

separator of G . Form a directed rooted tree by picking an arbitrary vertex as the root and orienting the remaining edges towards the root, and call this a *directed tree decomposition*. In the directed tree, each node X other than the root node has exactly one *successor* $S(B)$, i.e., there exists exactly one node $S(B)$ such that $(B; S(B))$ is an arc of the directed tree decomposition. If B is the root, we set $S(B) = \cdot$. We also say that B is a *predecessor* of $S(B)$ and notice that a bag may have multiple predecessors, and has none if and only if it is a leaf of T .

Dichotomic tree decomposition. Let G be a graph and $(T; B)$ a directed tree decomposition of G of width c . We can transform the directed tree decomposition of G into a directed tree decomposition of the same graph and width where every node has indegree at most 2. We call such a tree decomposition *dichotomic*. Indeed, suppose edges $(B_1; B); (B_2; B); \dots; (B_t; B)$ with $t \geq 3$ are part of the directed tree decomposition. Then we can create a copy \bar{B} of B , add edge $(\bar{B}; B)$, and split the edges entering B as follows: $(B_1; \bar{B}); \dots; (B_{dt=2c}; \bar{B})$ and $(B_{dt=2c+1}; B); \dots; (B_t; B)$. The indegree of X and \bar{X} is at most $bt=2c+1 < t$. Notice that the digraph we obtain is still a tree decomposition of G , since each edge of G is still covered by a bag, and the bags which contain any vertex $v \in G$ still form a continuous subtree.

If we repeat the above operation for a non-dichotomic tree once for each original node of indegree at least 3, the maximum indegree over all nodes of the tree goes from $t > 2$ to $bt=2c+1$, while the number of nodes is at most doubled. Hence, we can iterate the operation so as to obtain a tree decomposition with at most a quadratic number of vertices, all of which have indegree at most 2.

From here on, we assume without loss of generality that our directed tree decomposition is dichotomic. Let T^θ be a subtree of T , and $V(T^\theta)$ the set of nodes contained in at least a bag of T^θ . We say that T^θ is a *closed subtree* of T if $B \in V(T^\theta)$ and B^θ is a predecessor of B , then B^θ belongs to T^θ . By connectivity, for each node B of a closed subtree T^θ , $S(B)$ also belongs to T^θ , with the exception of at most one node

that we call the *head* of T^θ and denoted by $H(T^\theta)$. If $T^\theta \notin T$, the successor of $H(T^\theta)$ exists and it is also called the *successor of T^θ* and denoted by $S(T^\theta)$. Note that each bag B is the head of exactly one closed subtree of T , that we denote by T_B .

Remark 1.5.1. *Let $(T; B)$ be a dichotomic directed tree decomposition. Then the following holds: let T^θ be a closed subtree of T and B be the head of T^θ . The removal of B partitions $T^\theta \setminus B$ in at most 2 closed subtrees. The successor of each of those subtrees is B .*

Popular matching problems with edge weights

2.1 Introduction

Although there are efficient algorithms that compute popular matchings of minimum and maximum size in marriage instances, no approach has been developed to deal with more general classes of popular matching problems. This is probably due to the lack of understanding of the structure of popular matching when compared, for instance, to stable matchings. Popular matchings do not form a lattice wrt \subseteq , and it is not hard to come up with examples where they exhibit a bizarre behavior, e.g., the gap between popular matching of maximum and minimum size is arbitrarily large, while no popular matchings of intermediate size exist.

While understanding the structure of popular matchings, if any, seems to be still beyond reach, one can investigate specific algorithmic questions. A typical way to generalize max and min cardinality problems is to assign weights to edges, and then ask for a feasible matching that maximizes the weight function. Recall that all weighted stable matching problems can be solved efficiently for stable matchings using the lattice structure, see Theorem 1.3.3. Those weights can be used to model a number of constraints, such as forbidding or forcing edges, or finding an *egalitarian* or *minimum regret* matching (see, e.g., [41] for a list of those problems).

We can therefore pose the same question for the set of popular matchings: can we give an efficient algorithm that finds a popular matching of maximum weight? This problem and special cases have been posed as open questions in many papers, includ-

ing [12, 13, 41, 28, 33]. However, very little was known about this. The only notable exception arises from the problem where we ask for a popular matching containing a given edge, or decide that no such matching exists. This can be interpreted as the problem of maximizing a function over the edges with one strictly positive entry, and all other entries being equal to 0. A polynomial algorithm for this problem has been given by Cseh and Kavitha [13]. This results also easily implies that one can find a popular matching without a given edge (resp. with a given node) in polynomial time, or conclude that no such matching exists. To the best of our knowledge, no other special case had been solved.

In this section, we introduce a technique to show hardness results for popular matching problems. This technique is based on a reduction from 3-SAT to a new problem on *stable* matchings, also introduced in here. We show that this idea is very fertile, as it allows us to give very fine-grained answers to the complexity questions mentioned above. These include both NP-hardness results and sharp inapproximability results. We also show how other hardness results that have recently appeared in the literature can be proved with these techniques. This leads to much simpler proofs of those facts. Hence, the “unifying” approach that could not be obtained for fast algorithms, can instead be obtained for hardness results.

In Chapter 4, we will push the connection between 3-SAT and matching under preferences further. In particular, we will show that more hardness results related to stable, popular, and Pareto-optimal matchings can be deduced using this connection to 3-SAT and classical hardness of approximability results the latter.

The rest of the chapter is organized as follows. In Section 2.1, we formally define the main problems under consideration and state our results. The main technical contribution is given in Section 2.2, where a reduction from Monotone 3-SAT to the existence of certain stable matchings is given. In Section 2.3, we use the reduction from Section 2.2 to investigate the complexity of popular matchings with edge and

node weights. We conclude with the analysis of the complexity of finding a popular matching that is not stable or dominant in bipartite graphs, and finding a popular matching in general graphs in Section 2.4.

Problems and results

The main building block of all our hardness results is the following problem.

Problem 2.1.1. *Given:* A marriage instance $(G; <)$. *Decide:* If there exists a stable matching M such that G_M has no augmenting path.

For the problem above, we show the following.

Theorem 2.1.1. *Problem 2.1.1 is NP-Complete.*

We then consider several special cases of the problem of finding a popular matching of maximum weight.

Problem 2.1.2 (Popular matching with forbidden and forced elements problem – pmffe). *Given:* A marriage instance $(G; <)$, $U^{in}; U^{out} \subseteq V(G)$, $F^{in}; F^{out} \subseteq E(G)$, such that $U^{in} \cap U^{out} = \emptyset$, $F^{in} \cap F^{out} = \emptyset$, and if $uv \in F^{in} \cap F^{out}$ then $u, v \notin U^{in} \cup U^{out}$. *Decide:* If there exists a popular matching of $(G; <)$ that contains all vertices of U^{in} , all edges of F^{in} , no vertices of U^{out} and no edges of F^{out} , or conclude that such matching does not exist.

Problem 2.1.3 (Maximum Popular matching with positive weights – mwp). *Given:* A marriage instance $(G; <)$, weights $w : E(G) \rightarrow \mathbb{R}_+$. *Find:* A popular matching M of $(G; <)$ that maximizes $w(M)$.

Problem 2.1.4 (Minimum weight popular matching with positive weights – miwp). *Given:* A marriage instance $(G; <)$, weights $w : E(G) \rightarrow \mathbb{R}_+$. *Find:* A popular matching M of $(G; <)$ that minimizes $w(M)$.

Let M be a popular matching that is neither dominant nor stable. If the size of M is strictly less than the size of a dominant matching and is strictly larger than the size of a stable matching, then we call M a *middle* matching; if it equals to the size of a dominant matching, we call it *hidden*.

Problem 2.1.5 (Hidden matching – hm). **Given:** A marriage instance $(G; <)$. **Decide:** If there exists a hidden matching M of $(G; <)$.

For those problems, we show the following tight hardness, approximability, and hardness of approximation results.

Theorem 2.1.2. If $U^{out} = F^{in} = F^{out} = ;$ or $U^{in} = F^{in} = F^{out} = ;$, or $jU^{in} [U^{out} [F^{in} [F^{out}j \ 1$, then *pmffe* can be solved in polynomial time. Otherwise, for each other set of values for $jU^{in}j, jU^{out}j, jF^{in}j, jF^{out}j$, *pmffe* is NP-hard.

Theorem 2.1.3. *miwp* and *mwp* are NP-Hard. Unless $P=NP$, *miwp* cannot be approximated in polynomial time up to any factor, and *mwp* cannot be approximated in polynomial time better than a factor $\frac{1}{2}$. On the other hand, there is a polynomial-time algorithm that computes a $\frac{1}{2}$ -approximation to *mwp*.

We remark that the hardness of *mwp* has been independently shown in [34], but to the best of our understanding, results from [34] does not imply any hardness of approximation.

Theorem 2.1.4. *hm* is NP-complete.

Hardness results follow from a similar reduction as the one introduced to prove Theorem 2.1.1 and from the combinatorial characterization of popular matchings given by Theorem 1.3.5. The approximation algorithm builds on a decomposition result from [13] that partitions a popular matching into a stable and a dominant matching.

We then move to the following problems, whose hardness had been recently established in the literature [20, 17, 35].

Problem 2.1.6 (Middle matching – mm). **Given:** A marriage instance system $(G; <)$. **Find:** A middle matching M of $(G; <)$.

Problem 2.1.7 (Popular matching in general graphs – pmgg). **Given:** A preference system $(G; <)$. **Find:** A popular matching of $(G; <)$, or conclude that such matching does not exist.

These hardness results can be shown by simple modifications of our construction, showing that a uniform approach to all those problems is possible. Those proofs are, we believe, simpler than the original.

Theorem 2.1.5. *pmgg and mm are NP-hard.*

2.2 The reduction

From a Monotone 3-SAT formula to $G(\cdot)$

Consider a Monotone 3-SAT instance \mathcal{I} . We denote the set of variables by $X = \{x_1, \dots, x_n\}$, the set of clauses by $\mathcal{C} = \{c_1, \dots, c_m\}$, and the set of literals by $L = \{x_1, \dots, \bar{x}_n\}$, with $p = 3m$. Literals of the form x_i are called *positive* (occurrences of variable i), while those of the form \bar{x}_i are called *negative* (occurrences of variable i). We refer therefore to clauses of a Max Monotone 3-SAT instance as *positive clauses* and *negative clauses* accordingly. From now on, we assume without loss of generality that \mathcal{I} does not consist of only positive (resp. negative) clauses and that no variable only appears as a positive (resp. negative) literal.

Define the *positive gadget* to be the graph H constructed as in Fig. 2.1, left. Node a is called the *apex* of H . Node c is called the *consistency enforcer* of H . Node g is




Figure 2.1: The gadget associated to the positive (left) and negative (right) literals, and the corresponding canonical incomplete preference patterns.

called the *gateway* of H .

Define the *negative gadget* to be the graph N constructed as in Fig. 2.1, right. Similarly to the positive gadget, we call node a (resp. c , g) the *apex* (resp. *consistency enforcer*, *gateway*) of N . Moreover, we call edge eh the *evicted* edge of N . Fig. 2.1, right and right, also define incomplete preference patterns I_N , which we again call *canonical*.

To each positive (resp. negative) literal ℓ of \mathcal{C} , we associate one positive (resp. negative) gadget $H(\ell)$ (resp. $N(\ell)$). We will write $G(\ell) = (V(\ell); E(\ell))$ to denote the gadget associated to a literal ℓ , whether it is positive or negative. Nodes of $G(\ell)$ will be denoted by $a(\ell)$, $b(\ell)$, etc.

Fix an arbitrary order of clauses of \mathcal{C} and of literals in each clause. We define four graphs $G_s := G_s(\mathcal{C})$, $G_g := G_g(\mathcal{C})$, $G_m := G_m(\mathcal{C})$ and $G_h := G_h(\mathcal{C})$ as follows. Those four graphs will be used in different hardness reductions. They are all variations of a graph defined as follows. First, construct disjoint graphs $G(\ell)$ for each literal ℓ , and add nodes u and v . Now fix a clause c , and let $\ell_1; \ell_2; \ell_3$ be the literals from the clause, in this order (with ℓ_3 possibly not present). Identify the apex of $G(\ell_1)$ with u ; the apex of $G(\ell_i)$, $i = 2, 3$ with the gateway of $G(\ell_{i-1})$; and add edge $g(\ell_k)v$, where ℓ_k is the last literal of the clause. Repeat the operation for all clauses (note that vertices u and v are unique). Let $\bar{G} := G(\bar{V}; \bar{E})$ be the graph obtained. Let

$$G_s = G(\underbrace{\bar{V} [fs; tg;]}_{V(G_s)}; \underbrace{\bar{E} [E [fst; tug)]}_{E(G_s)}), \text{ where}$$

$E := [_{i=1, \dots, n} f c(\cdot) c(\cdot^{\theta}) : \cdot$ (resp. \cdot^{θ}) is a positive (resp. negative) occurrence of variable ig

and $s; t \not\cong \bar{V}$. Define graphs $G_g; G_m; G_h$ as follows:

$$V(G_g) := V(G_s) [f p_1; p_2; p_3; p_4; w; x; y g \text{ and } E(G_g) := E(G_s) [E_1^g, \text{ where}$$

$$E_1^g := f p_1 s; p_2 s; p_3 p_4; p_3 p_1; p_4 p_2; v w; w x; x y; y w g;$$

$$V(G_m) := V(G_s) [f w; x; y g \text{ and } E(G_m) := E(G_s) [E_1^m, \text{ where:}$$

$$E_1^m := f v w; w x; x y g;$$

$$V(G_h) := V(G_s) [f t_1; u_1; w; x; y g \text{ and } E(G_h) := E(G_s) [E_1^h, \text{ where:}$$

$$E_1^h := f t t_1; u u_1; v w; w x; x y; y v g;$$

See Fig. 2.3, 2.4, 2.5 for an example of graphs $G_g(\cdot)$, $G_m(\cdot)$, and $G_h(\cdot)$ respectively.

Lemma 2.2.1. $G_s(\cdot)$, $G_m(\cdot)$ and $G_h(\cdot)$ are bipartite.

Proof. First, we will show that $G_m(\cdot)$ is bipartite. We assign each vertex from $V := V(G_m(\cdot))$ to exactly one of two sets A or B , and prove that if $u; v \in A$ (or $u; v \in B$) then $uv \not\in E := E(G_m(\cdot))$. The proof can be followed in Fig. 2.4, where nodes are colored according to the set of the bipartition they belong to. Assign $s; u; w; y$ to A and $t; v; x$ to B . For any \cdot corresponding to a positive literal, assign $a(\cdot); d(\cdot); c(\cdot); g(\cdot)$ to A and $b(\cdot); e(\cdot); f(\cdot)$ to B . For any \cdot corresponding to a negative literal, assign $a(\cdot); d(\cdot); e(\cdot); g(\cdot)$ to A and $b(\cdot); c(\cdot); f(\cdot); h(\cdot)$ to B . Note that this assignment is consistent with the fact that $a(\cdot) \quad g(\cdot^{\theta})$ for some literals \cdot and

φ , and that some $a(\cdot)$ are identified with u (they are all assigned to A). One easily checks that edges in all gadgets connect nodes in different sets of the bipartition. Now consider $c(\cdot)c(\cdot)$. By construction, we can assume wlog that \cdot (resp. \cdot) is a positive (resp. negative) literal. Hence $c(\cdot) \in A$, $c(\cdot) \in B$, as required. Edges $st; tu; vw; wx; xy; yv$ have vertices in different sets of the bipartition as well.

$G_h(\cdot)$ is also bipartite, since we can use the vertex assignment of graph $G_m(\cdot)$ and additionally assign t_1 and u_1 to A and B respectively.

$G_s(\cdot)$ is also bipartite, since we can use the vertex assignment of graph $G_m(\cdot)$ and remove vertices $w; y$ and their assignments. \square

In the following, we omit explicit dependence on \cdot . Moreover, when a statement is true for all graphs G_s, G_g, G_m , and G_h , we drop the subscript and simply write it holds for G .




Figure 2.2: A monotone 3-SAT instance $(x_1 - x_2 - x_3) \wedge (x_1 - x_2 - x_4) \wedge (x_2 - x_4 - x_5)$ and the associated graph $G_s(\cdot)$ and preference system. Disk and circle labels of vertices provide a bipartition of $G_s(\cdot)$.




Figure 2.3: A monotone 3-SAT instance $\phi = (x_1 \wedge x_2 \wedge x_3) \wedge (\neg x_1 \wedge \neg x_2 \wedge \neg x_4) \wedge (x_2 \wedge x_4 \wedge x_5)$ and the associated graph G_ϕ and preference system. Disk and circle labels of vertices provide a bipartition of G_ϕ .

The preference system

We now define the preference system for the graphs G_s , G_g , G_m , and G_h defined in the previous section.

Consider the union I of all canonical incomplete preference patterns for gadgets corresponding to literals of ϕ . Recall that vertex sets of gadgets are disjoint subsets of V , with the exception of $g(\neg_1) \cap a(\neg_2)$, with \neg_2 being the literal following \neg_1 in some clause, and $u \in a(\neg)$, where \neg is the first literal in some clause. However, the canonical incomplete preference patterns assign I to edges incident to nodes $a(\neg)$. Hence, I is an incomplete preference pattern for G .

We will construct now preference patterns L_g , L_m and L_h , see Fig. 2.2, 2.3, 2.4, 2.5. (Recall that, for the sake of shortness, we write $L(v; v^\beta) = \emptyset$ when $L_s(v; v^\beta) = L_g(v; v^\beta) = L_m(v; v^\beta) = L_h(v; v^\beta) = \emptyset$). They are defined as follows:

$$L_s, L_g, L_m \text{ and } L_h \text{ agree with } I;$$




Figure 2.4: A monotone 3-SAT instance $\varphi = (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee \neg x_4) \wedge (x_2 \vee x_4 \vee x_5)$ and the associated graph $G_m(\varphi)$ and preference system. Disk and circle labels of vertices provide a bipartition of $G_m(\varphi)$.

For edges not contained in \bar{E} , set: $L(s; t) = 1$, $L_s(t; s) = L_g(t; s) = L_m(t; s) = 2$; $L_h(t; s) = 3$, $L(u; t) = L(t; u) = 1$, $L_g(s; p_1) = 2$, $L_g(s; p_2) = 3$, $L_g(p_1; s) = 2$, $L_g(p_2; s) = 1$, $L_g(p_1; p_3) = 1$, $L_g(p_3; p_1) = 2$, $L_g(p_3; p_4) = L_g(p_4; p_3) = 1$, $L_g(p_4; p_2) = L_g(p_2; p_4) = 2$, $L_h(t; t_1) = L_h(u; u_1) = 2$, $L_h(t_1; t) = L_h(u_1; u) = 1$, $L_h(u_1; t_1) = L_h(u_1; t_1) = 2$, $L_g(v; w) = j\Gamma(v)j$, $L_g(w; x) = 2$, $L(x; w) = 1$, $L(x; y) = 2$, $L(y; x) = 1$, $L_g(y; w) = 2$, $L_g(w; y) = 1$, $L_m(v; w) = L_h(v; w) = j\Gamma(v)j - 1$, $L_m(v; y) = L_h(v; y) = j\Gamma(v)j$, $L_m(w; v) = L_h(w; v) = 2$, $L_m(y; v) = L_h(y; v) = 2$, $L_g(w; v) = 3$, $L(x; w) = 1$, $L(x; y) = 2$, $L(y; x) = 1$;

we complete the preference list for the remaining edges as follows:

- $ub(\cdot)$, for all literals \cdot that are first in some clause, and $uc(\cdot)$ for all negative literals that are first in some clause. Since $L_g(u; t) = L_m(u; t) = 1$ and no other edge is incident to u , we can assign $L_g(u; v) = L_g(u; v)$ for $v = b(\cdot)$




Figure 2.5: A monotone 3-SAT instance $\Phi = (x_1 \wedge x_2 \wedge x_3) \wedge (\neg x_1 \wedge \neg x_2 \wedge \neg x_4) \wedge (x_2 \wedge x_4 \wedge x_5)$ and the associated graph $G_h(\Phi)$ and preference system. Disk and circle labels of vertices provide a bipartition of $G_h(\Phi)$.

or $v \geq f b(\cdot); c(\cdot)g$, through an arbitrary bijection to $f2; \dots; j\Gamma(v)jg$. Since $L_h(u; t) = 1; L_h(u; u_1) = 2$ and no other edge is incident to u , we can assign $L_h(u; v)$ for $v = b(\cdot)$ or $v \geq f b(\cdot); c(\cdot)g$, through an arbitrary bijection to $f3; \dots; j\Gamma(v)jg$.

- $g(\cdot)v$ and $vg(\cdot)$, for all literals \cdot that are last in some clause. Since all such $g(\cdot)$ have degree 2 in G , and $L(g(\cdot); f(\cdot)) = 1$, we set $L(g(\cdot); v) = 2$. Moreover, $\Gamma(v)$ is composed of all such $g(\cdot)$ and vertex w for G_g , to which we already assigned $L_g(v; w) = j\Gamma(v)j$. We assign therefore $L_g(v; g(\cdot))$ via an arbitrary bijective map to $f1; \dots; j\Gamma(v)j - 1g$. For G_m and G_h , $\Gamma(v)$ is composed of all such $g(\cdot)$, vertices w and y , to which we already assigned $L_m(v; w) = L_h(v; w) = j\Gamma(v)j - 1$ and $L_m(v; y) = L_h(v; y) = j\Gamma(v)j$. We assign therefore $L_m(v; g(\cdot)) = L_h(v; g(\cdot))$ via an arbitrary bijective map to $f1; \dots; j\Gamma(v)j - 2g$.

- $c(\cdot)c(\cdot^\theta)$ and $c(\cdot^\theta)c(\cdot)$, where \cdot is a positive occurrence and \cdot^θ is a negative occurrence of the same variable. Note that we already assigned $L(c(\cdot);e(\cdot)) = 1$, $L(c(\cdot);b(\cdot)) = j\Gamma(c(\cdot))j$. We assign $L(c(\cdot);c(\cdot^\theta))$ for a fixed $c(\cdot)$ via an arbitrary bijective map to $f_2;:::;j\Gamma(c(\cdot))j - 1g$. Similarly, we assign $L(c(\cdot^\theta);c(\cdot))$ for a fixed $c(\cdot^\theta)$ via an arbitrary bijective map to $f_3;:::;j\Gamma(c(\cdot^\theta))j - 1g$.
- $a(\cdot)b(\cdot)$ for all positive literals \cdot that are not first in some clause. Recall that we have $L(a(\cdot);f(\cdot^\theta)) = 1$ for some \cdot^θ . We set therefore $L(a(\cdot);b(\cdot)) = 2$.
- $a(\cdot)b(\cdot)$, $a(\cdot^\theta)b(\cdot)$ for all negative literals \cdot that are not first in some clause. Recall that we have $L(a(\cdot);f(\cdot^\theta)) = 1$ for some \cdot^θ . We set therefore $L(a(\cdot);b(\cdot)) = 3$, $L(a(\cdot^\theta);b(\cdot)) = 2$.

We now derive properties of matchings that are popular in any of the preference systems introduced above and satisfy some further hypothesis. Recall that we omit the subscript and denote by $(G; <)$ any of the instances defined above .

In the proof of the next and following lemmas, we often use the characterization of popular matchings given in Theorem 1.3.5, hence we omit referring explicitly to this theorem each time.

Properties of certain popular matchings of $(G; <)$

Lemma 2.2.2. *Let M be a popular matching of $(G; <)$ that does not contain any evicted edge. Then M does not contain any consistency edge, i.e., $M \setminus E = ;$.*

Proof. Suppose by contradiction that M contains a consistency edge, connecting a positive gadget $H(\cdot)$ and a negative gadget $N(\cdot^\theta)$. We focus now on nodes of $N(\cdot^\theta)$, hence omit the dependency on \cdot^θ . Since the consistency edge incident to c is matched, if d is not matched, then cd would be a $(+;+)$ edge adjacent to an unmatched

vertex, which contradicts popularity. Hence, $bd \not\subseteq M$, which in turn implies $ef \subseteq M$, otherwise be is a $(+; +)$ edge adjacent to an unmatched vertex, again a contradiction (recall that $eh \subseteq M$ by hypothesis). This implies that he is a $(+; -)$ edge, fg is a $(+; +)$ edge, and $h; e; f; g$ is an M -alternating path containing a $(+; +)$ edge from an unmatched vertex, a contradiction. \square

Lemma 2.2.3. *Let M be a popular matching of $(G; <)$ that contains tu . Then:*

1. *The gateway (resp. apex) of each gadget is matched to its favorite partner;*
2. *M does not contain any consistency or evicted edges;*
3. *In each positive gadget $H(\cdot)$, exactly one of the following is true (dependency on \cdot is omitted):*

a) $F(H) := fbd; ceg \subseteq M$, or

b) $T(H) := fbc; deg \subseteq M$.

Similarly, in each negative gadget $N(\cdot^0)$, exactly one of the following is true (dependency on \cdot^0 is omitted):

a) $F(N) := fcd; beg \subseteq M$, or

b) $T(N) := fce; bdg \subseteq M$.

4. *If \cdot is a positive and \cdot^0 a negative occurrence of the same variable, then we cannot have both $T(H(\cdot)) \subseteq M$ and $T(N(\cdot^0)) \subseteq M$.*

Proof. 1. We consider negative and positive gadgets separately. Note that it is enough to prove the statement for the gateway, since the one for the apex follows by the fact that each apex is either also a gateway, or it is node u (that is matched to its favorite partner t by hypothesis).

Pick first a negative clause and one of its literals \cdot . We omit the dependency on \cdot in nodes of $N(\cdot)$. The favorite partner of g is f . Hence, suppose by contradiction $fg \not\subseteq M$. Note that since g is also f 's favorite partner, fg is a $(+; +)$ edge, which

implies $fe \not\subseteq M$, else f is an unmatched node adjacent to a $(+;+)$ edge starting from an unmatched node, a contradiction. Edge eh is therefore labeled $(-;+)$. The path $h;e;f;g$ is therefore an M -alternating path from an unmatched vertex to a $(+;+)$ edge, a contradiction concluding the proof for this case.

Consider any positive clause, and take, among its literals, the first one that violates the thesis. Call this literal ℓ . Again, we omit the dependency on ℓ in nodes of $H(\ell)$, and observe that f is g 's favorite partner. Therefore, $fd \subseteq M$, else f is unmatched and fg is a $(+;+)$ edge adjacent to an unmatched node, a contradiction. If e is unmatched, then de is a $(+;+)$ adjacent to an unmatched node, again a contradiction. Since $\Gamma(e) = fd;cg$ and $df \subseteq M$, we conclude $ec \subseteq M$. A similar argument (with e replaced by b) implies $ab \subseteq M$. Now, if ℓ is the first literal of the clause, $a = u$, contradicting $tu \subseteq M$. Else, $a = g(\ell')$ for the literal ℓ' preceding ℓ in the clause, contradicting the choice of ℓ .

2. Suppose by contradiction M contains an evicted edge. Pick any clause c with a negative literal such that M contains one or more evicted edges, and let $N(\ell)$ be the gadget associated to the first literal in clause c with $e(\ell)h(\ell) \subseteq M$. We omit the dependency on ℓ in nodes of $N(\ell)$. Part 1 implies that a is matched to its favorite partner. The hypothesis implies that $eh \subseteq M$. Hence, if $bd \not\subseteq M$, then be is a $(+;+)$ edge incident to the unmatched node b , a contradiction. Hence $bd \subseteq M$. Since $a; d; e$ are not matched to c in M , if c is not matched through a consistency edge it is unmatched, and ce is $(+;+)$ edge, a contradiction.

Hence $cc(\ell') \subseteq M$ for some $H(\ell')$. Arguments similar to those above imply $d(\ell')e(\ell') \subseteq M$ and $a(\ell')b(\ell') \subseteq M$, contradicting part 1. This concludes the proof that M does not contain any evicted edge. The fact that M does not contain any consistency edge immediately follows from what just proved and Lemma 2.2.2.

3. Fix a gadget $G(\ell)$ and omit dependency on ℓ in nodes. By part 1 and 2, a is matched to its first choice, and M does not contain any consistency or evicted edges.

Hence, we can restrict our attention to the subgraph induced by $b; c; d; e$. Simple case checking shows that, if any of those is unmatched, then there is a $(+; +)$ edge incident to an unmatched node, a contradiction to popularity. Hence, each of these vertices must be matched, and since they form a cycle with 4 vertices, one of the possibilities from the thesis of the lemma must hold.

4. Suppose we have both $T(H(\cdot)) \cap M$ and $T(N(\cdot^0)) \cap M$. Then the consistency edge $c(\cdot)c(\cdot^0)$ is $(+; +)$. Hence, $h(\cdot^0); e(\cdot^0); c(\cdot^0); c(\cdot)$ is an M -alternating path from an unmatched node to a $(+; +)$ edge, a contradiction. \square

Corollary 2.2.1. *If M is a popular matching of $(G; <)$, and $tu \notin M$, then there are no edges labelled $(+; +)$ in the subgraph of G induced by \bar{V} .*

Proof. From Lemma 2.2.3 we know that the gateway of each gadget is matched to its favorite partner because of Lemma 2.2.3, parts 3 and 4. Using Lemma 2.2.3, part 3 and 4, one easily checks that no edge between nodes of \bar{V} can be labelled $(+; +)$. \square

Lemma 2.2.4. *Let M be a popular matching of $(G; <)$ such that $tu \notin M$ and there exists a clause c such that $F(\cdot) \cap M$ for all literals \cdot of c . Then there is M -alternating path from s to v in G_M .*

Proof. Since $tu \notin M$, by Lemma 2.2.3, apexes and gateways are matched to their favourite partner. We assume that c is given by exactly three literals – say $\cdot_1; \cdot_2; \cdot_3$. Similar arguments apply for the case with less than three literals. The construction of the M -augmenting path can be followed in Fig. 2.6.

Assume first \cdot to be a positive clause, let $H(\cdot_1); H(\cdot_2); H(\cdot_3)$ be the corresponding gadgets, and assume $F(\cdot_1); F(\cdot_2); F(\cdot_3) \cap M$. Then $st, a(\cdot)b(\cdot), d(\cdot)f(\cdot), g(\cdot_3)v$ are $(+; \cdot)$ or $(\cdot; +)$ edges for $\cdot = \cdot_1; \cdot_2; \cdot_3$, and the promised M -alternating path is:

$$s; t; u; b(\cdot_1); d(\cdot_1); f(\cdot_1); g(\cdot_1); b(\cdot_2); d(\cdot_2); f(\cdot_2); g(\cdot_2); b(\cdot_3); d(\cdot_3); f(\cdot_3); g(\cdot_3); v;$$




Figure 2.6: An alternating path from s to w when all literals are False in positive (above) and negative (below) clauses.

Now assume that \bar{c} is a negative clause, and let $N(\bar{c}_1); N(\bar{c}_2); N(\bar{c}_3)$ be the gadget corresponding to its literals. Then $st, a(\bar{c})b(\bar{c}), e(\bar{c})f(\bar{c}), g(\bar{c}_3)v$ are $(+; -)$ or $(-; +)$ edges for $\bar{c} = \bar{c}_1; \bar{c}_2; \bar{c}_3$, and the the augmented path is

$$s; t; u; b(\bar{c}_1); e(\bar{c}_1); f(\bar{c}_1); g(\bar{c}_1); b(\bar{c}_2); e(\bar{c}_2); f(\bar{c}_2); g(\bar{c}_2); b(\bar{c}_3); e(\bar{c}_3); f(\bar{c}_3); g(\bar{c}_3); v:$$

□

From matchings to feasible assignments

We are now ready to prove one direction of the relation between feasible assignments for \mathcal{C} and matchings in the graphs we constructed.

Lemma 2.2.5. *Let M be a stable matching without an $s \rightarrow v$ augmenting path in G_s (resp. a popular matching that does not contain wx and contains tu in G_m, G_h or G_g). Then there exists a feasible assignment in \mathcal{C} .*

Proof. Note that in all cases M is popular. In G_s , t and u are each other's favorite partners and M is assumed to be stable. We conclude that $tu \in M$ for all cases.

Define an assignment $\bar{x}(M)$ of the variables of \mathcal{M} as follows: if $T(\bar{c}) \in M$ for some positive (resp. negative) occurrence \bar{c} of variable x_i , then set $\bar{x}_i(M)$ to be true (resp. false). Else, set $\bar{x}_i(M)$ to true or false arbitrarily. Note that $\bar{x}(M)$ is well-defined, since by Lemma 2.2.3, part 4, we cannot assign the same variable to true and false simultaneously.

We now prove that $\bar{x}(M)$ is a feasible assignment. It suffices to show that, for each clause c , $T(\bar{c}) \in M$ for some literal \bar{c} in clause c . For G_s , the statement follows by Lemma 2.2.4.

In the cases of G_m and G_h , wx is a $(+; +)$ edge and s is unmatched. Assume by contradiction that there is a clause c^ℓ such that all literals are set to false. We consider cases with positive and negative literals separately. Without loss of generality assume that c^ℓ contains exactly three positive literals $\bar{c}_1; \bar{c}_2; \bar{c}_3$. Notice that by Lemma 2.2.3, part 3, $F(\bar{c}_1); F(\bar{c}_2); F(\bar{c}_3) \in M$, and $st, a(\bar{c}_1)b(\bar{c}_1), d(\bar{c}_1)f(\bar{c}_1), g(\bar{c}_3)v$ are $(+; -)$ or $(-; +)$ edges for $\bar{c} = \bar{c}_1; \bar{c}_2; \bar{c}_3$. There is an M -alternating path:

$$s; t; u; b(\bar{c}_1); d(\bar{c}_1); f(\bar{c}_1); g(\bar{c}_1); b(\bar{c}_2); d(\bar{c}_2); f(\bar{c}_2); g(\bar{c}_2); b(\bar{c}_3); d(\bar{c}_3); f(\bar{c}_3); g(\bar{c}_3); v; w; x;$$

Consider now a case with three negative literals $\bar{c}_1; \bar{c}_2; \bar{c}_3$. Then $st, a(\bar{c}_1)b(\bar{c}_1), e(\bar{c}_1)f(\bar{c}_1), g(\bar{c}_3)v$ are $(+; -)$ or $(-; +)$ edges for $\bar{c} = \bar{c}_1; \bar{c}_2; \bar{c}_3$. The following is an M -alternating path:

$$s; t; u; b(\bar{c}_1); e(\bar{c}_1); f(\bar{c}_1); g(\bar{c}_1); b(\bar{c}_2); e(\bar{c}_2); f(\bar{c}_2); g(\bar{c}_2); b(\bar{c}_3); e(\bar{c}_3); f(\bar{c}_3); g(\bar{c}_3); v; w; x; y;$$

In both case we obtain a contradiction to popularity, since the above alternating paths contain an unmatched node s and $wx = (+; +)$.

In the case of G_g , wx is a $(+; +)$ edge and p_2 is unmatched, otherwise $p_2s \in M$ or s is unmatched. In both cases this would contradict popularity. We conclude that s is matched, otherwise $sp_2 = (+; +)$ and p_2 is unmatched, which would contradict

popularity. Assume by contradiction that there is a clause c^ℓ such that all literals are set to false. We consider cases with positive and negative literals separately. Without loss of generality assume that c^ℓ contains exactly three positive literals $\ell_1; \ell_2; \ell_3$. Notice that by Lemma 2.2.3, part 3, $F(\ell_1); F(\ell_2); F(\ell_3) \in M$, and $p_3p_1, p_2p_4, st, a(\ell_1)b(\ell_1), d(\ell_1)f(\ell_1), g(\ell_3)v$ are $(+; -)$ or $(-; +)$ edges for $\ell = \ell_1; \ell_2; \ell_3$. There is an M -alternating path: $p_2; p_4; p_3; p_1; s; t; u; b(\ell_1); d(\ell_1); f(\ell_1); g(\ell_1); b(\ell_2); d(\ell_2); f(\ell_2); g(\ell_2); b(\ell_3); d(\ell_3); f(\ell_3); g(\ell_3); v; w; x$. Consider now a case with three negative literals $\ell_1; \ell_2; \ell_3$. Then $p_3p_1, p_2p_4, st, a(\ell_1)b(\ell_1), e(\ell_1)f(\ell_1), g(\ell_3)v$ are $(+; -)$ or $(-; +)$ edges for $\ell = \ell_1; \ell_2; \ell_3$. The following is an M -alternating path: $p_2; p_4; p_3; p_1; s; t; u; b(\ell_1); e(\ell_1); f(\ell_1); g(\ell_1); b(\ell_2); e(\ell_2); f(\ell_2); g(\ell_2); b(\ell_3); e(\ell_3), f(\ell_3); g(\ell_3); v; w; x; y$. Similarly to above in both case we obtain a contradiction to popularity, since the above alternating paths contain an unmatched node p_2 and $wx = (+; +)$.

Since in each clause there is at least an ℓ such that $T(\ell) \in M$, we conclude that is satisfied by the assignment we constructed. \square

From feasible assignments to matchings

Now suppose that we have a consistent assignment of the literals, i.e., an assignment of the variables, which satisfies \mathcal{C} . We show how to construct matchings M_s in G_s , M_m in G_m , M_g in G_g , and M_h in G_h . As before, we drop the subscript when a certain statement holds for all the four cases.

- add tu and, for each literal ℓ , $f(\ell)g(\ell)$ to M ;
- for for each positive literal ℓ , add $T(\ell)$ to M if the corresponding variable is set to true; add $F(\ell)$ to M if the corresponding variable is set to false;
- for each negative literal ℓ , add $T(\ell)$ to M if the corresponding variable is set to false; add $F(\ell)$ to M if the corresponding variable is set to true;
- add $vw; xy$ to M when $M \notin M_s$;




Figure 2.7: A matching M (bold edges) of G_m corresponding to a feasible assignment of $\varphi = (x_1 _ x_2 _ x_3) \wedge (\neg x_1 _ : x_2 _ : x_4) \wedge (x_2 _ x_4 _ x_5)$ from Fig. 2.4: $x_1; x_4; x_5$ are set to True and $x_2; x_3$ to False.

add $p_3 p_4, p_1 s$ to M_g ; add $t_1 u_1$ to M_h .

Fig. 2.7 gives an example of a matching M_m corresponding to a feasible assignment of the clause from Fig. 2.4 in G_m . We show the following.

Lemma 2.2.6. *M is popular in G . Moreover, M_s is a stable matching in G_s that does not contain an M -augmenting path between s and v in G_M .*

Proof. We start by investigated the location of $(+; +)$ edges in G .

Claim 2.2.1. *There are no edges labelled $(+; +)$ in M in the subgraph of G induced by \bar{V} .*

Proof. Let ij be a $(+; +)$ edge in M for some vertices i and j in the given subgraph. This cannot be incident to $t; u; a(\cdot); g(\cdot)$ for each \cdot , and $f(\cdot)$ for negative literals \cdot , since those are matched to their favorite partner. Since the neighbors of both s and v

are either t or $g(\cdot)$ for some \cdot , it cannot be incident to v and s as well. Moreover, in positive literals \cdot , we have that $d(\cdot)$ is always matched to someone that it prefers to $f(\cdot)$, hence no $(+;+)$ edge is incident to $f(\cdot)$. We are left with edges between nodes $b; c; d; e$ of the gadgets.

Suppose first, ij is a consistency edge, say $c(\cdot)c(\cdot^\theta)$, with \cdot positive literal. Then $c(\cdot)b(\cdot); c(\cdot^\theta)e(\cdot^\theta) \not\subseteq M$. By construction, this implies that $T(H(\cdot)); T(N(\cdot^\theta)) \subseteq M$, a contradiction to the fact that \bar{X} is an assignment to variables in \cdot .

Hence, ij must connect nodes $b; c; d; e$ within the same gadget. This cannot happen. \square

Claim 2.2.2. *The only edge labeled $(+;+)$ in $G_g; G_m$ and G_h is wx . There are no edges labeled $(+;+)$ in G_s .*

Proof. From Claim 2.2.1, it suffices to consider edges incident to at least one of $f; s; t; t_1; u_1; p_1, p_2; p_3; p_4; w; x; y; g$. One easily checks that none of these edges is labelled $(+;+)$, other than wx . This also implies the second statement. \square

Note that Claim 2.2.2 allows us to conclude that M_s is stable, and that conditions (i) and (ii) of Theorem 1.3.5 are satisfied for $M_h; M_g; M_m$. In order to conclude the proof, we need to understand M -alternating paths in G_M . Recall that, in G_M , $(\cdot; \cdot)$ edges are deleted. We started by investigating the location of $(\cdot; \cdot)$ edges in G .

Claim 2.2.3. *If a positive clause contains a true literal \cdot , then $a(\cdot); b(\cdot) = (\cdot; \cdot)$ in M . If a negative clause contains a true literal \cdot^θ , then $e(\cdot^\theta)f(\cdot^\theta) = (\cdot; \cdot)$ in M .*

Proof. Consider $H(\cdot)$ and $N(\cdot^\theta)$ as in the hypothesis of the claim. By construction, $b(\cdot)c(\cdot); d(\cdot)e(\cdot) \not\subseteq M$ and $c(\cdot^\theta)e(\cdot^\theta); b(\cdot^\theta)d(\cdot^\theta) \not\subseteq M$. This implies that $a(\cdot)b(\cdot) = (\cdot; \cdot)$ and $e(\cdot^\theta)f(\cdot^\theta) = (\cdot; \cdot)$. \square

Define an M -alternating path P in $G_g; G_m; G_h$ to be *malicious* if it starts from an unmatched node and contains a $(+;+)$ edge. Observe that, by construction, the only

unmatched nodes are p_2 in G_g , s in $G_m; G_h$, and $h(\cdot)$ in all graphs for all negative literals \cdot .

Claim 2.2.4. *For any clause c , there is no malicious path from s to wx in $G_m; G_h$ (from p_2 to wx in G_g) that is contained in $f s; t; u; v; w; x; p_1; \dots; p_4 g [\bigcup_{\cdot \in c} V(\cdot)$, and there is no M -augmenting path from s to v that is contained in $f s; t; v g [\bigcup_{\cdot \in c} V(\cdot)$.*

Proof. Since we have an assignment of the literals which satisfies \cdot , each clause contains at least one true literal. According to Claim 2.2.3, $a(\cdot)b(\cdot) = (\cdot; \cdot)$ and $e(\cdot)f(\cdot) = (\cdot; \cdot)$ for some literals \cdot and \cdot^θ in positive and negative clauses correspondingly. Since any path that satisfies the hypothesis with c a positive (resp. negative) clause must pass through edges $a(\cdot)b(\cdot)$ (resp. $e(\cdot)f(\cdot)$) for all $\cdot \in c$, the thesis follows. \square

Claim 2.2.5. *A malicious path starting at $h(\cdot)$ does not traverse $f(\cdot)$. In particular, there is no malicious path from $h(\cdot)$ to wx that is contained in $f s; t; t_1; u_1; p_1; p_2; p_3; p_4; u; v; w; x; y g [\bigcup_{\cdot \in c} V(\cdot)$ for any negative clause c .*

Proof. $h(\cdot)$ is unmatched and $e(\cdot)f(\cdot) \not\subseteq M$ regardless of the true or false assignment to literal \cdot . Hence P cannot take both $h(\cdot)e(\cdot)$ and $e(\cdot)f(\cdot)$, since none of them is in M . Since $f(\cdot)$ has degree 2, $f(\cdot) \not\subseteq P$. We conclude that there is no alternating path from $h(\cdot)$ to wx contained only in c . \square

Claim 2.2.6. *Let \cdot be a positive literal, and let P be a malicious path in $G_h; G_m$; or G_g or an M -augmenting path in G_s . Then $c(\cdot)c(\cdot^\theta)$ is an edge of P for at most one \cdot^θ and if that happens $f(\cdot) \not\subseteq P$.*

Proof. The first statement follows from the definition of path. Now assume there is \cdot^θ such that $c(\cdot)c(\cdot^\theta) \subseteq P$. We omit dependency on \cdot in nodes. By construction, $cc(\cdot^\theta) \not\subseteq M$. Suppose by contradiction $f \subseteq P$. Since f has degree 2, fd is an edge of P . We have two possibilities. First, if \cdot is true, then $bc; de \subseteq M$. Then P or its

inverse (i.e. the path obtained traversing nodes from P in opposite order) contains the subpath $f; d; e; c; c(\text{ }^\theta)$, a contradiction to the fact that P is M -alternating. Second, if ^θ is false, then $bd; ce \not\subseteq M$. Then P or its inverse contains the subpath $f; d; b; c; c(\text{ }^\theta)$, again a contradiction to the fact that it is M -alternating. \square

Claim 2.2.7. *Let ^θ be a negative literal, and let P be a malicious path in $G_h; G_m$; or G_g or an M -augmenting path in G_s . Then $c(\text{ }^\theta)c(\text{ }^\theta)$ is an edge of P for at most one ^θ and if that happens and $a(\text{ }^\theta) \notin u$, then $a(\text{ }^\theta) \not\subseteq P$.*

Proof. Assume there is ^θ such that $c(\text{ }^\theta)c(\text{ }^\theta) \subseteq P$. Again, we omit explicit dependency on ^θ in nodes, and the first part of the statement is immediate. If $a \notin u$, then $\Gamma(a) = ff(\text{ }^\theta); b; cg$ for some ^θ . We have two possibilities. First, if ^θ is true, then $ce; bd \subseteq M$. This implies that P or its inverse contains the subpath $c(\text{ }^\theta); c; e$. Hence if $a \subseteq P$, then ac is not an edge of P . Since $af(\text{ }^\theta) \subseteq M$ for some ^θ by construction and P is M -alternating, both ab and $bd \subseteq M$ are edges of P . But then we must have that cd is an edge of P , a contradiction to P being a path. Second, if ^θ is false, then $cd; be \subseteq M$. Hence, either P or its inverse contains the subpath $c(\text{ }^\theta); c; d; b; e$. Thus, $ac; ab$ are not edges of P , which implies that $a \not\subseteq P$. \square

We can now conclude the proof of the lemma. First, we show that G_s has no M -augmenting path between s and v . By contradiction, assume there is such path P in G_s . Based on Claim 2.2.4 there is a consistency edge that belongs to P . Take the first consistency edge $c(\text{ }^\theta)c(\text{ }^\theta)$ that P traverses, and suppose that $c(\text{ }^\theta)$ is traversed by P before $c(\text{ }^\theta)$. This means that $a(\text{ }^\theta)$ and z are two consecutive nodes of P , traversed in this order and before $c(\text{ }^\theta)$, for some $z \subseteq V(\text{ }^\theta)$ (possibly $z = c(\text{ }^\theta)$). Suppose first ^θ is a positive literal. Hence, $a(\text{ }^\theta)b(\text{ }^\theta)$ is an edge of P . By Claim 2.2.3, the variable corresponding to ^θ is set to false, hence $F(N(\text{ }^\theta)) \not\subseteq M$ by construction. Using again Claim 2.2.3, we deduce $f(\text{ }^\theta) \not\subseteq P$. Hence P , in order to reach v , must traverse $a(\text{ }^\theta)$ after $c(\text{ }^\theta)$. If $a(\text{ }^\theta) = u$, then P cannot traverse u again. Hence $a(\text{ }^\theta) \notin u$.

u and, by Claim 2.2.7, $a(\cdot^\theta) \not\geq P$. In both cases, we obtain a contradiction. If conversely \cdot is a negative literal, by Claim 2.2.7, $a = u(\cdot)$. If $F(N(\cdot)) \not\geq M$, then by construction $cd;be \geq M$. Hence, P contains the subpath $u; c; d$ or the subpath $u; b; e; c; d$ (we omitted dependency on \cdot). In both cases, $c(\cdot)c(\cdot^\theta)$ cannot be an edge of P , a contradiction. If conversely $fce;bdg = T(N(\cdot)) \not\geq M$, then P contains the subpath $u; c; e$ or $u; b; d; c; e$ (again, we omitted the dependency on \cdot). In both cases, $cc(\cdot^\theta)$ cannot be an edge of P , a contradiction.

In order to deduce that M_m, M_h, M_g are popular, it suffices to show that condition (iii) of Theorem 1.3.5 is satisfied. Suppose, by contradiction, there exists a malicious path P starting from s in $G_m; G_h$, or from p_2 in G_g . Then $P = s; t; u; \dots$ in $G_m; G_h$, or $P = p_2; p_4; p_3; p_1; s; t; u; \dots$ in G_g . Claim 2.2.4 implies that there is at least one consistency edge that belongs to P . Take the first consistency edge $c(\cdot)c(\cdot^\theta)$ that P traverses, and suppose that $c(\cdot)$ is traversed by P before $c(\cdot^\theta)$. This means that $a(\cdot)$ and z are two consecutive nodes of P , traversed in this order and before $c(\cdot)$, for some $z \geq V(\cdot)$ (possibly $z = c(\cdot)$). Suppose first \cdot is a positive literal. Hence, $a(\cdot); b(\cdot) \geq P$. By Claim 2.2.3, \cdot is false, hence \cdot^θ is true. Using again Claim 2.2.3, we deduce $f(\cdot^\theta) \not\geq P$. Hence P , in order to reach wX , must traverse $a(\cdot^\theta)$ after $c(\cdot^\theta)$. If $a(\cdot^\theta) = u$, then P cannot traverse u again. Hence $a(\cdot^\theta) \notin u$ and, by Claim 2.2.7, $a(\cdot^\theta) \geq P$. In both cases, P cannot reach wX , a contradiction. If conversely \cdot is a negative literal, by Claim 2.2.7 $a = u(\cdot)$. If \cdot is set to true, then by construction $cd;be \geq M$. Hence, P contains the subpath $u; b; c; d; e$ or the subpath $u; c; e$ (we omitted dependency on \cdot). In both cases, $c(\cdot)c(\cdot^\theta)$ cannot be an edge of P , a contradiction. If conversely \cdot is set to false, then P contains the subpath $u; b; e$ or $u; c; d$ (again, we omitted the dependency on \cdot). In the latter case, $cc(\cdot^\theta)$ cannot be an edge of P . In the former, the inverse of P must contain the subpath $c(\cdot^\theta); c; d; b$, a contradiction (b would have degree 3 in P).

Suppose now there exists a malicious path P starting from $h(\cdot^\emptyset)$. Since $tu \geq M$,

$u \not\geq P$. From Claim 2.2.5, P must traverse (in this order) a consistency edge $c(\lrcorner^0)c(\lrcorner)$ with \lrcorner^0 being a false literal. Pick the first such edge. By Claim 2.2.6, $f(\lrcorner) \geq P$. Hence, after traversing $c(\lrcorner)$, P visits only gadgets corresponding to literals preceding \lrcorner in the clause, until it traverses (in this order) another consistency edge $c(\lrcorner^{000})c(\lrcorner^{0000})$, since as we argued above, $u \not\geq P$. This contradicts Claim 2.2.6, since it must also be that $f(\lrcorner^{000}) \geq P$.

Since s and the $h(\lrcorner)$ with \lrcorner being any negative literal are the only unmatched vertices, we conclude that the condition (iii) is also satisfied, and $M_s; M_h; M_g$ are popular. \square

2.3 Proof of Theorems 2.1.1–2.1.3

We have now all facts needed to prove our first results.

Proof of Theorem 2.1.1. The problem is clearly in NP. Consider the graph $G = G_s$ constructed in Section 2.2 and the corresponding Monotone 3-SAT formula φ . Lemma 2.2.5 implies that, if G_s has a stable matching M with no M -augmenting path in G_M , then φ is satisfiable. Conversely, Lemma 2.2.6 implies that, if φ is satisfiable, then G_s contains a stable matching M with no augmenting path in G_M . The thesis then follows from Theorem 1.4.3.

Proof of Theorem 2.1.2. Suppose first $U^{in} = F^{in} = F^{out} = \emptyset$. Then the problem boils down to deciding if there exists a popular matching that does not contain a given set of nodes. By Lemma 1.3.3, it is enough to check if this is true for a(ny) stable matchings in $(G; <)$. It is well-known that this can be done in polynomial time, see e.g. [22].

Suppose now $U^{out} = F^{in} = F^{out} = \emptyset$. Then, again by Lemma 1.3.3, the problem corresponds to deciding if there exists a dominant matching of $(G; <)$ that contains a given set of nodes. This can be done efficiently due to Theorem 1.3.7.

Suppose now $U^{out} = U^{in} = F^{out} = ;$, and $F^{in} = feg$. This is the *popular edge problem*, that is shown in [13] to be solvable in polynomial time. We call the case when $U^{out} = U^{in} = F^{in} = ;$, and $F^{out} = feg$ the *unpopular edge problem*. We show that a solution to the latter follows from the solution to the popular edge problem [13]. Define E_s and \bar{E}_s as:

$$E_s = \{e \in E : \exists \text{ stable matching } M \text{ s.t. } e \in M\}$$

$$\bar{E}_s = \{e \in E : \exists \text{ stable matching } M \text{ s.t. } e \notin M\}$$

E_d, \bar{E}_d (resp. E_p, \bar{E}_p) are defined similarly, by replacing “stable” with “dominant” (resp. popular). It is proved in [13] that $E_p = E_s \cup E_d$. We now argue that $\bar{E}_p = \bar{E}_s \cup \bar{E}_d$.

Consider any $e \in \bar{E}_s \cup \bar{E}_d$. Since there is a stable or dominant matching that does not contain e , it follows that $e \in \bar{E}_p$. We conclude that $\bar{E}_p = \bar{E}_s \cup \bar{E}_d$. Consider any $e = ij \in \bar{E}_p$. There is a popular matching P s.t. $e \notin P$, and i or j is matched (if i and j are unmatched then $(i;j) = (+;+)$, and M is not popular). Wlog assume $jk \in M$. It follows from $E_p = E_s \cup E_d$ that there exists a stable or dominant matching M^0 s.t. $jk \in M^0$, hence $ij \notin M^0$. We conclude that $\bar{E}_p = \bar{E}_s \cup \bar{E}_d$.

Since $\bar{E}_p = \bar{E}_s \cup \bar{E}_d$, we can solve the forbidden edge problem by checking if $e \in \bar{E}_s$ or $e \in \bar{E}_d$. This can be done efficiently using Theorem 1.3.3 and Theorem 1.3.7.

We now move to NP-Complete cases. Testing whether a matching is popular can be performed in polynomial time, see Theorem 1.3.4. Hence all considered problems are in NP. Consider now a Monotone 3-SAT instance ϕ and the graph $G = G_m$ constructed as in Section 2.2. We claim that G admits a popular matching M with edges tu and xy if and only if ϕ is satisfiable. The *only if* case follows from Lemma 2.2.5. For the *if case*, recall that, in Section 2.2, starting from a feasible assignment to ϕ , we construct a matching M such that $tu; xy \in M$ (by construction) and M is popular

(by Lemma 2.2.6).

Moreover, note that, for a popular matching M of G , the following are equivalent:

$$tu \in M, \quad st \notin M, \quad s \notin V(M):$$

Similarly, the following are equivalent:

$$xy \in M, \quad wx \notin M, \quad y \in V(M):$$

Hence, we settled all NP-Complete cases with $|U^{in}| + |U^{out}| + |F^{in}| + |F^{out}| = 2$.

To show NP-Completeness when sets U^{out} and/or F^{out} have bigger size, we can add nodes $h(\cdot)$ to U^{out} and/or evicted edges to F^{out} , since we know from Lemma 2.2.3 that those nodes and edges are never in a popular matching that contains tu . If U^{in} and/or F^{in} have bigger sizes, we can repeatedly add nodes $z;k$, with z adjacent to k and k adjacent to $g(\cdot)$ for some literal \cdot such that $z;k$ are each other's favorite partner. Then clearly zk belongs to any popular matching, so it can be added to F^{in} , and z to U^{in} .

Proof of Theorem 2.1.3. Because of the case $|F^{in}| = 2, U^{out} = U^{in} = F^{out} = \emptyset$ of Theorem 2.1.2, it is NP-hard to decide whether a graph has a popular matching with two given edges. Give those edges weight 1, and every other edge weight 0, and solve mwp. Theorem 2.1.2 implies that mwp cannot be solved within a factor better than $1/2$.

Because of the case $|F^{out}| = 2, U^{out} = U^{in} = F^{in} = \emptyset$ of Theorem 2.1.2, it is NP-hard to decide whether a graph has a popular matching without two given edges. Give those edges weight 1, and every other edge weight 0, and solve miwp. Theorem 2.1.2 implies that it is NP-hard to decide if the optimum is 0, and the statement follows.

In order to prove the existence of a $\frac{1}{2}$ -approximation algorithm for mwp, consider the following claim.

Claim 2.3.1. Consider an instance $(G; <); w$ of *mwp*. Let M_{st} and M_{dom} be the maximum weight stable and dominant matchings of G respectively. Then $w(M) \leq 2W$ for any popular matching M , where $W = \max(w(M_{st}); w(M_{dom}))$.

Proof. Assume by contradiction that there exists a popular matching M such that $w(M) > 2W$. It was shown in [13] that any popular matching M can be partitioned into a stable matching M_0 restricted to a subgraph G^s and a dominant matching $M_1 = M \setminus M_0$ restricted to $G \setminus G^s$, and there always exist matchings M_0^s and M_1^d such that $M_{st}^s = M_0 \cup M_1^s$ is stable and $M_{dom}^d = M_0^d \cup M_1$ is dominant in G . First assume $w(M_0) \geq w(M_1)$. This implies $w(M_0) = \frac{1}{2}(w(M_0) + w(M_0)) + \frac{1}{2}(w(M_0) + w(M_1)) = \frac{1}{2}w(M)$. Consider $M_{st}^s = M_0 \cup M_1^s$, $w(M_{st}^s) = w(M_0) + w(M_1^s)$. Since weights are non-negative, $w(M_{st}^s) \geq w(M_0) + \frac{1}{2}w(M)$. We conclude that $W \geq w(M_{st}^s) \geq \frac{1}{2}w(M)$. This contradicts to the assumption that $w(M) > 2W$. If $w(M_1) \geq w(M_0)$, apply the same argument to M_{dom}^d . \square

The 1=2-approximation immediately follows from the previous claim, and the fact that a dominant and stable matching of maximum weight can be computed in polynomial time, see Chapter 1.

Next corollary deals with weighted popular matching problems where weights are given on the *nodes* instead of the edges.

Corollary 2.3.1. *The following problem*

Problem 2.3.1 (Node-weighted popular matching problem). **Given:** A marriage instance $(G; <)$, with weights w on nodes. **Find:** A popular matching M of $(G; <)$ that maximizes $w(M)$.

is NP-Hard and not approximable to any factor in polynomial time, unless $P=NP$.

Conversely, if $w \geq 0$ or $w \leq 0$, it can be solved in polynomial time.

Proof. Consider the instance $(G_m; <)$ as defined in Section 2.2, and w being the vector with -1 in the component corresponding to S , 1 in the components corresponding

to y , and 0 otherwise. As discussed in the proof of Theorem 2.1.2, it is NP-hard to decide if the optimum is strictly greater than 0. On the other hand, if $w = 0$ (resp. $w = 0$), Lemma 1.3.3 implies that the minimum (resp. maximum) of $w(M)$ with M popular is achieved at any stable (resp. dominant) matching, which can be easily computed, see again Chapter 1. \square

The next problem can be seen as an extension of pmffe when F^{out} is the only non-empty set.

Corollary 2.3.2. *The following problem:*

Problem 2.3.2 (Bounded intersection popular matching). **Given:** A marriage instance $(G; <)$, $F \subseteq E(G)$ and $q \in \mathbb{N}$. **Decide:** If there exists a popular matching M of $(G; <)$ such that $|M \cap F| = q$.

is NP-Complete.

Proof. Clearly the problem is in NP. Consider the preference system $(G_m; <)$ as defined in Section 2.2. As in the proof of Theorem 2.1.2, add q pairs of nodes $k; z$, such that kz is in any popular matching. Let F be the set given by this q pairs, plus st, wx . Then a popular matching M satisfies $|M \cap F| = q$ if and only if $M \cap F = \{st, wx\}$ if and only if φ is satisfiable. \square

Recall that testing whether a matching is popular can be performed in polynomial time, see Theorem 1.3.4. We conclude this section by showing that, on the other hand, deciding whether a subset of nodes can be matched among themselves only as to form a popular matching of the original instance is NP-Complete. We call this the exclusive popular set problem. It can also be seen as an extension of pmfee when U^{in} (or U^{out}) is the only non-empty set.

Corollary 2.3.3. *The following problem:*

Problem 2.3.3 (Exclusive popular set problem - eps). **Given:** A marriage instance $(G; <)$, $U \subseteq V(G)$ and $q \geq 2 \in \mathbb{N}$. **Decide:** If there exists a popular matching M of $(G; <)$ with $V(M) = U$.

is NP-Complete.

Proof. Clearly the problem is in NP. Consider the instance $G = G_m$ from Section 2.2, and set $U = (ft; u; v; w; x; yg \cup V(\cdot)) \cap [h(\cdot)]$, where the union ranges over all literals. Let M be a popular matching of $(G; <)$. We claim that $V(M) = U$ if and only if $st; wx \notin M$. The thesis then follows by the proof of Theorem 2.1.2. Let $V(M) = U$. Then by definition $st \notin M$, and $wx \notin M$, otherwise y is unmatched. Now suppose $st; wx \notin M$. Then $tu \in M$. By Lemma 2.2.3, we know that for each literal ℓ , $g(\ell)f(\ell) \in M$, $h(\ell)$ is M -exposed, and $c(\ell)c(\ell^0) \notin M$ for all ℓ^0 . Simple arguments (see e.g. the proof of Lemma 2.2.3, part 3), imply that, for a fixed ℓ , nodes $b(\ell); c(\ell); d(\ell); e(\ell)$ must be all matched, and all among themselves. This implies $V(M) = U$, as required. \square

2.4 Proof of Theorems 2.1.4 and 2.1.5

Proof of Theorem 2.1.5. We start by proving NP-completeness of pmgg. pmgg belongs to NP because of Theorem 1.3.4. We refer again to the graph G_g from Section 2.2.

Claim 2.4.1. *Let M be popular of $(G_g; <)$. Then $tu \in M$ and $wx \notin M$.*

Proof. Assume by contradiction that $tu \notin M$. $st \in M$, otherwise s is unmatched and adjacent to $tu = (+; +)$. $p_1; p_2$ must be matched, otherwise there is an alternating path $u; t; s; p_1$ or $u; t; s; p_2$ respectively that contains $tu = (+; +)$ and unmatched node p_1 or p_2 . This implies that $p_3p_4 \notin M$. In this case $t; u; s; p_2; p_4; p_3$ is an M -alternating path with $tu; p_4p_3 = (+; +)$, a contradiction to the popularity of M .

If $wx \in M$, then $wy = (+; +)$ is adjacent to unmatched node y . \square

Now suppose that the instance associated to G_g is satisfiable. Then, by Lemma 2.2.6, we conclude that G_g has a popular matching. Now let G_g has a popular matching M . By Claim 2.4.1, we conclude that M satisfies the hypothesis of Lemma 2.2.5, hence that is satisfiable. We conclude that pmgg is NP-complete.

In the remainder of the proof, we will use the modification of graph G_m to which edge vy has been added. We remark that all statements seen in Section 2.2 for G_m extend to this case. For simplicity of notation, we will still denote this new graph by G_m .

In order to conclude that mm (hm) is NP-complete, we cannot apply directly Lemma 2.2.5, since there are popular matchings that do not satisfy the hypothesis of the lemma.

For a popular matching M of $(G_m; <)$ ($(G_h; <)$), we therefore argue in Claims 2.4.2 - 2.4.8 that if M is a popular matching such that $wx \not\geq M$ and $tu \geq M$, then M is a middle (hidden) matching, while all other popular matchings in the instance are stable or dominant.

All popular matchings of $(G; <)$ can be partitioned in the following four sets:

$\mathcal{M}_1 = fM$ is a popular matching : $tu \geq M; xw \not\geq Mg$,

$\mathcal{M}_2 = fM$ is a popular matching : $tu \geq M; xw \geq Mg$,

$\mathcal{M}_3 = fM$ is a popular matching : $tu \not\geq M; xw \geq Mg$,

$\mathcal{M}_4 = fM$ is a popular matching : $tu \not\geq M; xw \not\geq Mg$.

In Claims 2.4.2 - 2.4.4 we consider $(G_m; <)$ and $(G_h; <)$ simultaneously.

Claim 2.4.2. *Let $M \geq \mathcal{M}_1$. Then M is neither stable nor dominant.*

Proof. Observe that $wx = (+; +)$ and s is exposed. Since $wx = (+; +)$ then M is not stable. Consider a negative gadget $N(\cdot)$ such that $a(\cdot) = u$. From Lemma 2.2.3 we know that either $c(\cdot)e(\cdot); b(\cdot)d(\cdot) \geq M$ or $c(\cdot)d(\cdot); b(\cdot)e(\cdot) \geq M$. In both cases

there is an augmenting path $s; t; u; c(\cdot); e(\cdot); h(\cdot)$ or $s; t; z; c(\cdot); d(\cdot); b(\cdot); e(\cdot); h$. This implies that M is not dominant by Theorem 1.3.6. \square

Claim 2.4.3. *Let $M \succeq M_2$. Then M is stable.*

Proof. We show the statement for G_h , the one for G_m following analogously. From Corollary 2.2.1 it follows that M restricted to $G \setminus fs; t; t_1; u_1; u; v; w; x; yg$ does not contain any $(+; +)$ edges. From Lemma 2.2.3 we know that $f(\cdot)g(\cdot) \succeq M$ for all \cdot . This implies that $g(\cdot)v \not\geq M$. We conclude that $vy \succeq M$, otherwise $vy = (+; +)$ and $v; y$ are unmatched. Note that $st; vw; yx; tt_1; uu_1 = (+; \cdot)$. M does not contain $(+; +)$ edges, hence it is stable. \square

Claim 2.4.4. *Let M be a popular matching and $tu \not\geq M$. Then*

1. *M does not contain $a(\cdot)b(\cdot)$ for any positive literal \cdot ;*
2. *M does not contain any consistency edges;*
3. *In each positive gadget $H(\cdot)$, exactly one of the following is true (dependency on \cdot is omitted):*

- a) *$fbd; ceg \succeq M$, or*
- b) *$fbc; deg \succeq M$.*

Moreover, $f(\cdot)g(\cdot) \succeq M$.

4. *If $f(\cdot)g(\cdot) \succeq M$ for some negative clause \cdot , then $e(\cdot^\theta)h(\cdot^\theta) \not\geq M$, for any literal \cdot^θ that follows literal \cdot in the same negative clause;*

Proof. 1. Suppose by contradiction that $a(\cdot)b(\cdot) \succeq M$ in some $H(\cdot)$, and let \cdot be the first literal in the clause that violates the thesis. If $a(\cdot) = u$, then $c(\cdot)$ must be matched. If $c(\cdot)e(\cdot) \succeq M$, then $t; u; b(\cdot); c(\cdot); e(\cdot); d(\cdot)$ is an alternating path with $tu; e(\cdot)d(\cdot) = (+; +)$, a contradiction to popularity. If $c(\cdot)c(\cdot^\theta) \succeq M$ for a negative literal \cdot^θ , then $d(\cdot^\theta)$ must be matched. We conclude $d(\cdot^\theta)b(\cdot^\theta) \succeq M$. Similarly to above, $e(\cdot^\theta)$ must be matched. Notice that $e(\cdot^\theta)f(\cdot^\theta) \not\geq M$, otherwise

$h(\cdot^0)e(\cdot^0)f(\cdot^0)g(\cdot^0)$ contains an unmatched node $h(\cdot^0)$ and $f(\cdot^0)g(\cdot^0) = (+; +)$. This implies that $e(\cdot^0)h(\cdot^0) \not\subseteq M$. In this case $t; u; b(\cdot); c(\cdot); c(\cdot^0); d(\cdot^0); b(\cdot^0); e(\cdot^0)$ contains $tu; b(\cdot^0)e(\cdot^0) = (+; +)$, a contradiction to popularity. If \cdot is not the first literal in the clause, let \cdot^{\emptyset} be the literal that immediately precedes \cdot in the clause. $f(\cdot^{\emptyset})d(\cdot^{\emptyset}) \subseteq M$, otherwise $f(\cdot^{\emptyset})$ is unmatched, and $f(\cdot^{\emptyset})g(\cdot^{\emptyset}) = (+; +)$. Similarly, we argue that $e(\cdot^{\emptyset})c(\cdot^{\emptyset}); b(\cdot^{\emptyset})a(\cdot^{\emptyset}) \subseteq M$. This contradicts to the assumption that \cdot is the first literal in the clause such that $b(\cdot)a(\cdot) \subseteq M$.

2. Assume that M contains a consistency edge, connecting a positive gadget $H(\cdot)$ and a negative gadget $N(\cdot^0)$. Edge $e(\cdot)d(\cdot) \subseteq M$, otherwise $c(\cdot)e(\cdot) = (+; +)$ is adjacent to unmatched node $e(\cdot)$. $a(\cdot)b(\cdot) \subseteq M$, otherwise $b(\cdot)d(\cdot) = (+; +)$ and $b(\cdot)$ is unmatched. This is a contradiction to part 1.

3. Consider a positive gadget $H(\cdot)$. We omit the dependency on \cdot in nodes of $H(\cdot)$. Suppose by contradiction that $cb; ce \not\subseteq M$. Recall from part 2 that M does not contain any consistency edges. This implies that c is unmatched and cb is a $(+; +)$ edge incident to an unmatched node b . If $db; de \not\subseteq M$, then $ec \subseteq M$. Since $ab \not\subseteq M$ from part 1, then an alternating path $dec b$ contains $de = (+; +)$ and unmatched node b . In both cases we obtain a contradiction to popularity. $fg \subseteq M$, since d is matched to b or e , otherwise $fg = (+; +)$ and f is unmatched, which contradicts popularity.

4. Assume $f(\cdot)g(\cdot) \subseteq M$ and $e(\cdot^0)h(\cdot^0) \subseteq M$, for a literal \cdot^0 that follows \cdot . Node $b(\cdot^0)$ must be matched. If $b(\cdot^0)d(\cdot^0) \subseteq M$, then $c(\cdot^0)a(\cdot^0) \subseteq M$ by part 2 and $f(\cdot^{\emptyset})a(\cdot^{\emptyset}) \not\subseteq M$. This contradicts to $f(\cdot)g(\cdot) \subseteq M$, if $\cdot^{\emptyset} = \cdot$, or to popularity of M , since in the later case $e(\cdot^{\emptyset})f(\cdot^{\emptyset}) \subseteq M$, and $h(\cdot^{\emptyset})$ is unmatched. If $b(\cdot^0)a(\cdot^0) \subseteq M$ with a similar argument we deduce that $a(\cdot^0); b(\cdot^0); d(\cdot^0); c(\cdot^0)$ is an alternating cycle with a $(+; +)$ edge, a contradiction. \square

Claim 2.4.5. *Let $M \subseteq M_3 \sqcup M_4$ of $(G_m; >)$. Then M is dominant.*

Proof. We first show that there is exactly one evicted edge $e(\cdot)h(\cdot) \subseteq M$. From Claim

2.4.4, part 1, we deduce that $M(u) \supseteq N(\cdot)$ for some \cdot . We omit the dependency on \cdot in nodes of $N(\cdot)$. If $uc \supseteq M$, then $db \supseteq M$, otherwise $t;u;c;d$ is an alternating path with $tu = (+;+)$ and unmatched node d . Similarly, e is matched. If $ef \supseteq M$, then $h;e;f;g$ is an alternating path with $fg = (+;+)$ and h being unmatched. We conclude that eh (hence fg) $\supseteq M$. From Claim 2.4.4, part 4, we know that all other evicted edges in this clause are unmatched.

If $ub \supseteq M$, then $cd \supseteq M$, otherwise $t;u;b;d$ is an alternating path with $tu = (+;+)$ and unmatched node d . As above, e is matched, $ef \not\supseteq M$. This implies again that eh (hence fg) $\supseteq M$.

Consider a different negative clause. Let \cdot^θ be its first literal. Similarly to above $e(\cdot^\theta)$ must be matched, and $e(\cdot^\theta)f(\cdot^\theta) \not\supseteq M$. If $e(\cdot^\theta)h(\cdot^\theta) \supseteq M$, then $b(\cdot^\theta)d(\cdot^\theta) \supseteq M$, and $c(\cdot^\theta)$ is unmatched, since M does not contain any consistency edges (Claim 2.4.4, part 2). We conclude that $c(\cdot^\theta)e(\cdot^\theta)$ is a $(+;+)$ edge incident to an unmatched node, a contradiction.

We now show that all unmatched nodes belong to the same bipartition class. From Claim 2.4.4, part 3, we know that for any positive clause $H(\cdot)$, $b(\cdot);c(\cdot);e(\cdot);d(\cdot)$ are matched. This implies that $f(\cdot)g(\cdot) \supseteq M$ for all positive gadgets. From what shown above, $b(\cdot);c(\cdot);d(\cdot);e(\cdot)$ are matched in all first literals for all negative clauses. In all other negative literals evicted edges are unmatched. This implies that $b(\cdot);c(\cdot);e(\cdot);d(\cdot)$ are matched in all gadgets and $f(\cdot)g(\cdot) \supseteq M$ for all last literals \cdot of a negative clause. If $xw \supseteq M$, then $vy \supseteq M$; if $xw \not\supseteq M$, then $vw;xy \supseteq M$. We conclude that $h(\cdot)$ are the only unmatched nodes, and they belong to the same bipartition class, see Lemma 2.2.1. Hence, so there is no augmenting path. tu is a $(+;+)$ edge, so M is not stable. We conclude that M is dominant. \square

Claim 2.4.6. *Let $M \supseteq M_1$ of $(G_m; >)$. Then M is middle.*

Proof. From Claim 2.4.2 we know, that M is neither stable nor dominant. We will

show that there is a dominant matching of size larger than $|M|$. Consider a matching M^θ that contains the following edges:

$st; vy; wx;$

$a(\cdot^\theta)c(\cdot^\theta); b(\cdot^\theta)d(\cdot^\theta); e(\cdot^\theta)h(\cdot^\theta)$ for the first literal \cdot^θ in some negative clause such that such that $L(u; (c(\cdot^\theta))) = 1;$

$f(\cdot)g(\cdot)$ for all gadgets;

all literals, except \cdot^θ (from above), are set to False: $b(\cdot)d(\cdot); c(\cdot)e(\cdot)$ for all positive gadgets $H(\cdot)$, and $c(\cdot)d(\cdot); b(\cdot)e(\cdot)$ for all negative gadgets $H(\cdot)$.

See Figure 2.8 for an example of such assignment. $vw; yx = (+; -)$, and tu is the only $(+; +)$ edge. Vertices $h(\cdot)$, for all negative gadgets $H(\cdot)$, except $H(\cdot^\theta)$, are the only unmatched nodes. M^θ is popular, since there is no alternating path between tu and $h(\cdot)$. M^θ leaves unmatched only nodes from one set of the bipartition (nodes $h(\cdot)$). On the other hand, M leaves s unmatched (which belongs to the other set of the bipartition, see Lemma 2.2.1), hence $|M| < |M^\theta|$. We conclude that M^θ is dominant, and $|M^\theta| > |M| + 1$. \square

Claim 2.4.7. *Let $M \geq M_3 \sqcup M_4$ of $(G_h; >)$. Then M is dominant and leaves unmatched exactly node s and nodes $h(\cdot)$ for all literals \cdot .*

Proof. We first observe that $tt_1, uu_1 \geq M$. Suppose this is not the case. Then $t_1u_1 \geq M$ and $u; u_1; t_1; t$ is an alternating path containing two $(+; +)$ edges, a contradiction.

Consider first the first literal \cdot in some negative clause, omitting dependency on \cdot for nodes of $N(\cdot)$. We claim that e is matched to either c or b . Suppose this is not the case. From Claim 2.4.4, part 2, we deduce that c is either unmatched, or matched to d . The former implies that e is incident to the $(+; +)$ edge ce , a contradiction. Hence the latter applies. Since u is matched to u_1 , we deduce that b is unmatched and incident to the $(+; +)$ edge bd , a contradiction. Hence ec or $eb \geq M$. This implies:




Figure 2.8: A dominant matching M^0 (bold edges) corresponding to an example from Fig. 2.4.

$eh; ef \not\subseteq M$, hence $fg \subseteq M$, thus (via Claim 2.4.4, part 4) $eh \not\subseteq M$ for all literals \cdot in the same clause as \cdot . We can therefore repeat the arguments above and deduce that, within each negative clause $N(\cdot)$, nodes $b(\cdot); c(\cdot); d(\cdot); e(\cdot)$ are matched among themselves, and $f(\cdot)g(\cdot) \subseteq M$. Because of Claim 2.4.4, part 3, the same apply to gadgets in positive clauses. This implies that v is not matched to any node $g(\cdot)$.

Suppose first $wx \not\subseteq M$. Then $xy \subseteq M$ (since x is the first choice of y), hence $vw \subseteq M$. If conversely $wx \subseteq M$, then $vy \subseteq M$. We deduce that the only nodes left unmatched by M are s and $h(\cdot)$, for all clauses \cdot . Note that nodes $h(\cdot)$ belong to the same set of the bipartition (see Lemma 2.2.1). Moreover, the only nodes that can be reached from s via an M -alternating path in G_m are t and t_1 , for $t_1; u_1$ is a $(\cdot; \cdot)$ edge. Using Theorem 1.3.6, we deduce that M is dominant. \square

Claim 2.4.8. *Let $M \subseteq M_1$ of $(G_h; >)$. Then M is hidden.*

Proof. First, notice that all matchings from \mathcal{M}_1 have the same size. This follows from the following. Since $tu \in M$ then $t_1u_1 \in M$. Similarly, $vw, xy \in M$. From Lemma 2.2.3, we conclude that s and $h(\cdot)$, for all $N(\cdot)$, are the only unmatched nodes in any $M \in \mathcal{M}_1$.

We will show that a dominant matching has size jMj . Consider a matching M^0 such that contains the following edges:

$$tt_1; uu_1; vy; wx;$$

$$f(\cdot)g(\cdot) \text{ for all gadgets};$$

All edges that correspond to False assignment of all literals: $b(\cdot)d(\cdot); c(\cdot)e(\cdot)$ for all positive gadgets $H(\cdot)$, and $c(\cdot)d(\cdot); b(\cdot)e(\cdot)$ for all negative gadgets $H(\cdot)$.

tu is the only $(+; +)$ edge. Vertices s and $h(\cdot)$ for all \cdot are the only unmatched vertices. Since $t_1u_1 = (-; -)$ and $uu_1 \in M$, there is no alternating path from an unmatched node to $(+; +)$, we conclude that M^0 is popular. Notice that if M^0 is popular, then $st \notin M^0$. If $st \in M^0$ then $tu \notin M^0$ and $t_1u_1 \in M^0$. In this case tt_1u_1u is an alternating path with $tt_1; uu_1 = (+; +)$. If $st \notin M^0$, then either $tu; t_1u_1 \in M^0$ or $tt_1; uu_1 \in M^0$. Since all $h(\cdot)$ for all \cdot belong to the same bipartition class, Lemma 2.2.1, we conclude that M^0 is the largest popular matching. M^0 is dominant, since there is no augmenting path between s and $h(\cdot)$ for any \cdot . $jM^0j = jMj$. From Claim 2.4.2 we conclude that M is a hidden matching. \square

Popular matchings in graphs of bounded treewidth

In Chapter 2, we proved several hardness results on popular matchings. In particular, in the roommate case, even the problem of deciding if a popular matching exists is NP-Complete, see Theorem 2.1.5.

In this chapter, we show that polynomial-time solvability can be regained if we assume that the input graph has bounded treewidth. This applies with no other restriction on the input, i.e., we allow the graph to be non-bipartite. More precisely, we prove the following.

Theorem 3.0.1. *Let ℓ be a fixed constant. There exists a function $f(n) = O(n^{3\ell+7})$ such that, given as input a graph G with n nodes with strict preference lists \prec , treewidth at most ℓ , and a cost vector c on the edges of $(G; \prec)$, in time $f(n)$ we can either conclude that $(G; \prec)$ has no popular matching, or find a popular matching of $(G; \prec)$ of minimum cost.*

3.1 Introduction

As discussed in Chapter 1, bounded treewidth is a classical assumption that often turns intractable problems into tractable ones. A typical example is Maximum Independent (Stable) Set (MIS), for which a polynomial-time algorithm exists in bounded treewidth graphs [9]. MIS enjoys two nice properties, often shared by problems for which the bounded treewidth approach is successful. The first is *monotonicity*: if S is an independent set in a graph $G(V; E)$ and G' is a subgraph of G , then the

solution induced by S on G^∂ is also feasible. The second is *locality*: in order to check if S is an independent set, it suffices to verify, for each node of S , if any node of its neighborhood also belongs to S .

Interestingly, similar properties do not hold for popular matchings. Indeed, popularity is not a local condition, since it may depend on how nodes far away in the graph are matched. Moreover, if we take a graph G and a popular matching M , the subset of M contained in an induced subgraph of G may not be popular. Examples with both those features can be easily constructed by building on the characterization of popular matchings given by Theorem 1.3.5.




Figure 3.1: An example of popular vs locally-popular. In each graph, let M be the matching given by bold edges. *Top graph.* M is not popular, since there is an M -alternating path in G_M with a $(+; +)$ edge and an unmatched node. But if we take $X = \overline{v_1; v_2}g$ and $S = \overline{v_3; v_4}g$, matching $M_{X \setminus S}$ is $(S; X)$ -locally popular. *Bottom graph.* M is popular, hence it is $(X; S)$ -locally popular for every choice of S and X . But if we restrict M to the subgraph G^∂ of G induced by $v_1; v_2; v_3; v_4; v_5$, the matching we obtain is not popular in G^∂ , since there is a $(+; +)$ edge incident to the now unmatched node v_5 .

Our technique to prove the tractability of the minimum cost popular matching problem in the bounded treewidth case is as follows. We assume wlog that all matchings have different costs. This can be achieved efficiently by standard perturbation techniques. Given a vertex separator S of G and a connected component X of $G[V \setminus S]$, we define the family of $(S; X)$ -locally popular matchings. This contains all match-

ings that may potentially be extended to a popular matching in the whole graph by adding edges not incident to $X \setminus S$ (a formal definition is given in Section 3.1). As $(S; X)$ -locally popular matchings can be exponentially many even in graphs of bounded treewidth, we cannot store all of them. Instead, we divide them in classes (which we call *tipping points*), and show that, if a matching M of a class can be completed to a $(S^\theta; X^\theta)$ -locally popular matching for some sets with $X^\theta \supseteq X$ and $X^\theta \setminus S^\theta \supseteq X \setminus S$ by adding some matching M^θ not incident to $X \setminus S$, then all matchings of the same class can be extended via the same matching M^θ to a $(S^\theta; X^\theta)$ -locally popular matching. Hence, it is enough to keep only a representative for each class — the one of minimum cost, which we call the *leader*. Finally, we show how to iteratively construct tipping points and their leaders by building on a tree decomposition of the graph (see Algorithm 3). In particular, if the graph has bounded treewidth, a tree decomposition of bounded width can be computed in linear time [10], there will be only polynomially many tipping points and leaders, and they can be built in polynomial time, see Theorem 3.3.1.

Additional definitions

Recall that our input graph is $G = (V; E)$. For $S \subseteq V$, we denote by $G \setminus S$ the subgraph of G induced by $V \setminus S$. If $G \setminus S$ is disconnected, S is said to be a *vertex separator*. Given a matching M on G and a set $U \subseteq V$, the matching *induced* by M on U is given by $E[U] \setminus M$, and it is denoted by $M[U]$.

Now also fix $U \subseteq V$. Given a matching M of G , we say that M is a *U -matching* if for all edges $(u; v) \in M$, we have $fu; vg \setminus U \neq \emptyset$. Given a matching M of G , the *U -matching induced by M* is the set $M_U = \{(u; v) \in M : fu; vg \setminus U \neq \emptyset\}$.

Again, since all graphs we deal with are simple, we represent paths and cycles as ordered set of nodes, with the first node of a cycle coinciding with the last. This also allows us to distinguish between the *rst* and *last* node of a path. We will say

that e is an edge of P if it is an edge between two consecutive nodes of P . A U -path P is a path in G where either the first or the last node of P (or possibly both) belongs to U , and all other vertices of P do not lie in U (i.e., if $P = (v_1; \dots; v_k)$, then $v_i \notin U \quad \forall v_1; v_k g$).

Two paths $P = (v_1; \dots; v_k)$, $P^\theta = (v_1^\theta; \dots; v_k^\theta)$ are called U -disjoint if $P \setminus P^\theta \setminus U = \emptyset$. P, P^θ are said to be *internally vertex-disjoint* if, for all $i \geq 1$ and $j \geq 1$, $v_i \neq v_j^\theta$ with possibly the exception of $(i; j) \in \{(1; 1); (1; k^\theta); (k; 1); (k; k^\theta)\}$. Consider paths $P_1; P_2; \dots; P_q$ of G , $q \geq 2$, whose union contains at least two distinct vertices and with the following properties:

$$P_i \setminus P_j = \emptyset, \text{ unless } j = i + 1 \text{ or } i; j \in \{(1; 1); (q; q)\};$$

If $q \geq 3$, for $i = 1; \dots; q - 1$, the last node of P_i is the first node of P_{i+1} , and $P_i; P_{i+1}$ are otherwise disjoint. Moreover, P_1 and P_q are disjoint, with possibly the exception of the last node of P_q coinciding with the first node of P_1 ;

If $q = 2$, the last node of P_1 coincides with the first node of P_2 , the last node of P_2 may coincide with the first node of P_1 , and paths P_1 and P_2 are otherwise disjoint.

The *juxtaposition* of $P_1; \dots; P_q$ is the path (if the first node of P_1 and the last of P_q are different) or cycle (otherwise) defined from $(P_1; P_2; \dots; P_q)$ as above by removing consecutive repeated vertices.

Recall that we assume that no two matchings of G have the same cost.

Locally popular matchings, Configurations, and Tipping points

In this section, we fix a graph $G = (V; E)$ together with strict rankings of the neighbors of each node in V , a vertex separator S of G , and a connected component X of $G \setminus S$ (possibly $S = \emptyset$ and $X = V$).

Let M be a $(X [S)$ -matching of G . We will extensively work with graph $G_M[X [S]$ – that is, the subgraph of G_M induced by $X [S$. For a node v in S that is matched in M to a neighbor outside $X [S$, labels of $G_M[X [S]$ incident to v are a function of the edge $(v; M(v))$, however note that the edge $(v; M(v))$ is not in $G_M[X [S]$.

We say that M is $(S; X)$ -locally popular if none of the structures (i), (ii), and (iii) from Theorem 1.3.5 is a subgraph of $G_M[X [S]$. We remark that a node that is not matched in $G_M[X [S]$ may still be M -covered (hence not M -exposed). If this happens, such a node cannot be the M -exposed node in the path (iii) from Theorem 1.3.5. The following simple proposition relates the definitions of popularity and $(S; X)$ -local popularity.

Lemma 3.1.1. *Let $G; S; X$ be as above, S^θ be a vertex separator of G , and X^θ be one of the connected components of $G \setminus S^\theta$ such that $X^\theta [S^\theta = X [S$. Then:*

1. *Let M be a matching in G . Then M is $(; V)$ -locally popular if and only if it is popular.*
2. *Let M^θ be a $(S^\theta; X^\theta)$ -locally popular matching. Let $M := M^\theta_{X [S}$. Then M is an $(S; X)$ -locally popular matching.*

Proof. 1. It follows by definition and from the fact that $G_M[; [V] = G_M$.

2. Suppose M is not $(S; X)$ -locally popular. Then, one of the structures from Theorem 1.3.5 is a subgraph of $G_M[X [S]$ — call it P . We claim that P is a forbidden structure in $G_{M^\theta}[X^\theta [S^\theta]$ (again, in the sense of Theorem 1.3.5), hence M^θ is not $(S^\theta; X^\theta)$ -locally popular, a contradiction.

Indeed, $X [S = X^\theta [S^\theta$, and none of the edges of $M^\theta \setminus M$ is incident to $X [S$ by definition. We deduce $G_M[X [S] = G_{M^\theta}[X [S]$ and a node of $X [S$ is M -exposed if and only if it is M^θ -exposed. This concludes the proof. \square

We now introduce the concepts of *con guration* and *tipping point* in order to partition the family of $(S; X)$ -locally popular matchings into a (small) number of

classes. These definitions are related to the existence of certain structures that are not forbidden in a popular matching, but restrict the capability of extending locally popular matchings to popular matchings of the whole graph.

From now on, we also fix M to be a $(X \sqcup S)$ -matching of G . Let P_M be the set of M -alternating paths of $G_M[X \sqcup S]$. We associate to each $P \in P_M$ a 2-dimensional *parity* vector π , where the first component of π is defined to be 0 if the first edge of P is a matching edge, 1 otherwise. Similarly, the second component of the parity vector π takes values 0 or 1, depending on whether the last edge of the path is a matching edge or not. We also associate to P a two-dimensional *level* vector $\ell = (e; p)$, where $e \in \{0; 1; 2\}g$ is the number of M -exposed nodes of P , and $p \in \{0; 1\}g$ denotes the number of $(+, +)$ edges of P . If $p = 2$, then we say that P is of level 1. Note that not all parity-level combinations are possible. Moreover, if P is of level 1, then M is not $(S; X)$ -locally popular. Let

$$U = ((u_1; v_1); (u_2; v_2); \dots; (u_k; v_k)); \quad (3.1)$$

for some $k \in \mathbb{N}$, where $u_1; v_1; u_2; \dots; v_k \in S \sqcup \{f; g\}$, under the additional condition that $u_i \neq v_i$ for all $i \in [k]$, all pairs are different, and each node of S can appear at most twice in the collection. Moreover, let $L = (\ell_1; \dots; \ell_k)$, $\Pi = (\pi_1; \pi_2; \dots; \pi_k)$ and, for $i \in [k]$, $\ell_i \in \{0; 1; 2\}g - \{0; 1\}g$ and $\pi_i \in \{0; 1\}g - \{0; 1\}g$. The triple $\mathcal{C} = (U; L; \Pi)$ is called an $(S; X)$ -*configuration*.

We say that M is *active at* \mathcal{C} if there exist pairwise X -disjoint S -paths $P^1; \dots; P^k \in P_M$, such that, for $i \in [k]$, P^i is of level ℓ_i and parity π_i ; starts at $u_i \in S$ if $u_i \neq f$; and at some node of X otherwise; ends at $v_i \in S$ if $v_i \neq g$; and at some node of X otherwise. We call those paths the *certificate of* \mathcal{C} at M ¹.

¹Note that, if M is active at \mathcal{C} , it is also active at the configurations e.g. obtained by permuting entries of U (and of L and Π accordingly). This causes some redundancy, yet this will not affect our analysis, so we do not eliminate it.

Let $f(U_i; L_i; \Pi_i)g_{i=1, \dots, q}$ be the collection of all $(S; X)$ -configurations at which M is active. We call

$$(f(U_i; L_i; \Pi_i)g_{i=1, \dots, q}; M_S)$$

the $(S; X)$ -tipping point of M .

Lemma 3.1.2. *Let $G; X; S$ be as above. Let M be an $(S; X)$ -locally popular matching. The $(S; X)$ -tipping point of M is uniquely defined. On the other hand, there exists a function $g : \mathbb{N}^S \rightarrow \mathbb{N}$ such that the number of $(S; X)$ -configurations is bounded by $g(jS_j)$, and the collection of $(S; X)$ -tipping points of M , where M ranges over all $(S; X)$ -locally popular matchings, has size at most $g(jS_j)jV^{jS_j}$.*

Proof. The first statement follows by definition. For the second and third: the number of $(S; X)$ -configurations $(U; L; \Pi)$ is a function of the size of S only, since all pairs from U are different, and each node of S can appear in at most two pairs from U . Moreover, the number of S -matchings of G is upper bounded by jV^{jS_j} . \square

Lemma 3.1.3. *Let $G; X; S$ be as above. Let M be an $(S; X)$ -locally popular matching and T be the $(S; X)$ -tipping point of M . Let N be another $(S; X)$ -locally popular matching whose $(S; X)$ -tipping point is also T . Let $e \in E(v)$ for some $v \in S$. Then:*

1. $e \in M$ if and only if $e \in N$.
2. Suppose $e \notin M$ and let $\ell \in f^+(v)$; g be the label of e at v wrt M . Then $e \notin N$ and ℓ is also the label of e at v wrt N .

Proof. Both statements immediately follow from $M_S = N_S$, which holds by definition of tipping point. \square

3.2 Leaders

Fix $G; S; X; M$ as in the previous section, and let T be the $(S; X)$ -tipping point of M . M is said to be the T -leader if it is the one of minimum cost among all $(S; X)$ -locally

popular matchings whose $(S; X)$ -tipping point is T . Due to the initial perturbation of costs, note that there is at most one T -leader for each tipping point T .

The following crucial lemma shows that, in order to find a min cost popular matching, it suffices to consider matchings that induce $(X \setminus S)$ -matchings that are T -leaders, for some $(S; X)$ -tipping point T . The proof of Lemma 3.2.1 is given in Section 3.4.

Lemma 3.2.1. *Suppose we are given $G; X; S; M$ as above. Let T be the $(S; X)$ -tipping point of M , and assume that M is not the T -leader. Let S^θ be a vertex separator of G , X^θ a connected component of $G \setminus S^\theta$ with the property that $X^\theta \cap X$ and $S^\theta \setminus X^\theta \subseteq S \setminus X$ (possibly $S^\theta = \emptyset$ and $X^\theta = V$). Let M^θ be a $(S^\theta; X^\theta)$ -locally popular matching such that $M^\theta_{X \setminus S} = M$. Let T^θ be the $(S^\theta; X^\theta)$ -tipping point of M^θ . Then M^θ is not a T^θ -leader.*

3.3 The algorithm

We now give our algorithm for computing a minimum weight popular matching, see Algorithm 3. Note that Algorithm 3 relies on the subroutine Update described in Algorithm 4, and the implementations of some other subroutines are not completely defined. We give a formal description of those together with a complexity analysis in the proof of Theorem 3.3.1.

Algorithm 3 takes as input a graph G with strict preference lists as usual, and a dichotomic directed tree decomposition $(T; B)$ of G (see Section 1.5). It iteratively constructs the sequence of closed subtrees T_B of T , with the first T_B corresponding to a leaf B of T , and the last to the root. For each T_B , let $S = B \setminus S(B)$ and $X = V(T_B) \setminus S$. The algorithm constructs and stores in L_B all the pairs $(M; T)$, where T is an $(S; X)$ -tipping point and M is the T -leader. This set can be found by building on the corresponding sets for the (at most two) predecessors of B . Finally,

of all matchings M such that $(M; T) \succeq L_B$ – with B being the root – the one of minimum cost is output. If at the end of any iteration we have $L_B = \cdot$; for the current B , we deduce that G has no popular matching.

With a little abuse of notation, we will write $M \succeq L_B$ if $(M; T) \succeq L_B$ for some T . Similarly, we write $L_B = L_B \sqcup fMg$ to mean that $(M; T)$ is added to L_B for an appropriate T , and similarly for $L_B \cap fMg$.

Algorithm 3:

Require: A graph G , together with, for each node $v \in V$, a strict ranking of the neighbors of v . A dichotomic directed tree decomposition $(T; B)$ of G .

Ensure: A popular matching of minimum cost in G .

- 1: For all $B \in \mathcal{B}$, label B as *unflagged*.
- 2: Choose an unflagged bag B whose predecessors are flagged, and flag B .
- 3: Let T_B be the closed subtree of T whose head is B , and set $L_B = \cdot$.
- 4: Let $S = B \setminus S(B)$, $X = V(T_B) \cap S$.
- 5: **if** B has no predecessor in T **then**
- 6: **for** all B -matchings M of G **do**
- 7: **if** M is an $(S; X)$ -locally popular matching **then**
- 8: Update(M , L_B)
- 9: **end if**
- 10: **end for**
- 11: **else**
- 12: Let $S_1 = B \setminus B_1$ and (possibly) $S_2 = B \setminus B_2$, where B_1 and (possibly) B_2 are the predecessors of B .
- 13: **for** all B -matchings M of G , all $M_1 \succeq L_{B_1}$ and (possibly) $M_2 \succeq L_{B_2}$ **do**
- 14: **if** $(M_1)_{S_1} = M_{S_1}$ and (possibly) $(M_2)_{S_2} = M_{S_2}$ **then**
- 15: Let $M = M \sqcup M_1 \sqcup M_2$.
- 16: **if** M is an $(S; X)$ -locally popular matching of G **then**
- 17: Update(M , L_B)
- 18: **end if**
- 19: **end if**
- 20: **end for**
- 21: **end if**
- 22: **if** $L_B = \cdot$ **then**
- 23: output: G has no popular matching.
- 24: **else if** there is a bag $B \in \mathcal{B}$ that is unflagged **then**
- 25: Go to Step 2.
- 26: **end if**
- 27: Let B be the head of T . Output the matching of minimum cost from L_B .

Algorithm 4: Update

Require: M, L_B
 Let T be the $(S; X)$ -tipping point of M .
if there exists $M^0 \in L_B$ whose tipping point is T **then**
 if $c(M) < c(M^0)$ **then**
 Set $L_B = L_B \cap fM^0g \cup fMg$.
 end if
else
 Set $L_B = L_B \cup fMg$.
end if

Analysis

The goal of this section is to prove the following result.

Theorem 3.3.1. *Algorithm 3 is correct. If the treewidth of $G = (V; E)$ is upper bounded by a constant t , then it can be implemented to run in time $O(n^{\beta^t + 7})$.*

We assume throughout the proof that every bag B that is not a leaf has two predecessors, as the (simpler) case where some B has only one predecessor follows in a similar fashion. For a bag B we write $S := B \setminus S(B)$ and $X := V(T_B) \cap S$; if B is not a leaf, we denote by B_1 and B_2 its predecessors, and write $S_i := B \setminus B_i$, $X_i = V(T_{B_i}) \cap S_i$ for $i = 1; 2$. We start by proving the correctness of Algorithm 3. We show that, at the end of the iteration where bag B is flagged,

- () L_B contains exactly all T -leaders, for all $(S; X)$ -tipping points T for which a T -leader exists.

Suppose () is proved. Then, when B is the root, L_B contains only popular matchings (by Lemma 3.1.1), and the one of minimum cost among those is the popular matching of minimum cost. Again by Lemma 3.1.1, if $L_B = \emptyset$ at the end of the iteration where B is flagged, then G has no popular matching, and the output of Algorithm 3 is again correct.

The proof of () is by induction on the number of nodes n_B of T_B . If $n_B = 1$, then the condition from the **if** statement in Step 5 is verified. In this case, the statement is immediate, since we enumerate over all possible B -matchings of G , check those that are $(S; X)$ -locally popular, and for each $(S; X)$ -tipping point T , keep the $(S; X)$ -locally popular matching of minimum cost active at T .

Now suppose $n_B > 1$. By induction hypothesis, for $i = 1;2$, all matchings that are stored in L_{B_i} are exactly all T -leaders, for all $(S_i; X_i)$ -tipping points T for which a T -leader exists (i.e., there is at least a $(S_i; X_i)$ -locally popular matching active at T). Now let T be an $(S; X)$ -tipping point for which a T -leader \hat{M} exists. It suffices to show that one of the matchings M constructed at Step 15 is indeed \hat{M} .

Since \hat{M} is $(S; X)$ -locally popular and $X_i \sqcap S_i \sqcap X \sqcap S$, matching $M_i := \hat{M}_{X_i \sqcap S_i}$ is also $(S_i; X_i)$ -locally popular by Lemma 3.1.1. By Lemma 3.2.1, M_i is a $(S_i; X_i)$ -leader. By induction, $M_i \in L_{B_i}$. On the other hand, $M := \hat{M}_B$ is a B -matching of G such that $M_{S_i} = (M_i)_{S_i}$ for $i = 1;2$. Since we enumerate all B -matchings of G , as well as all matchings from L_{B_1} and L_{B_2} , matching \hat{M} is eventually enumerated.

Running time analysis. We now bound the running time of Algorithm 3. We will use the following general fact proved in the claim below: if we are given an $(X \sqcap S)$ -matching M in a graph H and $jX \sqcap Sj$ is bounded, then one can check efficiently $(S; X)$ -local popularity of M . If M is $(S; X)$ -locally popular, then we can efficiently find the $(S; X)$ -tipping point at which M is active.

Claim 3.3.1. *Let $H = (U; F)$ be a graph, S a vertex separator of H , X a connected component of $H \setminus S$. Let M be a $(X \sqcap S)$ -matching of H . Assume $jX \sqcap Sj$ is upper bounded by a constant. Then in $O(jFj)$ time one can:*

1. *check if M is $(S; X)$ -locally popular in H and, if it is,*
2. *find the $(S; X)$ -tipping point (in H) at which M is active.*

Proof. Building graph $H_M[X \sqcap S]$ takes $O(jFj)$ time. Since $jX \sqcap Sj$ is upper bounded

by a constant, it takes constant time to enumerate all paths and cycles of $H_M[X \sqcup S]$, and to check if any of these violates the definition of $(S; X)$ -local popularity. In case M is indeed $(S; X)$ -locally popular, for any family of paths P in $H_M[X \sqcup S]$ and $(S; X)$ -configurations \mathcal{C} , one can check in constant time if P is a certificate for M at \mathcal{C} . Since by Lemma 3.1.2 the number of $(S; X)$ -configurations is upper bounded by $g(\ell + 1)$, which is a constant by hypothesis, the claim follows. \square

Observe that $jB_j \leq \ell + 1$, so enumerating all B -matchings takes time $O(jVj^{\ell+1})$. Because of Claim 3.3.1 and Step 7 (in Algorithm 3), finding the tipping point of any B -matching M can be performed in time $O(jEj) = O(jVj^2)$. Moreover, once the tipping point of M is computed, the Update function can be implemented to run in time $O(jVj^{\ell+1})$, since L_B at each step will contain at most one matching per tipping point, and by Lemma 3.1.2, there are at most $g(\ell + 1)jVj^{\ell+1}$ many tipping points.

We conclude that, when $n_B = 1$, the iteration when B is flagged runs in time $O(jVj^{\ell+2})$. We now prove, by induction on n_B , that the iteration when any B is flagged can also be implemented to run in time $O(jVj^{\beta^{\ell+5}})$, the base case having been just proved. Multiplying this by $O(jVj^2)$ (the number of bags of the tree decomposition), we obtain the desired bound.

Assume $n_B > 1$. Because of what was discussed above, given an $(S; X)$ -configuration \mathcal{C} and a matching $M = M \sqcup M_1 \sqcup M_2$, where M is a B -matching of G and for $i = 1; 2$ $M_i \subseteq L_{B_i}$, it is enough to give an upper bound on the time needed to decide if M is $(S; X)$ -locally popular and if it is active at \mathcal{C} . Recall that the number of B -matchings is $O(jVj^{\ell+1})$, while $jL_{B_i}j = O(jVj^{\ell+1})$ and the number of $(S; X)$ -configuration is at most $g(\ell + 1)$ by Lemma 3.1.2.

The condition from Step 14 can be verified in time $O(jEj)$. Hence, we assume that $(M_1)_{S_1} = M_{S_1}$ and $(M_2)_{S_2} = M_{S_2}$. We start with a simple claim.

Claim 3.3.2. *For $i = 1; 2$, let $\mathcal{C}_i = (U_i; L_i; \Pi_i)$ be a configuration at which M_i is*

active, and let P_i be a certificate of C_i at M_i . Then all paths from $P_1; P_2$ are pairwise internally vertex-disjoint.

Proof. By definition of certificate, paths from P_1 are pairwise internally vertex-disjoint, and similarly so are paths from P_2 . Now let $P_1 \geq P_1, P_2 \geq P_2$. For $i = 1; 2, P_i \subseteq V(T_{B_i})$ and P_i is an S_i -path. Since $V(T_{B_1}) \setminus V(T_{B_2}) \subseteq S_1 \sqcup S_2$, the claim follows. \square

Now fix an $(S_1; X_1)$ -configuration $C_1 = (U_1; L_1; \Pi_1)$ in G at which M_1 is active, and an $(S_2; X_2)$ -configuration $C_2 = (U_2; L_2; \Pi_2)$ in G at which M_2 is active. Consider the graph H and matching M_H obtained as follows. Start from $H = G_M[B]$ and the corresponding matching $M_H = M[B]$, and let $(u; v)$ be the first pair from U_1 . Assume $u; v$ are matched by $M_H, \ell_1 = (0; 1), \ell_2 = (0; 0)$, the other cases following in a similar fashion. Add to H and M_H new nodes $u^\ell; w^\ell; z^\ell; v^\ell$, matching edges $(u^\ell; w^\ell), (z^\ell; v^\ell)$, $(+; +)$ edge $(w^\ell; z^\ell)$, and $(+; -)$ edges $(u; u^\ell)$ and $(v; v^\ell)$. Note that $(u; u^\ell; w^\ell; z^\ell; v^\ell; v)$ is a M_H -alternating S_1 -path of parity ℓ_1 and level ℓ_2 starting at u and ending at v . We call it the *shortcut* of the path with endpoints $u; v$. Repeat this for all pairs from U_1 and U_2 , adding at each time new nodes and an appropriate path. Note that this adds to H a bounded number of nodes. This means that the graph H and matching M_H can be constructed in time $O(jEj)$.

Claim 3.3.3. *Let M be an $(S \sqcup X)$ -matching. M is not an $(S; X)$ -locally popular matching in G if and only if there exist an $(S_i; X_i)$ -configuration C_i at which M_i is active for $i = 1; 2$, such that, if we construct graph H and matching M_H as above, then M_H is not $(S; V(H) \setminus S)$ -locally popular in H .*

Proof. Consider the forbidden subgraph P (from Theorem 1.3.5) of $G_M[X \sqcup S]$, whose existence certifies that M is not $(S; X)$ -locally popular in G , and take P that is minimal with this property. Using an immediate modification to Algorithm 5 (given in Section 3.4), we can write P as the juxtaposition of $(P_1; P_2; \dots; P_k)$ where each P_j

is: either an S_i -path contained in $G[X_i \cup S_i]$ for $i \in \{1, 2\}$, or a path contained in $G[B]$. Collect all such S_1 -paths from $(P_1; \dots; P_k)$ not contained in $G[B]$ in P_1 , and all S_2 -paths from $(P_1; \dots; P_k)$ not contained in $G[B]$ in P_2 . One easily checks that P_i form the certificate of a certain $(S_i; X_i)$ -configuration C_i at M_i .

Now consider the graph H and matching M_H obtained from configurations C_1 and C_2 . By replacing each path from P_1 and P_2 with the corresponding shortcut, we deduce that M_H is not an $(S; V(H) \cap S)$ -locally popular matching in H .

The opposite direction follows in a similar (inverse) fashion: start from a forbidden path in H that certifies that M_H is not $(S; V(H) \cap S)$ -locally popular in H , write it as $(P_1; \dots; P_k)$, replace each of those subpaths with the appropriate certificate paths so as to obtain an M -alternating forbidden path by Claim 3.3.2. \square

The proof of the following claim follows in a similar fashion to the proof of Claim 3.3.3.

Claim 3.3.4. *Let M be an $(S; X)$ -locally popular matching in G and let C be an $(S; X)$ -configuration. Then M is active at C if and only if there exist a $(S_i; X_i)$ -configuration C_i at which M_i is active for $i = 1; 2$ such that if we construct graph H and matching M_H , then M_H is active at C (here we interpret C as an $(S; V(H) \cap S)$ -configuration in H , which is allowed since $S \subseteq V(H)$).*

Because of Claim 3.3.3, we can check if M is an $(S; X)$ -locally popular matching in G by checking, for each pair of $(S_i; X_i)$ -configurations C_i at which M_i is active for $i = 1; 2$, if the corresponding M_H is $(S; V(H) \cap S)$ -locally popular in H . For $i = 1; 2$, the number of $(S_i; X_i)$ -configurations is at most $g^{(l+1)}$. Since the number of nodes of H is bounded (we start with a bounded set of nodes and we add a bounded number of new nodes), testing if M_H is $(S; V(H) \cap S)$ -locally popular in H can be done in time $O(jEj) = O(jVj^2)$ by Claim 3.3.1. Similarly, if M is $(S; X)$ -locally popular in G , we can find its tipping point in time $O(jVj^2)$ by repeatedly applying Claims 3.3.4 and

3.3.1. Hence, the iteration when B is flagged can be performed in time $O(jVj^{\beta'+5})$. This concludes the proof of Theorem 3.3.1.

3.4 Proof of Lemma 3.2.1

Let N be the $(S; X)$ -locally popular matching that is the T -leader. This implies that $M_S = N_S$, and N is a $(X \upharpoonright S)$ -matching. Define $N^\theta := N \upharpoonright (M^\theta \cap M)$, and notice this is a disjoint union. Hence, $N^\theta_{X \upharpoonright S} = N$. We will now show that N^θ is a $(S^\theta; X^\theta)$ -locally popular matching whose $(S^\theta \upharpoonright X^\theta)$ -tipping point is T^θ . We then have:

$$c(N^\theta) = c(N) + c(N^\theta \cap M) < c(M) + c(M^\theta \cap M) = c(M^\theta);$$

concluding the proof. We start with some claims, the first two of which immediately follow from construction.

Claim 3.4.1. $G_{M^\theta}[X \upharpoonright S] = G_M[X \upharpoonright S]$ and $G_{N^\theta}[X \upharpoonright S] = G_N[X \upharpoonright S]$.

Claim 3.4.2. $M^\theta = M \upharpoonright (N^\theta \cap M)$, and the latter is a disjoint union.

Claim 3.4.3. Let $e \in E$ not be incident to a node of X .

- a) $e \in N^\theta$ if and only if $e \in M^\theta$.
- b) Let $e \notin N^\theta$. Then $e \in E(G_{N^\theta})$ if and only if $e \in E(G_{M^\theta})$. Moreover, if $e \in E(G_{N^\theta})$, then it has the same labels in G_{N^θ} and in G_{M^θ} .

Proof. a) Suppose $e \in N^\theta$. If $e \in N$, then $e \in N_S = M_S = M = M^\theta$, where equality holds by hypothesis and inclusions by definition. Else, $e \in (N^\theta \cap M) = M^\theta$. The argument can be reversed to show the opposite direction.

b) Suppose $e = (u; v) \notin N^\theta$. By part a), $e \notin M^\theta$. If $u \notin S$, then $N^\theta(u) = M^\theta(u)$ by part a). We deduce that the label of e at u coincides in M^θ and N^θ . If instead $u \in S$,

then the label of e at u in G_M and G_N coincide by Lemma 3.1.3. The claim follows since $M_S^\theta = M_S = N_S = N_S^\theta$. \square

Claim 3.4.4. N^θ is an $(X^\theta [S^\theta)$ -matching of G .

Proof. Let us first show that it is a matching of G . Recall that $M_{X[S]}^\theta = M$. Hence, $M^\theta n M$ has no edge incident to $X [S$. Since by hypothesis N is an $(X [S)$ -matching, N has no edge incident to nodes other than $X [S$. Hence N^θ is the union of two node-disjoint matchings of G , hence a matching.

Now by hypothesis, $S [X = S^\theta [X^\theta$. Hence, N has no edge in $G n (S^\theta [X^\theta)$. Moreover, M^θ is by hypothesis an $(S^\theta; X^\theta)$ -locally popular matching, so it is an $(S^\theta [X^\theta)$ -matching. Hence, N^θ is an $(S^\theta [X^\theta)$ -matching. \square

Claim 3.4.5. Let $k \geq 2$ and $P^1; \dots; P^k \subseteq P_{M^\theta}$ be pairwise internally vertex-disjoint and X -disjoint. Assume that, for $j \in [k]$, path P^j is of level $\ell_j \in \{0; 1; 2g - 1; g\}$ and parity $\Pi_j \in \{0; 1g - 1; g\}$, that every node appears at most twice in $P^1; \dots; P^k$, no pair of paths have exactly the same endpoints, and no path is contained in X . Then there exist pairwise internally vertex-disjoint and X -disjoint paths $Q^1; \dots; Q^k \subseteq P_{N^\theta}$ with the following properties:

1. For $j \in [k]$, Q^j is of level ℓ_j and parity Π_j , and $P^j \setminus S^\theta = Q^j \setminus S^\theta$.
2. For $j \in [k]$, let u_j (resp. v_j) be the first (resp. last) node of P^j . If u_j (resp. v_j) $\notin X$, then u_j (resp. v_j) is the first (resp. last) node of Q^j . If u_j (resp. v_j) $\in X$, the first (resp. last) node of Q^j also belongs to X .

Moreover,

3. 1,2 also hold if we switch the roles of M^θ and N^θ .

Proof. The proof of the claim will only assume that M and N are $(S; X)$ -locally popular matchings with the same $(S; X)$ -tipping point. Recall that $N^\theta = N[(M^\theta n M)$, $M^\theta = M[(N^\theta n N)$ (by Claim 3.4.2), and that those unions are disjoint. Therefore, the

roles of M and N (hence those of M^0 and N^0) can be exchanged and the conclusions preserved. This proves 3 (assuming 1,2).

Let us consider any path P from $P^1; \dots; P^k$, where we omit the superscript for the sake of readability. If $P \setminus X = \emptyset$, we write $P_1 := P$. Else, we write P as the juxtaposition of $P_0; P_1; P_2; \dots; P_q; P_1$, defined through the procedure described in Algorithm 5. Note that some of those paths may consist of a single node. We call them *trivial*. Note that a trivial path can be removed from the juxtaposition without changing the resulting P . Hence, in the following, we assume that all paths are non-trivial (without changing the subscripts of the remaining non-trivial paths), and we let q be the subscript of the last non-trivial path (other than 1).

The following fact immediately holds by construction.

Claim 3.4.6. *Let $fP_i g_{i \in [q] \setminus \{0,1\}}$ be the output of Algorithm 5 on input $P = P^j$ for some $j \in [k]$. Then: (i) P is the juxtaposition of $P_0; P_1; \dots; P_q; P_1$. (ii) For all $i \in [q] \setminus \{0,1\}$, $P_i \in P_{M^0}$. (iii) For $i \in \mathbb{N}$, $i \geq 2$, the first and last nodes of P_i belong to S . (iv) For $i \in \mathbb{N}$, i odd, P_i is an S -path and $P_i \cap X \neq \emptyset$. (v) For i even, $P_i \setminus X = \emptyset$. (vi) Assume $q \geq 1$. Then P_1 ends at a node of S , and starts either at a node of X (if the **if** condition in Step 2 of Algorithm 5 is satisfied), or at a node of S (otherwise). (vii) If $P_1 \setminus X \neq \emptyset$, then P_1 is an S -path of $G_{M^0}[X \setminus S]$ ending at some node of X .*

Let us call a subpath P_i of P *hidden* if $i \not\equiv 1 \pmod{2}$ or $i = 1$ and $P_1 \setminus X \neq \emptyset$ (i.e., condition (vii) above is verified). Now consider the collection of paths $fP_i^j g$ defined above, for $j \in [k]^2$.

Fact 3.4.1. *Consider the family of paths $fP_i^j g$ with $j \in [k]$. Then:*

²We introduced the superscripts back, to distinguish the paths produced by applying Algorithm 5 to P^1, \dots, P^k .

Algorithm 5:

- 1: Let v be the first vertex of P .
- 2: **if** $v \in X$ **then**
- 3: Let $P_0 = fvg$ and $u = v$.
- 4: **else**
- 5: Let u^j be the first vertex of P that is contained in X , and u the predecessor of u^j in P . Note that $u \in S$. Let P_0 be the subpath of P between v and u (possibly $u = v$ and $P_0 = fug$).
- 6: **end if**
- 7: Let $i = 1$.
- 8: Starting from u , traverse P until the first node $v \in S$, $v \neq u$ is encountered.
- 9: **if** such v does not exist **then**
- 10: Set $q = i - 1$ and go to Step 21.
- 11: **else**
- 12: Let P_i be the subpath between u and v . Set $i = i + 1$ and $u = v$.
- 13: Starting from u , traverse P until the first node $u^j \in X$ is encountered.
- 14: **if** such u^j does not exist **then**
- 15: Set $q = i - 1$ and go to Step 21.
- 16: **else**
- 17: Let P_i be the subpath of P between u and the predecessor v of u^j in P . Note that $v \in S$ and possibly $v = u$.
- 18: Set $u = v$, $i = i + 1$, and go to Step 8.
- 19: **end if**
- 20: **end if**
- 21: Let P_1 be the subpath of P between u and the last node of P .
- 22: Output $P_0; P_1; \dots; P_q; P_1$.

1. $\{P_i^j\}$ is a collection of pairwise internally vertex-disjoint paths and each node appears at most twice in the collection.
2. The restriction of the collection $\{P_i^j\}$ to hidden paths is a family of pairwise X -disjoint S -paths from P_M , and the level of each of those path is not 1.
3. If $P_i^j \setminus P_{i^0}^{j^0} \neq \emptyset$, then either (a) $j = j^0$ and ($i = i^0 - 1$ or $i = i^0 \in \{q; 1\}$), or (b) i (resp. i^0) is the first or last path of the juxtaposition leading to P^j (resp. P^{j^0}).

Proof. Part 1 follows from hypothesis, the fact that paths P^j , $j \in [k]$ form a collection of pairwise internally vertex-disjoint paths and Fact 3.4.6.

We now prove part 2. Restrict the collection $fP_i^j g$ to hidden paths. We already argued in Fact 3.4.6 that those are S -paths. As they are subpaths of M^θ -alternating paths, they are also M^θ -alternating. Using again Fact 3.4.6, they are contained in $X \setminus S$. Using Claim 3.4.1, we deduce they are paths from P_M . The fact that they are X -disjoint follows by hypothesis for $j \neq j^\theta$, and by Fact 3.4.6 for $j = j^\theta$. The fact that M^θ is $(S^\theta; X^\theta)$ -locally popular implies that they cannot be of level $\neq 1$. This proves part 2.

We now prove part 3. Let $P_i^j \setminus P_{i^\theta}^{j^\theta} \in \mathcal{P}$. If $j = j^\theta$, by Fact 3.4.6 the only possibility is that $P_i^j, P_{i^\theta}^{j^\theta}$ are consecutive paths in the juxtaposition leading to P . This is case (a) in 3. Else, $j \neq j^\theta$ and case (b) follows from the fact that P^j and P^{j^θ} are internally vertex-disjoint, concluding the proof of 3. \square

Consider the collection of hidden paths $fP_i^j g$ for $j \in [k]$, and let u_i^j (resp. v_i^j) be the first (resp. last) node of each of those paths. Consider $C = (U; L; \Pi)$ defined as follows: U is the ordered collection of pairs:

- $(u_1^j; v_1^j)$ if the first node of P_1^j belongs to S , and $(; ; v_1^j)$ otherwise.
- $(u_i^j; v_i^j)$ for $i \in [3]$ odd and
- $(u_1^j; ; ;)$ if P_1^j is hidden.

L is the ordered collection of levels $\neq j$ of the hidden paths P_i^j , while Π is the ordered collection of their parities. Using Claim 3.4.6, Claim 3.4.1 and the hypothesis, one easily verifies the following.

Claim 3.4.7. *C is an $(S; X)$ -configuration. M is active at C , and the collection of hidden paths $fP_i^j g$ for $j \in [k]$ is a certificate of C at M .*

By hypothesis, N is also active at C . Hence, for all pairs $(i; j)$ such that P_i^j is hidden, we can find pairwise X -disjoint S -paths $Q_i^j \subseteq P_N$, with each P_i^j starting (resp. ending) at node u_i^j (resp. v_i^j) if this belongs to S , and at a node of X otherwise, of the appropriate level and parity.

For $j \geq 2$, recall that P^j is the juxtaposition of $(P_0^j; P_1^j; \dots; P_q^j; P_1^j)$, from which we removed the trivial paths. Consider the path Q^j obtained by replacing, in the juxtaposition above, each hidden path P_i^j with the corresponding Q_i^j . We conclude the proof of the claim by showing that the collection of $Q^j, j \geq 2$, satisfies the claim.

Fix $j \geq 2$. Let us first argue that Q^j is a path in G that satisfies property 2 from the statement of the claim. Note that, by construction, for all hidden P_i^j , the first (resp. last) vertex of P_i^j coincides with the first (resp. last) vertex of Q_i^j , with possibly the exception of the first vertex of P_0^j (resp. last vertex of P_1^j) if it belongs to X . Moreover, since all Q_i^j are pairwise X -disjoint S -paths with no edge in $G \cap X$, each Q^j is a path in G . The statement follows.

Let us now argue that Q^j is N^θ -alternating. Since $N^\theta[(X^\theta \sqcup S^\theta) \cap X] = M^\theta[(X^\theta \sqcup S^\theta) \cap X]$ and all P_i^j that are not hidden are M^θ -alternating and do not intersect X , we conclude that all such P_i^j are also N^θ -alternating. For each hidden P_i^j , Q_i^j is an N -alternating path in $G_N[X \sqcup S]$ by construction. Since $G_N[X \sqcup S] = G_{N^\theta}[X \sqcup S]$ by Claim 3.4.1, each Q_i^j is also N^θ -alternating. By construction, the parity of each hidden P_i^j coincides with that of the corresponding Q_i^j . We conclude that each Q^j is N^θ -alternating.

Let us now show that $Q^j \geq P_{N^\theta}$, and it is of the same level and parity of P^j . Let P_i^j be non-hidden. By Fact 3.4.6, it is contained in $G[X^\theta \sqcup S^\theta] \cap X$. By Claim 3.4.3, $P_i^j \setminus N^\theta = P_i^j \setminus M^\theta$, and each edge from $P_i^j \cap N^\theta$ has the same labels in $G_{M^\theta}[X^\theta \sqcup S^\theta]$ and $G_{N^\theta}[X^\theta \sqcup S^\theta]$. Now consider a hidden path P_i^j . By construction, Q_i^j is a path of $G_N[X \sqcup S] = G_{N^\theta}[X \sqcup S]$ of the same parity and level of P_i^j . The statement follows.

Now observe that property 1 of the claim holds, since the hypothesis $X \cap X^\theta$ implies $S^\theta \setminus X = \emptyset$, hence we have:

$$Q^j \setminus S^\theta = (Q^j \setminus (S^\theta \setminus S)) \sqcup (Q^j \setminus (S^\theta \cap S)) = (P^j \setminus (S^\theta \setminus S)) \sqcup (P^j \setminus (S^\theta \cap S)) = P^j \setminus S^\theta.$$

Finally, let us observe that $Q^1; \dots; Q^k$ are internally vertex-disjoint and X -disjoint. The X -disjointness follows from the fact that all Q_i^j corresponding to hidden P_i^j are X -disjoint, while other P_i^j do not intersect X by Fact 3.4.6. The fact that $Q^1; \dots; Q^k$ are pairwise internally vertex-disjoint follows from Fact 3.4.1. \square

In order to conclude the proof of the lemma, we need to prove that N^θ is a $(S^\theta; X^\theta)$ -locally popular matching whose tipping point is T^θ .

($S^\theta; X^\theta$)-local popularity. Let us first show that N^θ is $(S^\theta; X^\theta)$ -locally popular. Suppose it is not, then there exists a forbidden path or cycle P as in item (i), (ii), or (iii) from Theorem 1.3.5 that is contained in $G_{N^\theta}[S^\theta \setminus X^\theta]$. Since (by Claim 3.4.1) $G_N[X \setminus S] = G_{N^\theta}[X \setminus S]$ and $v \succeq X \setminus S$ is N -exposed if and only if it is N^θ -exposed, we deduce $P \not\subseteq X \setminus S$, else N would not be $(S; X)$ -locally popular, a contradiction.

Suppose P is a path: take one that is inclusionwise minimal among all paths that certify that N^θ is not $(S^\theta; X^\theta)$ -locally popular. We can then write P as the juxtaposition of $(P^1; P^2)$ of levels $\ell_1; \ell_2 \notin \ell$, so that the first node of P^2 (which coincides with the last node of P^1) does not belong to X . Clearly $P^1; P^2$ satisfy the hypothesis of Claim 3.4.5 with the roles of N^θ and M^θ exchanged and $k = 2$. Let $Q^1; Q^2$ be the paths whose existence is guaranteed by the claim. We claim that the juxtaposition of $(Q^1; Q^2)$ is a forbidden path in $G_{M^\theta}[X^\theta \setminus S^\theta]$, a contradiction.

First, observe that the last node of Q^1 coincides with the last node of P^1 (since the latter does not belong to X) which coincides with the first node of P^2 , which coincides with the first node of Q^2 (again, since it does not belong to X). We now argue that P^1 and P^2 have no other node in common. We know they are internally vertex-disjoint and X -disjoint. Hence, if the first node of P^1 or the last node of P^2 belongs to X , we are done. Suppose not. Then, by Claim 3.4.5, the first node of Q^1 is the first node of P^1 , which is different by hypothesis (P is a path) from the last node of P^2 , which coincides, again by Claim 3.4.5, with the last node of Q^2 . Hence, the juxtaposition of $(Q^1; Q^2)$ is a path of G , that we denote by Q . Again by Claim

3.4.5, $Q^1, Q^2 \not\subset P_{M^\theta}$. Moreover, the parity and level of Q^1 (resp. Q^2) coincide with the parity and level of P^1 (resp. P^2). Hence $Q \not\subset P_{M^\theta}$ is the required forbidden path, obtaining the desired contradiction.

The argument when P is a cycle follows in a similar fashion, writing P as the juxtaposition of $(P^1; P^2; P^3)$ with $P^1 \setminus X = P^2 \setminus X = ;$, and the endpoints of P_3 lying in S (note that we can assume that P contains exactly one $(+; +)$ edge, else it also contains a path with two $(+; +)$ edges and we are back to the previous case).

Tipping point. Let us now show that T^θ is the tipping point of N^θ . First, observe that $X \setminus X^\theta$ implies $S^\theta \setminus X = ;$. For $v \in S^\theta \setminus S$, we know that $N^\theta(v) = N(v) = M(v) = M^\theta(v)$. For $v \in S^\theta \cap S$, we have $N^\theta(v) = M^\theta(v)$ by construction. Hence $N_{S^\theta}^\theta = M_{S^\theta}^\theta$. Now let $C = fU; L; \Pi g$ be a $(S^\theta; X^\theta)$ -configuration at which M^θ is active, with U as in (3.1). We will show that N^θ is also active at C . A symmetric argument shows that, if N^θ is active at C , then so is M^θ , concluding the proof.

Take paths $P^1; \dots; P^k$ that are the certificate of C at M^θ . We claim that those paths satisfy the hypothesis of Claim 3.4.5. They are, by construction, S^θ -paths that are pairwise X^θ -disjoint. This implies that they are pairwise internally vertex-disjoint and X -disjoint (using again $X \setminus S^\theta = ;$). The X^θ -disjointness implies that each node from X^θ appears at most once. Moreover, again by construction, each node from S^θ appears at most twice. Hence, every node appears at most twice in $P^1; \dots; P^k$. Moreover, by definition of U , no two paths have the same endpoints.

Let $Q^1; \dots; Q^k \subset P_{N^\theta}$ be the paths whose existence is guaranteed by Claim 3.4.5. We show that $Q^1; \dots; Q^k$ are the certificate of C at N^θ , concluding the proof. Fix $j \in [k]$. The parity and level of Q^j is the same as that of P^j . Hence, the parity of Q^j is Π_j , and its level is j . Recall that Q^j is an S^θ -path. By property 2 of Claim 3.4.5, Q^j starts (resp. ends) at the same node as P^j when this belongs to S^θ , and at some node of X^θ otherwise. Using property 1 of Claim 3.4.5, we deduce that Q^j is a S^θ -path that starts (resp. ends) at u_j (resp. v_j) when $u_j \notin ;$ (resp. $v_j \notin ;$), and at a node

$x \supseteq X^0$ otherwise.

We are left to show that paths $fQ^jg_{j=1,\dots,k}$ are X^0 -disjoint. Suppose not, then there exists $v \supseteq X^0$, $j \notin j^0$ such that $v \supseteq Q^j \setminus Q^{j^0}$. Note that v cannot be an internal node of Q^j or Q^{j^0} , since those paths are pairwise internally vertex-disjoint by construction. On the other hand, if v is an endpoint of both Q^j and Q^{j^0} , then the statement either follows from the X -disjointness of Q^j and Q^{j^0} , or by property 2 of Claim 3.4.5. □

*The distance of stable and popular matchings from the
Pareto-optimal frontier*

4.1 Introduction

As we have seen in Chapter 1, stability is the pristine concept used in two-sided matching markets. However, in many contexts, one side of the market (typically schools, when the other side of the bipartition are students) is seen as just a good to be distributed to the other side. Using this point of view, preferences of schools play a minor role, hence stability is not an appropriate solution concept anymore. The focus is therefore moved towards *Pareto-optimality* for students. Pareto-optimal matchings can be obtained by *serial dictatorship* but are, in general, less tractable than stable matchings. For instance, unless $P=NP$, a Pareto-optimal matching of minimum size cannot be found efficiently, though a Pareto-optimal matching of maximum size can be found in polynomial time [4].

The goal of much of the modern literature on economic aspects of two-sided markets is to somehow resolve this tension between stability and Pareto-optimality. Of particular interest are therefore instances where there is a matching that is both stable and Pareto-optimal. This situation has been studied in [3]. Because of the lattice structure of stable matchings, the only candidate to lie in this set is the man-optimal stable matching. Since this matching can be found efficiently, and Pareto-optimality can also be checked efficiently [4], we conclude the following fact.

Theorem 4.1.1. *Given a marriage instance, one can decide in polynomial time if there is a stable matching that is Pareto-optimal.*

Unfortunately, in typical instances, a matching that is both Pareto-optimal and stable does not exist. However, this fact alone does not convey how contradicting the properties of Pareto-optimality and stability are for a given instance. Moreover, as the families of both stable and Pareto-optimal matchings can be very large, a market designer is often left with the dilemma of which matching to choose. A stable matching that is “close” (in some well-defined sense) to being Pareto-optimal could be favoured over one that is “far away” (and similarly exchanging Pareto-optimality and stability). But do such matchings exist and, if they do, can they be computed efficiently?

In this section, we introduce and study parameters measuring the distance between the sets of stable/popular and Pareto-optimal matchings. As we will see, those parameters can be interpreted in two ways. First, as a structural property of the instance, describing in a natural way how far stable/popular matchings are from the Pareto-optimal frontier. Then, as lower bounds on the number of “changes” certain well-known algorithms need to perform when moving from a stable to a Pareto-optimal matchings.

Δ , Λ , Γ , and our results.

In order to introduce the first parameter, fix a marriage instance $(G; <)$ with bipartition $(A; B)$ and for $M; M^0 \in \mathcal{M}(G)$ define

$$\Delta(M; M^0) := \sum_{m \in A} |M^0(m) \setminus M(m)|$$

Hence, for each pair of matchings M and M^0 , we have $\Delta(M; M^0) \in \{0, 1, \dots, |A|\}$, with the extreme values achieved when $M = M^0$ and when $M \setminus M^0 = \emptyset$, respectively. For

an instance $(G; <)$, we let:

$$\Delta(G; <) := \min_{M \text{ stable}, M^\theta \text{ Pareto opt}} \Delta(M; M^\theta):$$

Hence, $\Delta(G; <)$ resembles somehow the classical Hamming distance between binary vectors. It measures the minimum number of men that have to change partner, in order to move from a stable to a Pareto-optimal matching, and is a natural measure of the distance between the set of Pareto-optimal and stable matchings. We will omit the dependency on $(G; <)$ in Δ when the underlying graph is clear from the context.

It will be useful to work with matching-dependent functions as well. For a stable matching M , we let

$$\Delta(M) := \min_{M^\theta \text{ Pareto-optimal}} \Delta(M; M^\theta):$$

For a Pareto-optimal matching M^θ , we let

$$\Delta(M^\theta) := \min_{M \text{ stable}} \Delta(M; M^\theta):$$

The notation is well-defined, since when M is both Pareto-optimal and stable, we have $\Delta(M) = 0$.

Many algorithm studied in theory and used in practice, such as Kesten’s EADAM [37] and the Top Trading Cycle [50], take as input an initial (stable) matching M , and produce a Pareto-Optimal matching M^θ that every man weakly prefers to M (i.e., M^θ Pareto-dominates M). This is motivated by the fact that no man would accept a “globally better solution”, if this solution is less competitive for him. Since departing from stability is often countered by protests of agents, it is therefore important to understand how to achieve a Pareto-optimal matching with a minimum number of changes.

To capture this setting, we define therefore $PO(M)$ to be the set of Pareto-




Figure 4.1: On the top: a marriage instance $(G; <)$. In the middle: the only stable matching M of the instance. It contains exactly one coalition cycle C . One can easily check (or resort to Lemma 4.2.2) that the only matching in $PO(M)$ is $M \setminus C$. Hence, $\Lambda(M) = k + 2$. On the bottom: the matching realizing $\Delta(M) = 1$.

optimal matchings that Pareto-dominate M . Note that, for $M^\theta \in PO(M)$, we can alternatively compute $\Delta(M; M^\theta)$ as follows.

Observation 4.1.2. *For a matching M and a Pareto-optimal matching $M^\theta \in PO(M)$, we have:*

$$\Delta(M; M^\theta) = \text{JFM} \setminus A : M^\theta(m) \notin M(m) \text{GJ} = \text{JFM} \setminus A : M^\theta(m) >_m M(m) \text{GJ} : \quad (4.1)$$

We define therefore

$$\Lambda(G; <) := \min_{M \in S; M^\theta \in PO(M)} \Delta(M; M^\theta):$$

Again, the associated local function is

$$\Lambda(M) := \min_{M^\theta \in PO(M)} \Delta(M; M^\theta)$$




Figure 4.2: The Hasse diagram of the poset of matchings of a marriage instance with 4 students and schools. Stable matchings are represented by unfilled circles, while Pareto-optimal matchings correspond to elements of the poset without predecessors. There is exactly one Pareto-optimal matching M_0^θ (unfilled square) that dominates the student-optimal stable matching M_0 . Hence $\Gamma(M_0) = \Gamma(M_0; M_0^\theta)$. There is a second Pareto-optimal matching M_1^θ (filled large circle) that dominates the other stable matching M_1 . Hence $\Gamma(M_1) = \min\{\Gamma(M_1; M_0^\theta); \Gamma(M_1; M_1^\theta)\}g$ and $\Gamma(G; <) = \min\{\Gamma(M_0; M_0^\theta); \Gamma(M_1; M_0^\theta); \Gamma(M_1; M_1^\theta)\}g$. On the other hand, in order to compute $\Delta(G; <)$, we must also consider the other Pareto-optimal matchings (filled squares).

Clearly, one has $\Delta = 0$ if and only if $\Lambda = 0$, and $\Lambda \leq \Delta$. It is maybe counterintuitive that the latter inequality can be strict – that is, the matching realizing $\Delta(M)$ may not Pareto-dominate M . Actually, it is easy to construct examples where the distance between Λ and Δ can be arbitrarily large, see Fig. 4.1. See Fig. 4.2 for a pictorial representation of Λ vs. Δ .

Many algorithms produce a Pareto-optimal matching M^θ by starting from an already existing matching M , and then iteratively take the symmetric difference of M with *Pareto-improving* (or *coalition*) cycles, see Section 1.3 for definitions. Recall that these are M -alternating cycles where every man prefers his neighbor not in M to the one in M . We call such an Algorithm a *Pareto-improving* algorithm. At each step, multiple choices for C are possible. As we argue in Lemma 4.2.3, when M is stable (and, more generally, popular) $\Lambda(M)$ is exactly the minimum number of men

that need to be in some Pareto-improving cycle when going from M to $M^0 \in PO(M)$. Hence, it gives a lower bound to the number of changes in the matching that any such algorithm must perform.

Recall that the existence of a stable matching that is Pareto-optimal (equivalently, that $\Delta = \Lambda = 0$) can be verified efficiently by building on the lattice structure of the set of stable matchings.

However, the general problem of computing or even approximating Δ is very hard. We show in fact the following.

Theorem 4.1.3. *There exists a universal constant c such that the following holds. Unless $P=NP$, for each $\epsilon > 0$, there is no polynomial-time algorithm that, given a marriage instance $(G(A \cup B; E); <)$, accepts if $\Delta(G) \leq \epsilon |A|$, rejects if $\Delta(G) \geq c|A|$, and is free to accept or reject if none of these cases apply. In particular, computing $\Delta(G)$ is NP-hard.*

The proof of this theorem and of related hardness results is presented in Section 4.5. We counterbalance this results by showing that deciding whether Δ or Λ is a constant can be decided in polynomial time. The proof of next result appears in Section 4.4. We denote by PO the set of Pareto-optimal matchings.

Theorem 4.1.4. *Let k be a constant. There exists a polynomial-time algorithm that, given a marriage instance $(G; <)$, either concludes $\Delta \leq k$ (resp. $\Lambda \leq k$) and outputs $M \in S$, $M^0 \in PO$ (resp. $M \in S$, $M^0 \in PO(M)$) with $\Delta(M; M^0) \leq k$, or correctly concludes that $\Delta > k$ (resp. $\Lambda > k$).*

A corollary of Theorem 4.1.3 is that there exist instances where the distance of any stable matching from the Pareto-optimal frontier is very large. One could therefore think of relaxing stability and hopefully moving “closer” to the set of Pareto-optimal

matchings. Hence, we define the parameter $\Gamma(G; <)$, that is the analogous of $\Delta(G; <)$ when the set of stable matchings is replaced by that of popular matchings.

Formally, for a marriage instance $(G; <)$, we let:

$$\Gamma(G; <) := \min_{M \text{ popular}, M^\theta \text{ Pareto opt}} \Delta(M; M^\theta):$$

It will also be useful to define, for a marriage instance $(G; <)$ and a Pareto-Optimal matching M^θ

$$\Gamma(M^\theta) := \min_{M \text{ popular}} \Delta(M; M^\theta):$$

and similarly for a popular matching M

$$\Gamma(M) := \min_{M^\theta \text{ Pareto-optimal}} \Delta(M; M^\theta):$$

The notation is well-defined, since when M is both Pareto-optimal and popular, we have $\Gamma(M) = 0$.

As one can imagine, it can in fact be the case that Γ is much smaller than Δ . In Section 4.3, we give an example of an instance where $\Gamma = 0$ and $\Delta = |A| - 2$. Maybe not surprisingly (given complexity results on popular matching presented in this thesis), when we move from stability to popularity, even deciding whether there is a matching that is both popular and Pareto-optimal becomes hard.

Theorem 4.1.5. *Let $(G; <)$ be a marriage instance. It is NP-Complete to decide if $\Gamma(G) = 0$.*

On the other hand one could expect that, by relaxing stability to popularity, it would be easier to find a matching that is reasonably close to the Pareto-optimal frontier. Our next results shows this is not the case.

Theorem 4.1.6. *There exists a universal constant c such that the following holds. For each $\epsilon > 0$, unless $P=NP$, there is no polynomial-time algorithm that, given a preference system $(G(A [B; E); <)$, accepts if $\Gamma = 0$, rejects if $\Gamma \geq c|A|$, and is free to accept or reject if none of these cases apply.*

The techniques

Recall the general approach developed in Chapter 2 to prove hardness results on stable/popular matchings. Starting from an instance ϕ of Monotone 3-SAT, a marriage instance $(G; <)$ is constructed, with the property that ϕ is satisfiable iff $(G; <)$ contains the required stable or popular matching.

In order to prove results from this section we push this connection further. We illustrate the general idea for Theorem 4.1.3, which is also followed by most other hardness results in this section (although technical details are quite different). The reduction between ϕ and $(G; <)$ is defined so that there is a bijection between stable matchings in G and *consistent* literal assignments in ϕ (see Section 1.4 for definitions). We then show that one can define a map q from *feasible* assignments to Pareto-optimal matchings, in such a way that the closer a feasible literal assignment α is to be consistent (in terms of number of satisfied in close), the closer (in terms of Δ) the Pareto-optimal matching $q(\alpha)$ is to be stable. If we could show that q is surjective over the set of Pareto-optimal matchings, then we would be done by Theorem 1.4.2 on the hardness of approximation of Monotone 3-SAT. However, this is not the case: there are Pareto-optimal matchings that are not in the codomain of q . However, we show that none of those Pareto-optimal matchings is the closest (again, in terms of Δ) to any stable matching, concluding the proof.

For the polynomial-time results, we heavily build on known structural and algorithmic facts on stable matchings, introduced in Chapter 1.

Organization of the chapter. We start with settling notation and basic facts in 4.2. In Section 4.3, we present a number of examples which illustrate that the parameters introduced can be very different from each other, and that more generally show the complexity of the problem we investigate. Our polynomiality results are presented in Section 4.4. Theorem 4.5 and related results are proved in Section 4.5. Theorem 4.1.5 and Theorem 4.1.6 are proved in section 4.5.

4.2 Further notation and basic facts

If H is a subgraph of a graph G and $M; M^0$ are matchings on G , we let $\Delta_H(M; M^0) := \Delta(M_H; M^0_H)$, where for a matching N and a subgraph H of G , N_H is the restriction of N to nodes of H .

Lemma 4.2.1. *Let M be a popular matching and $M^0 \succeq PO(M)$. Then $V(M) = V(M^0)$ and $M \setminus M^0$ is a collection of vertex-disjoint coalition cycles for M .*

Proof. We first observe $V(M) = V(M^0)$. Since M^0 is Pareto-optimal, we have

$$A(M^0) \geq A(M); \tag{4.2}$$

Suppose first by contradiction that there exists $b \in B(M^0) \cap B(M)$. Ask all agents to vote between M and M^0 . Let A_M (resp. A_{M^0} , A_0) be the number of men voting for M (resp. M^0 , abstaining). Define B_M, B_{M^0}, B_0 similarly. By definition of Pareto-dominance,

$$A_M = 0; A_{M^0} = \#\{a \in A : M(a) \neq M^0(a)\};$$

Clearly b contributes to B_{M^0} (and not to B_M). Hence we have:

$$\begin{aligned}
A_{M^0} + B_{M^0} &= \sum_{a \in A} |M(a) \setminus M^0(a)| + 1 \\
&= \sum_{b \in B} |M(b) \setminus M^0(b)| + 1 \\
&> B_M = A_M + B_M;
\end{aligned} \tag{4.3}$$

contradicting the popularity of M . This shows

$$B(M^0) = B(M); \tag{4.4}$$

Now assume there exists $a \in A(M^0) \setminus A(M)$. Because of (4.2), we deduce that there exists $b \in B(M^0) \setminus B(M)$, and we are back to the previous case, which we know cannot happen. Together with (4.2), this implies $A(M^0) = A(M)$. Hence $|B(M)| = |B(M^0)|$, and (4.4) implies $B(M) = B(M^0)$. We conclude $V(M) = V(M^0)$.

Hence, the symmetric difference of M and M^0 is given by a collection of vertex-disjoint cycles. Since $M^0 \in PO(M)$, those cycles must be coalition cycles for M , and the thesis follows. \square

A *Pareto-improving algorithm* A starts from a matching M that is maximal and trade-in free, chooses a coalition cycle C of M , and replaces M with $M \Delta C$. Note that $M \Delta C$ is maximal and trade-in free. When M has no coalition cycle, then the algorithm outputs M which by Lemma 1.3.4 is Pareto-optimal. The *number of displaced men* $d(A; M)$ is the number of men for which the partner in the input matching M and output matching differ.

Lemma 4.2.2. *Let M be a popular matching and $M^0 \in PO(M)$. Then there is Pareto-improving algorithm A that on input M , outputs M^0 . Conversely, any Pareto-improving algorithm A on input M outputs a matching from $PO(M)$.*

Proof. Let M be a popular matching and $M^0 \in PO(M)$. Since a popular matching is

maximal and trade-in free, we deduce that M is a well-defined input to any Pareto-improving algorithm. By Lemma 4.2.1, there is a collection of vertex-disjoint cycles $C_1; \dots; C_k$ such that $M \oplus M^\theta = fC_1; \dots; C_kg$. Let now $i \geq [k]$. Since $C_1; \dots; C_i$ are vertex disjoint, $C_{i+1}; \dots; C_k$ are coalition cycles in $((M \oplus C_1) \oplus \dots) \oplus C_i$. Hence, there is a Pareto-improving algorithm that, for $i \geq [k]$, at step i selects C_i , and this algorithm outputs M^θ .

The converse statement follows from the fact that, for any matching M and coalition cycle C , $M \oplus C \oplus M$. □

Lemma 4.2.3. *Let M be a popular matching. Then*

$$\Lambda(M) = \min d(A; M);$$

where the minimum is taken over all Pareto-improving algorithms A .

Proof. Let M^θ realize $\Lambda(M)$. Then $M^\theta \geq PO(M)$ and, by Lemma 4.2.2 and its proof, there exists a Pareto-Improving algorithm A that outputs M by selecting, at each step, one of the k coalition cycles from $M \oplus M^\theta$. Hence, using Observation 4.1.2, we have:

$$\min d(A; M) \leq d(A; M) = \# \{a \geq A : M(a) \neq M^\theta(a)\} = \Lambda(M);$$

showing one side of the inequality.

Conversely, consider the execution of a Pareto-improving algorithm A on input M , and let M^θ be the matching it outputs on input M . From Lemma 4.2.2, $M^\theta \geq PO(M)$. Hence, using again Observation 4.1.2, we deduce:

$$\Lambda(M) \leq \Delta(M; M^\theta) = \# \{a \geq A : M(a) \neq M^\theta(a)\} = d(A; M);$$

showing the other side of the inequality and concluding the proof. □

Lemma 4.2.4. *Let M be a popular matching that contains k vertex-disjoint coalition cycles. Then $\Lambda(M) = \Delta(M) = k$.*

Proof. $\Lambda(M) = \Delta(M)$ follows by definition. Let M^θ be a Pareto-optimal matching that realizes $\Delta(M)$. Consider a coalition cycle C of M . Since M^θ is Pareto-optimal, then $C \not\preceq M^\theta$. This implies that at least one man in C has changed his partner. Since there are k vertex-disjoint coalition cycles, and each cycle has at least one man that changed his partner, then $\Delta(M) = k$. \square

Lemma 4.2.5. *For matchings M, M^θ , we have:*

$$\frac{1}{2}jM = M^\theta j = \Delta(M; M^\theta) = jM = M^\theta j:$$

Proof. Assume there are k men, $m_1; m_2; \dots; m_k$, that change partners between M and M^θ . Then

$$M = M^\theta = f m_1 M(m_1); m_1 M^\theta(m_1); m_2 M(m_2); m_2 M^\theta(m_2); \dots; m_k M(m_k); m_k M^\theta(m_k)g:$$

If $m_i, i \geq 1 \dots k$ is unmatched in M (or M^θ), then $m_i M(m_i) = ;$ (or $m_i M^\theta(m_i) = ;$). We conclude that $jM = M^\theta j = 2k$, or $\frac{1}{2}jM = M^\theta j = k = \Delta(M; M^\theta)$.

On the other hand for each $i = 1 \dots k$ at least one of $M(m_i)$ and $M^\theta(m_i)$ is non-empty, call it $M^{\theta\theta}(m_i)$. This implies that

$$M = M^\theta = f m_1 M^{\theta\theta}(m_1); m_1 M^\theta(m_1); \dots; m_k M^{\theta\theta}(m_k)g:$$

We conclude that $jM = M^\theta j = k = \Delta(M; M^\theta)$. \square

Lemma 4.2.6. *For each marriage instance $(G(A [B]; <)$ we have $\Gamma(G) = \Delta(G) = jA j - 1$, and there are infinitely many instances with $\Gamma(G) = jA j - 2$ (resp. $\Delta(G) = jA j - 1$).*

Proof. Kesten’s EADAM algorithm [37] produces a Pareto-optimal matching that agrees with the man-optimal stable matching in at least one edge. This shows $\Gamma \geq \Delta - |jA_j| - 1$. The example with $\Delta = |jA_j| - 1$ is presented in Section 4.3. The example with $\Gamma = |jA_j| - 2$ is given in Section 4.3. \square

4.3 Phenomenology

In this section, we provide some examples. The goal is twofold: First, we want to illustrate the complexity of the problems investigated in this chapter. Second, we want to show some extreme scenarios, showing that Δ , Λ , and Γ can be very different, and can essentially cover all the spectrum, from being 0 to linear in $|jA_j|$.

$\Delta = |jA_j| - 1$ and $\Gamma = 0$

In this section we provide an infinite family of instances, such that $\Gamma = 0$ and $\Delta = |jA_j| - 2$. This construction is based on a modification of Irving and Leather’s instances [29] given by Faenza and Zhang [15].

Algorithm 6 and Algorithm 7 construct matrices A_n and B_n that represent complete preferences lists for agents from A and B respectively, with $|jA_j| = |jB_j| = 2^n + 1$. An element c_{ij} of a matrix indicates that entity c_{ij} is in position j of i th preference list. We denote the resulting marriage instance by $(G_n; <_n)$. Let $I(i)$ be a matrix of size i by i with all elements being 1. Let $A(i; j; k; l)$ be a submatrix of A with all entries of A at intersection of rows $i; i + 1; \dots; j$ and columns $k; k + 1; \dots; l$. $A(i;)$ indicates the i th row of matrix A .

Algorithm 7: Construction of B_n

$B_n(1 : 2^n; 1) = A_n(1 : 2^n; 2^n);$ $B_n(1 : 2^n; 2) = (2^n + 1; \dots; 2^n + 1);$ for ($i = 3 : \dots : 2^n + 1$) f <div style="margin-left: 20px;"> $B_n(1 : 2^n; i) = A_n(1 : 2^n; 2^n + 2 - i);$ </div> g $B_n(2^n + 1;) = (2^n + 1; 1; 2; 3; \dots; 2^n);$

Algorithm 6: Construction of A_n

initialization $C_1 = \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix};$
for ($i = 2; \dots; n$) *f*
 $C_{i-1}^g = C_{i-1} + 2^{i-1}I(2^{i-1});$
 $C_i = \begin{pmatrix} C_{i-1} & C_{i-1}^g \\ C_{i-1}^g & C_{i-1} \end{pmatrix};$
g
 $A_n(1 : 2^n; 1 : 2^n) = C_n;$
 $A_n(2^n + 1; \cdot) = (1; 2; \dots; 2^n + 1);$
 $A_n(\cdot; 2^n + 1) = (2^n + 1; \dots; 2^n + 1)$

Note that, in the instance $(G_n; <_n)$, for $i \geq [2^n]$, man i 's top preference is i and his second-to-least favorite woman is $2^n - i + 1$. Every man's least favorite woman is $2^n + 1$. Conversely, for $i \geq [2^n]$, woman i 's favorite partner is man $2^n - i + 1$, her least favorite is man i . Woman $2^n + 1$'s favorite partner is man $2^n + 1$.

Table 4.1 shows an example of the construction for $n = 3$.

1	1	2	3	4	5	6	7	8	9	1	8	9	7	6	5	4	3	2	1
2	2	1	4	3	6	5	8	7	9	2	7	9	8	5	6	3	4	1	2
3	3	4	1	2	7	8	5	6	9	3	6	9	5	8	7	2	1	4	3
4	4	3	2	1	8	7	6	5	9	4	5	9	6	7	8	1	2	3	4
5	5	6	7	8	1	2	3	4	9	5	4	9	3	2	1	8	7	6	5
6	6	5	8	7	2	1	4	3	9	6	3	9	4	1	2	7	8	5	6
7	7	8	5	6	3	4	1	2	9	7	2	9	1	4	3	6	5	8	7
8	8	7	6	5	4	3	2	1	9	8	1	9	2	3	4	5	6	7	8
9	1	2	3	4	5	6	7	8	9	9	9	1	2	3	4	5	6	7	8

Table 4.1: An example of the construction from Algorithm 6 (left) and Algorithm 7 (right) for $n = 3$. The unique stable matching is circled.

In the following, we fix $n \geq \mathbb{N}$ and the instance $(G_n; <_n)$. We denote a matching of $(G_n; <_n)$ by an ordered tuple with $2^n + 1$ elements, where man i is matched to the woman corresponding to the i -th woman in the tuple. We let $S_n := (2^n; 2^n - 1; \dots; 1; 2^n + 1)$ and $P_n := (1; 2; \dots; 2^n + 1)$.

Lemma 4.3.1. S_n is the only stable matching.

Proof. One readily verifies that Gale-Shapley algorithm with men proposing outputs S_n . On the other hand, as each woman is matched to her favorite partner in S_n , it is also the woman-optimal stable matching. Hence, $(G_n; <_n)$ only has one stable matching. \square

Lemma 4.3.2. P_n is Pareto-optimal and popular.

Proof. P_n is obtained by serial dictatorship algorithm using the following order of men: $f1; 2; \dots; 2^n + 1g$. This implies that P_n is Pareto-optimal by Lemma 1.3.9.

We now consider graph G_{P_n} obtained by labelling and removing edges as discussed in Chapter 1. Notice that:

$$\begin{aligned} (i; j) &= (; +) \text{ for } i = 1; \dots; 2^n \text{ and all possible } j; \\ (2^n + 1; j) &= (+; +) \text{ for } j = 1; \dots; 2^n; \\ (k; 2^n + 1) &= (;) \text{ for } k = 1; \dots; 2^n. \end{aligned}$$

This implies that all $(+; +)$ edges are adjacent. We conclude that there is no alternating path with two $(+; +)$ edges. All adjacent edges to $2^n + 1 \in B$ are $(;)$, this implies that there is no alternating cycle that contains $(+; +)$. Since all vertices are matched, using Theorem 1.3.4 we conclude that P_n is popular. \square

Lemma 4.3.3. $jS_n \setminus Mj = 1$ for all Pareto-optimal matchings M .

Proof. Let M be a Pareto-optimal matching of $(G_n; <_n)$. Assume first that the $2^n + 1$ -th men and women are matched to each other in M . Then $M = P_n$, else P_n Pareto-dominates M , a contradiction. The statement then follows since $jS_n \setminus P_nj = 1$.

Hence, we can assume $M(2^n + 1) \notin 2^n + 1$ and, by contradiction, that $jM \setminus S_nj = 2$. Then we have $M(i) = 2^n - i + 1$ and $M(j) = 2^n - j + 1$ for some $i \neq j \in [2^n]$. Construct matching M' starting from M swapping the partner of men i and j – that is, $M'(i) = 2^n - j + 1$ and $M'(j) = 2^n - i + 1$. Clearly $M' \succ M$, hence M is not Pareto-optimal. \square

Lemma 4.3.2 implies that $\Gamma = 0$. Lemma 4.3.1 and 4.3.3, together with the fact that preference lists of $(G_n; <_n)$ are complete, imply that $\Delta = |A| - 1$.

$\Gamma = 0$, $\Delta(M) = \Theta(|A|)$ for any dominant or stable matching M

It is a natural question whether Γ can be computed or well-approximated just by restricting the search over the set of stable or dominant matchings. We have seen in Section 1 that those are well-behaved classes of popular matchings and in Section 2.3 that the best polynomial-time approximation to a popular matching of maximum weight (with nonnegative weights) can be computed by picking the stable or dominant matchings of maximum weights.

In this section, we show this is not the case. In fact, We construct an infinite family of instances, such that $\Delta = 0$ is achieved on a *middle* matching (i.e., neither stable nor dominant), and $\Delta = \Omega(n)$ for any other non middle matching, where n is the number of men.

Our example of an instance $(G(A \cup B; E); <)$ is constructed from the same kind of blocks (or gadgets) that are connected with additional edges. One block i consists of five men $m_i^0; m_i^1; m_i^2; m_i^3; m_i^4$ and four women $w_i^0; w_i^1; w_i^2; w_i^3$, Fig. 4.3. Their preferences lists are the following:

$$\begin{array}{ll}
 m_i^0 : & w_i^0 & w_i^0 : & m_i^1; m_i^0 \\
 m_i^1 : & w_i^0; w_i^1 & w_i^1 : & m_i^1; m_i^2 \\
 m_i^2 : & w_i^3; w_i^1; w_i^2 & w_i^2 : & m_i^2 \\
 m_i^3 : & w_i^3 & w_i^3 : & m_i^2; m_i^3; m_i^4 \\
 m_i^4 : & w_i^3 & &
 \end{array}$$

Consider an instance I with $k \geq 2$ blocks with $5k$ men and $4k$ women. Add $2k(k-1)$ edges: $(m_i^1; w_j^0)$ and $(m_i^3; w_j^3)$ for any $i; j = 1; \dots; k$, $i \neq j$. For each man m_i^1 and m_i^3 add the $k-1$ women to the beginning of his preference list in arbitrary




Figure 4.3: One block i consists of five men $m_i^0; m_i^1; m_i^2; m_i^3; m_i^4$ and four women $w_i^0; w_i^1; w_i^2; w_i^3$ with their preferences labeled on the edges.

order. For each women w_i^0 and w_i^3 add $k - 1$ men to the end of her preferences list in arbitrary order.

Lemma 4.3.4. *Edges $(m_i^1; w_j^0)$, $(m_i^3; w_j^3)$, $(m_i^4; w_j^3)$ and $(m_i^2; w_j^1)$ are not contained in any popular matching for any $i; j = 1; \dots; k$, $i \neq j$.*

Proof. If one of the edges $(m_i^1; w_j^0)$, $i; j = 1; \dots; k$, $i \neq j$ belongs to a matching \mathcal{M} then $(m_j^0; w_j^0) = (+; +)$ is adjacent to an unmatched node m_j^0 . This implies that \mathcal{M} cannot be popular.

Similarly, if one of the edges $(m_i^3; w_j^3)$, $i; j = 1; \dots; k$, $i \neq j$ belongs to a matching \mathcal{M} then $(m_j^4; w_j^3) = (+; +)$ is adjacent to an unmatched node m_j^4 . This implies that \mathcal{M} cannot be popular.

Now we will prove that no popular matching contains one of the edges $(m_i^4; w_j^3)$, $i = 1; \dots; k$. By contradiction, assume there is a popular matching \mathcal{M} that contains an edge $(m_i^4; w_j^3)$, for some $i \in \{1; \dots; k\}$. From above $(m_i^3; w_j^3)$ is unmatched in any popular matching for $i; j = 1; \dots; k$, $i \neq j$. This implies that $(m_i^3; w_j^3) = (+; +)$ is

adjacent to an unmatched node m_i^3 and contradicts popularity of M .

In the last step we show that no popular matching contains one of the edges $(m_i^2; w_i^1)$, $i = 1; \dots; k$. By contradiction, assume there is a popular matching M that contains an edge $(m_i^1; w_i^1)$, for some $i \in 1; \dots; k$. From above $(m_i^4; w_i^3) \notin M$, for any $i = 1; \dots; k$. Notice that $(m_i^2; w_i^3) = (+; +)$, and $(m_i^3; w_i^3) \in M$, otherwise $(m_i^2; w_i^3) = (+; +)$ is adjacent to an unmatched node w_i^3 . Node m_i^1 has to be matched otherwise there is a path $m_i^1; w_i^1; m_i^2; w_i^3$ that contains an unmatched node m_i^1 and an edge $(m_i^2; w_i^3) = (+; +)$. We arrive to a contradiction that M is popular, since a path $m_i^0; w_i^0; m_i^1; w_i^1; m_i^2; w_i^3$ contains an unmatched node m_i^0 and an edge $(m_i^2; w_i^3) = (+; +)$. \square

Lemma 4.3.5. *There is exactly one dominant matching $M_0 = \{(m_i^0; w_i^0); (m_i^1; w_i^1); (m_i^2; w_i^2); (m_i^3; w_i^3) : i = 1; \dots; k\}$ in the instance I .*

Proof. First, we show that matching M_0 is popular. Notice that $(m_i^1; w_i^0)$ and $(m_i^2; w_i^3)$ are the only $(+; +)$ edges, and m_i^4 are the only unmatched nodes for $i = 1; \dots; k$. Edges $(m_i^1; w_i^0)$ and $(m_i^2; w_i^3)$ do not belong to any alternating path within any block i . Similarly, node m_i^4 does not belong to any alternating path with a $(+; +)$ edge within any block i . Since the edge $(m_i^1; w_i^0)$ is unmatched for any $i; j = 1; \dots; k$, $i \neq j$, then there is no alternating path that contains two $(m^1; w^0)$ edges from different blocks $i = i; j$. We apply the same argument for edges $(m^3; w^3)$ edges from different blocks $i = i; j$. Since there is no alternating path that connects edges $(m_i^1; w_i^0)$ and $(m_i^2; w_i^3)$ within one block, we argue that there is no alternating path between any two blocks that contains edges $(m_i^1; w_i^0)$ and $(m_j^2; w_j^3)$ from different blocks i and j . Similarly, we show that there is no alternating path that contains any $(+; +)$ edge from one block and unmatched node from another block.

Since nodes m_i^0 and w_i^2 have degree one for all blocks $i = 1; \dots; k$, and edge $(m_i^4; w_i^3)$ is unmatched in any popular matching for all $i = 1; \dots; k$, and dominant matching has the largest number of matched nodes, we conclude that $M_0 =$

$f(m_i^0; w_i^0); (m_i^1; w_i^1); (m_i^2; w_i^2); (m_i^3; w_i^3) : i = 1 :: k$ is the only popular matching with the largest number $8k$ matched nodes, and, therefore, is the only dominant matching. \square

Lemma 4.3.6. *There is exactly one stable matching $M_1 = f(m_i^1; w_i^0); (m_i^2; w_i^3) : i = 1 :: k$ in the instance I .*

Proof. For any block $i = 1 :: k$ edges $(m_i^0; w_i^0); (m_i^1; w_i^1); (m_i^2; w_i^1); (m_i^3; w_i^3); (m_i^4; w_i^3)$ are either $(+; -)$ or $(-; +)$. All edges between the blocks $(m_i^1; w_j^0); (m_j^3; w_i^3)$ are $(+; -)$ for $i; j = 1 :: k, i \neq j$. We conclude that there are no $(+; +)$ edges, and M_1 is stable. Similarly to Lemma 4.3.4, we argue that $(m_i^1; w_j^0), (m_j^3; w_i^3)$ are not contained in any stable matching for any $i; j = 1 :: k, i \neq j$. Recall that all stable matchings have the same people matched. We conclude that there are no any other options to match $m_i^1; w_j^0; m_j^3; w_i^3$ for $i = 1 :: k, i \neq j$, and M_1 is the only stable matching. \square

Lemma 4.3.7. *Matching $M_2 = f(m_i^0; w_i^0); (m_i^1; w_i^1); (m_i^2; w_i^3) : i = 1 :: k$ is middle and Pareto-optimal.*

Proof. First, we show that M_2 is popular. Notice that there is exactly one $(m_i^1; w_i^0) = (+; +)$ edge within any block $i = 1 :: k$. Since edge $(m_i^1; w_j^0)$ is unmatched for $i; j = 1 :: k, i \neq j$, we conclude that there is no alternating path with two $(+; +)$ edges. Similarly, since $(m_i^2; w_i^1) = (-; -)$ for all $i = 1 :: k$, there is no alternating path that contains unmatched node m_i^4 and $(m_i^1; w_i^0) = (+; +)$ edge for any $i = 1 :: k$. This implies that M_2 is popular. Since $jM_1j < jM_2j < jM_0j$, we conclude that M_2 is a middle matching.

Second, since nodes $m_i^4; m_i^3$ for all $i = 1 :: k$ are unmatched, they cannot be part of any coalition cycle. Since m_i^2 is matched to its top choice for all $i = 1 :: k$, and nodes $m_i^0, i = 1 :: k$ have degree one, they cannot be in any coalition cycle. Consider nodes $m_i^1, i = 1 :: k$. Assume there is $i \geq 1 :: k$ such that m_i^1 is in

a coalition cycle C . Since $(m_i^1; w_i^0); (m_j^1; w_j^0)$ are unmatched for all $j = 1; \dots; k$, we conclude that $(m_i^1; w_i^0); (m_i^1; w_i^1) \succeq C$, and consequently $(m_i^2; w_i^1) \succeq C$. Since $(m_i^2; w_i^1) = (m_i^1; w_i^1)$, C is not a coalition cycle. This implies that M_2 is Pareto-optimal. \square

Lemma 4.3.8. *Any Pareto-optimal matching M contains at most one edge from $f(m_i^1; w_i^0) : i = 1; \dots; kg$ and at most one edge from $f(m_i^3; w_i^3) : i = 1; \dots; kg$.*

Proof. Assume by contradiction there is a Pareto-optimal matching M that contains two edges $(m_i^1; w_i^0); (m_j^1; w_j^0)$ from $f(m_i^1; w_i^0) : i = 1; \dots; kg$. Since w_i^0 and w_j^0 are the most preferred choices for m_i^1 and m_j^1 correspondingly, then there is a coalition cycle $m_i^1; w_i^0; m_j^1; w_j^1$. This contradicts the Pareto-optimality of M .

Similarly, assume by contradiction there is a Pareto-optimal matching M that contains two edges $(m_i^3; w_i^3); (m_j^3; w_j^3)$ from $f(m_i^3; w_i^3) : i = 1; \dots; kg$. Since w_i^3 and w_j^3 are the most preferred choices for m_i^3 and m_j^3 correspondingly, then there is a coalition cycle $m_i^3; w_i^3; m_j^3; w_j^3$. This contradicts the Pareto-optimality of M . \square

Lemma 4.3.7 implies that $\Gamma = 0$ achieved on a middle matching M_2 . From Lemma 4.3.6 and Lemma 4.3.8 we conclude that $\Delta(\mathcal{M}_1)$ is linear in k , which is linear in the number of men n . Similarly, from Lemma 4.3.5 and Lemma 4.3.8 we conclude that $\Delta(\mathcal{M}_0)$ is linear in n . Since \mathcal{M}_1 and \mathcal{M}_0 are the only stable and dominant matchings, we conclude that they are the only not middle matchings.

$$\Gamma = |A_j| - 2$$

We now construct an infinite family of instances, each with exactly one popular matching \mathcal{M}_0 , and such that \mathcal{M}_0 has at most two common edges with any Pareto-optimal matching. Hence, for this instance, $\Gamma = |A_j| - 2$.

Our example will be constructed from the same kind of blocks (or gadgets) that are connected with additional edges. One block i consists of three men $m_i^0; m_i^1; m_i^2$ and two women $w_i^0; w_i^1$, Fig. 4.4. Their preferences lists are the following:

$$\begin{array}{ll}
 m_i^0 : & w_i^1; w_i^0 \\
 m_i^1 : & w_i^0; w_i^1 \\
 m_i^2 : & w_i^1; w_i^0 \\
 w_i^0 : & m_i^0; m_i^1; m_i^2 \\
 w_i^1 : & m_i^1; m_i^2; m_i^0
 \end{array}$$




Figure 4.4: One block i consists of three men $m_i^0; m_i^1; m_i^2$ and two women $w_i^0; w_i^1$ with their preferences labeled on the edges.

Consider an instance I with $k \geq 2$ blocks with $3k$ men and $2k$ women. Add $2k(k-1)$ edges: $(m_i^0; w_j^0)$ and $(m_i^1; w_j^1)$ for any $i; j = 1; \dots; k$, $i \neq j$. For each man m_i^0 and m_i^1 add the $k-1$ women to his preference lists between his first and second choices in the gadget, in arbitrary order. For each women w_i^0 and w_i^1 add $k-1$ men to the end of her preferences lists in arbitrary order.

Lemma 4.3.9. *There is exactly one popular matching $M_0 = f(m_i^0; w_i^0); (m_i^1; w_i^1) : i = 1; \dots; kg$ in the instance I .*

Proof. First, we will prove that no popular matching contains one of the edges $(m_i^0; w_j^0)$ or $(m_i^1; w_j^1)$, $i \neq j \in [k]$.

By contradiction, assume there is a popular matching M that contains an edge $(m_i^1; w_j^1)$, $i \neq j \in [1; \dots; kg]$. Since all new edges during the construction were added in the end of w_j^1 list, then w_j^1 prefers m_j^0 , m_j^1 and m_j^2 . We

conclude that $(m_j^2; w_j^0) \succeq M$, otherwise $(m_j^2; w_j^1) = (+; +)$ is adjacent to an unmatched node m_j^2 , which would contradict popularity. A path $m_j^1; w_j^0; m_j^2; w_j^1$ contains $(m_j^1; w_j^0); (m_j^2; w_j^1) = (+; +)$. This contradicts popularity of M and implies that $(m_i^1; w_i^1) \not\preceq M$, for $i \notin j \succeq [k]$ and all popular matchings M .

Similarly, by contradiction, assume there is a popular matching M that contains edge $(m_i^0; w_i^0)$, for some $i \notin j \succeq [k]$. Since all new edges during the construction were added in the end of w_j^0 list, then w_j^0 prefers m_j^0, m_j^1 and m_j^2 to m_j^0 . We conclude that $(m_j^2; w_j^1) \succeq M$, otherwise $(m_j^2; w_j^0) = (+; +)$ is adjacent to an unmatched node m_j^2 , which would contradict popularity. Consider an edge $(m_j^1; w_j^1)$. In the first part we concluded that $(m_j^1; w_i^1) \not\preceq M$ for all $i \notin j$. This implies that $(m_j^1; w_j^0) = (+; +)$ is adjacent to an unmatched node m_j^1 , which contradicts popularity of M and implies that $(m_i^0; w_i^0) \not\preceq M$, for $i \notin j \succeq [k]$ and all popular matchings M .

We next show that there is no popular matching that contains an edge $(m_i^0; w_i^1)$ or $(m_i^2; w_i^1)$ for any $i \succeq [k]$.

By contradiction, assume there is a popular matching M that contains an edge $(m_i^0; w_i^1)$, for some $i \succeq [k]$. We conclude that $(m_i^2; w_i^0) \succeq M$, otherwise $(m_i^2; w_i^1) = (+; +)$ is adjacent to an unmatched node m_i^2 , which would contradict popularity. Observe that $(m_i^0; w_i^0)$ is a $(-; +)$ edge. Cycle $m_i^2; w_i^0; m_i^0; w_i^1$ contains $(m_i^2; w_i^1) = (+; +)$, which contradicts popularity.

Similarly, by contradiction, assume there is a popular matching M that contains an edge $(m_i^2; w_i^1)$, for some $i \succeq f1; \dots; kg$. We conclude that $(m_i^1; w_i^0) \succeq M$, otherwise $(m_i^1; w_i^1) = (+; +)$ is adjacent to an unmatched node m_i^1 , which would contradict popularity. Notice that $(m_i^0; w_i^0) = (+; +)$ and m_i^0 is unmatched, since we proved earlier that $(m_i^0; w_i^0) \not\preceq M$, for all $i \notin i$. This contradicts popularity. We conclude that $(m_i^0; w_i^1); (m_i^2; w_i^1) \not\preceq M$ for any $i \succeq f1; \dots; kg$ and any popular matching M .

Notice that w_i^1 is matched in any popular matching M for all $i \succeq f1; \dots; kg$, otherwise $(m_i^0; w_i^1) = (+; +)$ is adjacent to unmatched node w_i^1 , which would contradict

popularity. From above $(m_j^1; w_j^1) \notin M$ for all $i; j \in \{1, \dots, k\}$ and all popular matchings M , and $(m_i^2; w_i^1) \notin M$, we conclude that $(m_i^1; w_i^1) \in M$ for all $i \in \{1, \dots, k\}$. Similarly, $(m_i^0; w_i^0) \in M$ for all $i \in \{1, \dots, k\}$, otherwise $(m_i^0; w_i^0) = (+; +)$ is adjacent to unmatched node m_i^0 .

We conclude that edges $(m_i^0; w_i^0); (m_i^1; w_i^1); i = 1 \dots k$ belong to all popular matchings. Since there is no any other combinations for matched edges, and popular matchings exist in every bipartite instance, there is only one popular matching M_0 . This concludes the proof. \square

Lemma 4.3.10. *Any Pareto-optimal matching M_1 contains at most one edge from $f(m_i^0; w_i^0) : i = 1 \dots k$ and at most one edge from $f(m_i^1; w_i^1) : i = 1 \dots k$.*

Proof. Assume by contradiction there is a Pareto-optimal matching M that contains two edges $(m_i^0; w_i^0); (m_j^0; w_j^0)$ from $f(m_i^0; w_i^0) : i = 1 \dots k$. Since w_i^0 and w_j^0 are the least preferred choice for m_i^0 and m_j^0 correspondingly, then there is a coalition cycle $m_i^0; w_i^0; m_j^0; w_j^0$. This contradicts to Pareto-optimality of M .

Similarly, assume by contradiction there is a Pareto-optimal matching M that contains two edges $(m_i^1; w_i^1); (m_j^1; w_j^1)$ from $f(m_i^1; w_i^1) : i = 1 \dots k$. Since w_i^1 and w_j^1 are the least preferred choice for m_i^1 and m_j^1 correspondingly, then there is a coalition cycle $m_i^1; w_i^1; m_j^1; w_j^1$. This contradicts to Pareto-optimality of M . \square

Lemma 4.3.9 and 4.3.10 implies that the popular matching M_0 has at most two common edges with any Pareto-optimal matching. Since all men are matched in M_0 , we conclude that $|\Gamma \cap jA_j| \leq 2$.

4.4 Two easy cases

In this Theorem, we show polynomial-time results. We start with the proof of Theorem 4.1.4.

Proof. We first deal with the computation of Δ . Consider the following algorithm:

1. For $t; \ell \geq [k]$:
 - a) Guess two disjoint sets of edges, say E^{in} and E^{out} , of size respectively t and ℓ , such that the number of men incident to $E^{\text{in}} \cup E^{\text{out}}$ is at most k ;
 - b) Compute the man-optimal stable matching M in $S(E^{\text{in}}; E^{\text{out}})$, see Lemma 1.3.1 and Corollary 1.3.1.
 - c) if $M^\theta := (M \cap E^{\text{in}}) \cup E^{\text{out}}$ is a Pareto-optimal matching, output $(M; M^\theta)$ and stop.
2. output $\Delta(G) > k$ and stop.

We now show that the algorithm above is correct and runs in time $jA \cup Bj^{O(k^2)}$.

The running time follows from the fact that there are $jA \cup Bj^{O(k)}$ sets of edges of size at most k .

For the correctness, suppose first $\Delta(G) \leq k$. Then, there exists a stable matching M and a Pareto-optimal matching M^θ such that $\Delta(M; M^\theta) \leq k$. Now let $t := jM \cap M^\theta j \leq k$ and $\ell = jM^\theta \cap M j \leq k$ and let $E^{\text{in}} := M \cap M^\theta$, $E^{\text{out}} := M^\theta \cap M$. Note that $t; \ell; E^{\text{in}}; E^{\text{out}}$ is a valid guess of the algorithm.

We claim that M is a man-optimal stable matching in $S(E^{\text{in}}; E^{\text{out}})$. Suppose not, then, by Lemma 1.3.1, there exists a matching $\bar{M} \in S(E^{\text{in}}; E^{\text{out}})$ with $\bar{M} > M$. Because of Theorem 1.3.1, graph $G(A \cup B; \bar{M} \cap M)$ is a collection of singletons and cycles whose edges alternate between M and \bar{M} . Take one such cycle C . Note that C does not contain nodes adjacent to edges of $E^{\text{in}} \cup E^{\text{out}}$. Since $\bar{M} > M$, $\bar{M}(m) > M(m)$ for every $m \in V(C) \cap A$. Hence, C is a coalition cycle for M , and, as it contains no node adjacent to edges of $E^{\text{in}} \cup E^{\text{out}}$, it is also a coalition cycle for M^θ , contradicting the fact that M^θ is Pareto-optimal.

Therefore, M is the man-optimal stable matching in $S(E^{\text{in}}; E^{\text{out}})$. Hence, the pair $M; M^\theta$ is constructed during some iteration of the algorithm.

Conversely, suppose the algorithm outputs a pair $(M; M^0)$. Then by construction M is stable, while M^0 is Pareto-optimal since the algorithm verifies so. We have:

$$\Delta(M; M^0) = \# \text{ men that have a different partner in } M \text{ and } M^0 \quad k$$

by construction, hence the output of the algorithm is correct.

Now for the computation of Λ , repeat the algorithm above after modifying (c) as follows.

(c') if $M^0 := (M \cap E^{\text{in}}) \cup E^{\text{out}} \in PO(M)$, output $(M; M^0)$ and stop.

The correctness of the algorithm follows as above, after observing that, if $\Delta(M; M^0) = \Lambda$, then M is again a man-optimal stable matching in $S(E^{\text{in}}; E^{\text{out}})$. \square

We next show that, for a fixed Pareto-optimal matching M , $\Delta(M)$ can be computed efficiently.

Lemma 4.4.1. *There is a polynomial-time algorithm that, given a marriage instance $(G; <)$ and a Pareto-optimal matching M , computes $\Delta(M)$ and the stable matching achieving it.*

Proof. Consider a preference system $(G(A \cup B; E); <)$ and a Pareto-optimal matching M . Assign weight 0 to all edges that belong to M , and weight 1 to all other edges.

We now claim that

$$\arg \min_{M^0 \in S} \Delta(M; M^0) = \arg \min_{M^0 \in S} w(M^0): \quad (4.5)$$

Indeed, let $M^0 \in S$. Note that:

$$\Delta(M; M^0) = \# \text{ men that have a different partner in } M \text{ and } M^0 = w(M^0) + jA(M \cap M^0)j:$$

As by Theorem 1.3.1, $jA(M \cap M^0)j$ does not depend on the choice of M^0 , and the claim follows. The thesis follows from the fact that a stable matching of minimum weight can be computed in polynomial time, see Theorem 1.3.3. \square

4.5 Proof of Theorem 4.1.3

In this section, we prove Theorem 4.1.3. For the reader's convenience, the theorem is restated below.

Theorem 4.1.3. *There exists a universal constant c such that the following holds. Unless $P=NP$, for each $\epsilon > 0$, there is no polynomial-time algorithm that, given a marriage instance $(G(A \cup B; E); <)$, accepts if $\Delta(G) \leq \epsilon |A|$, rejects if $\Delta(G) \geq c\epsilon |A|$, and is free to accept or reject if none of these cases apply. In particular, computing $\Delta(G)$ is NP-hard.*

The main idea of the proof is to construct a map from instances of Monotone 3-SAT to marriage instance $(G(A \cup B; E); <)$ such that if ϕ is satisfiable then Δ is small, but if ϕ is not satisfiable, there are many disjoint coalition cycles that results in Δ being (relatively) large.

For a fixed ϵ , the instance of the marriage problem is constructed from gadgets corresponding to literals and some edges that connect them. The following is the description of a gadget that corresponds to one literal. Positive and negative literals have the same associated gadgets, but the true and false assignments correspond to different stable matchings in those gadgets. Nodes of the gadget that correspond to literal ℓ_i are labeled as $m_j(\ell_i)$ or $w_k(\ell_i)$, for integers j and k . Sometime we omit the dependency on ℓ_i if it is clear to which gadget the nodes belong.

One gadget consists of $2l + 8$ men and $l + 6$ women (for some l to be fixed later on), see Fig. 4.5. Fig. 4.6 and Fig. 4.7 provide an example of the construction.

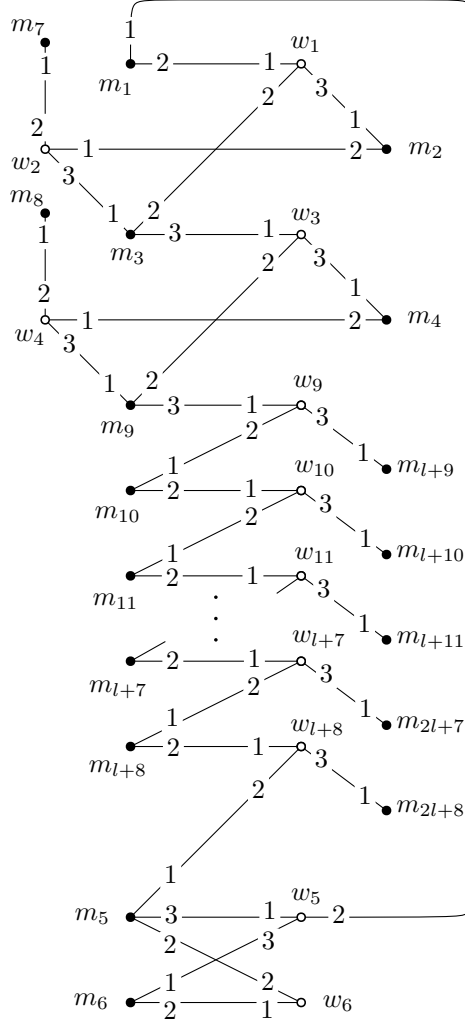


Figure 4.5: One gadget.

If \bar{i} and \bar{j} are literals that correspond to a variable and its negation respectively, then the corresponding gadgets are connected by extra edges $m_1(\bar{i})w_5(\bar{j})$. For man $m_1(\bar{i})$, add woman $w_5(\bar{j})$ to the beginning of his preference list, and for a woman $w_5(\bar{j})$, add man $m_1(\bar{i})$ after her the first choice $m_5(\bar{j})$.

For a clause with exactly two positive literals $\bar{1}; \bar{2}$, add edges $m_{i+1}(\bar{1})w_i(\bar{2})$ and $m_{i+1}(\bar{2})w_i(\bar{1})$ for $i = 9; 10; \dots; l + 7$. For men $m_{i+1}(\bar{1}); m_{i+1}(\bar{2})$ add women $w_i(\bar{1}); w_i(\bar{2})$, $i = 9; 10; \dots; l + 7$, to the top of their preferences list respectively. For women $w_i(\bar{1}); w_i(\bar{2})$ add men $m_{i+1}(\bar{1}); m_{i+1}(\bar{2})$, $i = 9; 10; \dots; l + 7$, to the end of their preferences list respectively. If the clause has three positive literals $\bar{1}; \bar{2}; \bar{3}$,

add edges $m_{i+1}(\lrcorner_1)w_i(\lrcorner_3)$, $m_{i+1}(\lrcorner_2)w_i(\lrcorner_1)$ and $m_{i+1}(\lrcorner_3)w_i(\lrcorner_2)$ for $i = 9;10;\dots;l+7$. The preference lists are modified similarly to the case with two literals in a clause.

For a clause with exactly two negative literals $\lrcorner_1; \lrcorner_2$, add edges $m_i(\lrcorner_1)w_i(\lrcorner_2)$ and $m_i(\lrcorner_2)w_i(\lrcorner_1)$ for $i = 9;10;\dots;l+8$. For men $m_i(\lrcorner_1); m_i(\lrcorner_2)$ add women $w_i(\lrcorner_1); w_i(\lrcorner_2)$, $i = 10;\dots;l+8$, after their top choice in the preference lists respectively. For women $w_i(\lrcorner_1); w_i(\lrcorner_2)$ add men $m_i(\lrcorner_1); m_i(\lrcorner_2)$, $i = 9;10;\dots;l+8$, to the end of their preferences list respectively. If the clause has three negative literals $\lrcorner_1; \lrcorner_2; \lrcorner_3$, add edges $m_i(\lrcorner_1)w_i(\lrcorner_3)$, $m_i(\lrcorner_2)w_i(\lrcorner_1)$ and $m_i(\lrcorner_3)w_i(\lrcorner_2)$ for $i = 9;10;\dots;l+8$. The preference lists are modified similarly to the case with two literals in a clause.

Lemma 4.5.1. *Any stable matching M of $(G(A [B; E]; <))$ does not contain any edges between the gadgets. Moreover, the set of stable nodes of G in each gadget is $[_{i=1;\dots;6,9;\dots;l+8} f m_i; w_i g$.*

Proof. First, we prove that edges between the gadgets in the same clause are not stable. Assume there is an edge $m_i(\lrcorner_j)w_i(\lrcorner_k) \in M$ between two negative literals \lrcorner_j and \lrcorner_k . This implies that $m_{i+1}(\lrcorner_k)w_i(\lrcorner_k) = (+; +)$ and M is not stable. Similarly, we argue that edges between the gadgets in a positive clause are not stable.

Notice, that $M^0 = [_{i=1;2;\dots;6;9;\dots;l+8} f m_i; w_i g$ for all gadgets is stable, since every woman is matched to her favorite partner. We conclude that the set of stable nodes in each gadget is $[_{i=1;\dots;6,9;\dots;l+8} f m_i; w_i g$. Edges m_3w_2 and m_9w_4 are not stable in all gadgets, since otherwise $m_7w_2 = (+; +)$ and $m_8w_2 = (+; +)$.

Second, we prove that edges connecting two literals that correspond to a variable and its negation are not stable. Assume by contradiction that there are two literals \lrcorner_1 and \lrcorner_2 that correspond to a variable and its negation respectively, and $m_1(\lrcorner_1)w_5(\lrcorner_2) \in M$. Since nodes $m_6(\lrcorner_2); m_5(\lrcorner_2); m_{l+8}(\lrcorner_2); \dots; m_9(\lrcorner_2); m_3(\lrcorner_2)$ are matched in M^0 , and edges between the gadgets and $m_3w_2; m_9w_4$ are not stable, then $m_5(\lrcorner_2)w_{l+8}(\lrcorner_2); m_{l+8}(\lrcorner_2)w_{l+7}(\lrcorner_2); \dots; m_{10}(\lrcorner_2)w_9(\lrcorner_2); m_9(\lrcorner_2)w_3(\lrcorner_2); m_3(\lrcorner_2)w_1(\lrcorner_2) \in M$.

This leaves node $m_1(\cdot_2)$ unmatched, a contradiction, since $m_1(\cdot_2)$ is matched in M^θ , and all stable matchings match the same set of nodes, see Theorem 1.3.1. \square

Lemma 4.5.2. *Let M be a stable matching of $(G(A [B; E]; <)$. Then in each positive gadget H (the one that corresponds to a positive literal), exactly one of the following is true (dependency on \cdot is omitted):*

$$T(H) := \{m_i w_i : i = 1; 2; \dots; 6; 9; \dots; l + 8\}, \text{ or}$$

$$F(H) := \{m_6 w_6; m_5 w_{l+8}; m_{l+8} w_{l+7}; \dots; m_{10} w_9; m_9 w_3; m_3 w_1; m_1 w_5; m_2 w_2; m_4 w_4\}.$$

Similarly, in each negative gadget N (the one that corresponds to a negative literal), exactly one of the following is true (dependency on \cdot is omitted):

$$F(N) := \{m_i w_i : i = 1; 2; \dots; 6; 9; \dots; l + 8\}, \text{ or}$$

$$T(N) := \{m_6 w_6; m_5 w_{l+8}; m_{l+8} w_{l+7}; \dots; m_{10} w_9; m_9 w_3; m_3 w_1; m_1 w_5; m_2 w_2; m_4 w_4\}.$$

Proof. In each gadget the set of stable nodes is $\{m_i w_i : i = 1; \dots; 6, 9; \dots; l + 8\}$, Lemma 4.5.1. There are two options for matching node w_5 regardless of the type of the gadget: $m_5 w_5$ and $m_1 w_5$. The corresponding stable matchings are $\{m_i w_i : i = 1; 2; \dots; 6; 9; \dots; l + 8\}$ and $\{m_6 w_6; m_5 w_{l+8}; m_{l+8} w_{l+7}; \dots; m_{10} w_9, m_9 w_3; m_3 w_1; m_1 w_5; m_2 w_2; m_4 w_4\}$. \square

Lemma 4.5.3. *Let M be a stable matching of $(G(A [B; E]; <)$. Let \cdot_1 be a positive and \cdot_2 a negative occurrence of the same variable. Then exactly one of the following is true:*

$$(a) T(H(\cdot_1)) \cap M;$$

$$(b) T(N(\cdot_1)) \cap M.$$

Conversely, each maximal matching M that satisfies exactly one of (a)-(b) above for each pair of positive \cdot_1 and negative \cdot_2 occurrence of the same variable is stable.

Proof. Proof by contradiction. Assume $T(H(\cdot_1)) \cap M$ and $T(N(\cdot_1)) \cap M$. This implies that $m_7(\cdot_1) w_5(\cdot_2) = (+; +)$, contradicting the stability of M .

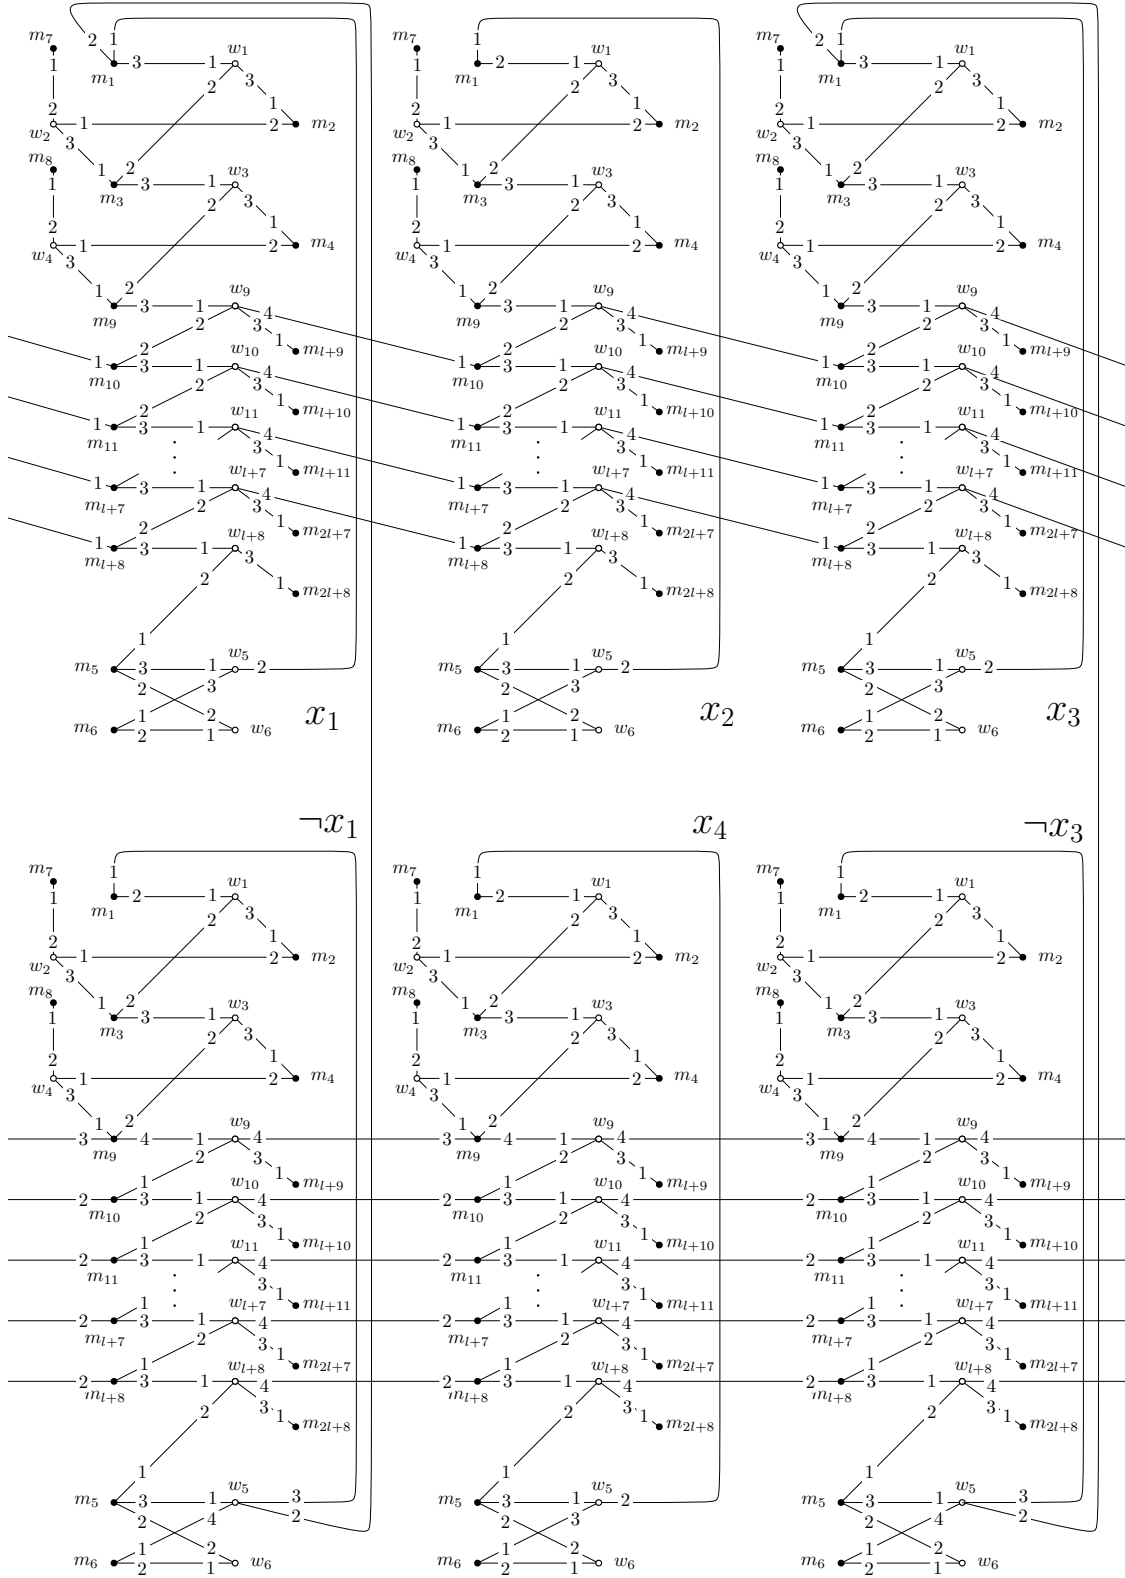


Figure 4.6: Corresponding graph for $(\mathcal{X}_1 - \mathcal{X}_2 - \mathcal{X}_3) \wedge (: \mathcal{X}_1 - : \mathcal{X}_4 - : \mathcal{X}_3)$.

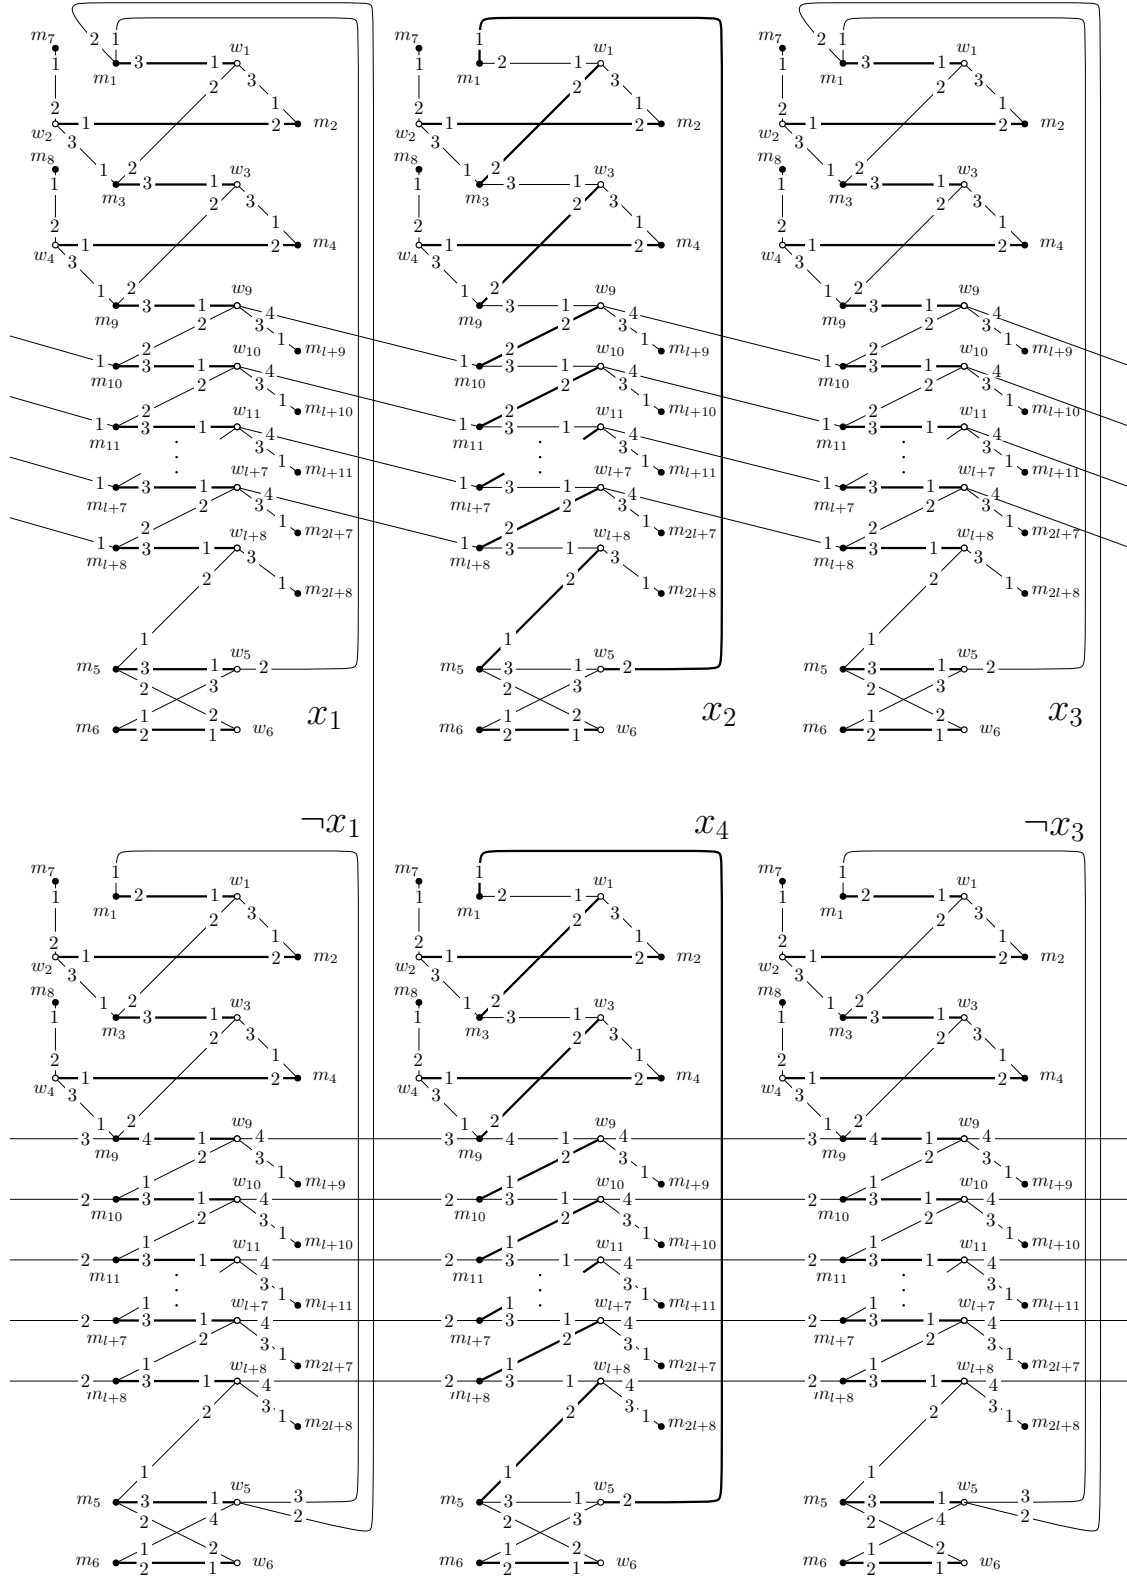


Figure 4.7: Corresponding stable matching (bold edges) for $(X_1 - X_2 - X_3) \wedge (X_1 - X_4 - X_3)$, where $X_1; X_3 = T; X_2; X_4 = F$.

Let M be a maximal matching that satisfy exactly one of (a)-(b) above for each pair of positive ℓ_1 and negative ℓ_2 occurrence of the same variable. We show that there no $(+; +)$ edges in this case. Assume there is a blocking edge. For any ℓ , $T(H(\ell))$ and $F(H(\ell))$ do not contain $(+; +)$ edges. For all gadgets $w_9; w_{10}; \dots; w_{l+8}$ prefer the least men from other gadgets. We conclude that only edge $m_1(\ell_1)w_5(\ell_2)$ could be $(+; +)$, when both ℓ_1 and ℓ_2 are assigned to true. This contradicts to the statement that M satisfy exactly one of (a)-(b) for each pair of positive ℓ_1 and negative ℓ_2 occurrence of the same variable. \square

To a consistent assignment σ , we associate a matching $M = M(\sigma)$ defined as follows: $T(H(\ell)) \in M$ if literal ℓ is set to true, $F(H(\ell)) \notin M$ otherwise; similarly, $T(N(\ell)) \in M$ if literal ℓ is set to true, $F(N(\ell)) \notin M$ otherwise. Conversely, to a stable matching M , because of Lemma 4.5.3, we can associate an assignment σ where a positive (resp. negative) literal ℓ is set to true if and only if $T(H(\ell)) \in M$ (resp. $T(N(\ell)) \in M$).

Corollary 4.5.1. *Let σ be a consistent assignment, then the corresponding matching M is stable. Conversely, let M be a stable matching of $(G(A \sqcup B; E); <)$, then the corresponding assignment σ is consistent.*

Associate to an assignment σ the set $M^\sigma(\ell) = M^\sigma = [\ell \cdot M^\sigma]$, where:

If variable x is set to true, then for each literal ℓ (both positive and negated) corresponding to x , set

$$M^\sigma := \{m_i w_j : i = 1; 2; \dots; 4; 9; \dots; l + 8; m_5 w_6; m_6 w_5\}$$

(dependence on σ is omitted from the nodes);

If variable x is set to false, then for each literal ℓ (both positive and negated) corresponding to x , set

$$M^\ell := \{m_6w_6; m_5w_{l+8}; m_{l+8}w_{l+7}; \dots; m_{10}w_9; m_9w_3; m_3w_1; m_1w_5; m_7w_2; m_8w_4\}$$

(dependence on ℓ is omitted from the nodes).

One easily verifies that M^ℓ is a matching of G and it is maximal.

Lemma 4.5.4. *Let σ be a satisfying assignment. Then $M^\sigma(\sigma)$ is Pareto-optimal.*

Proof. First, we prove that there are no coalition cycles that are contained in one gadget. Consider a matching M^ℓ that corresponds to a variable x that is set to false. In this case the statement follows from the fact that there are only two edges $m_9w_4; m_6w_5 = (+; -)$ that are adjacent to a matched men m_9 and m_6 that prefer w_4 and w_5 to their partners respectively. There is also one edge $m_2w_2 = (+; +)$ adjacent to unmatched node m_2 , so it can not be a part of any coalition cycles. Suppose by contradiction that m_9w_4 is part of a coalition cycle C . Then also $m_8 \succ C$, which is impossible. Similarly, if m_6w_6 is part of a coalition cycle C , then also $m_5 \succ C$, also impossible. Consider a matching M^ℓ that corresponds to a variable x that is set to true. Similarly to above, there are only six edges $m_1w_5; m_3w_2; m_3w_1; m_9w_4; m_9w_3; m_5w_{l+8} = (+; -); (+; +)$ that can be used to form a coalition cycle, but there is no cycle that these edges form.

Now we prove that edges between two literals in different clauses cannot be in any coalition cycle. This follows from the fact that in a matching M^ℓ that corresponds to a variable x that is set to false in positive clause, $m_1w_5 = (-; +)$, and in a matching M^ℓ that corresponds to a variable x that is set to true in negative clause ($\neg x$ is set to false), $m_6w_6 = (-; +)$.

Next, we prove that there are no coalition cycles that are contained in gadgets corresponding to one clause. Consider a positive clause with three literals $\ell_1; \ell_2; \ell_3$

in this order. Without loss of generality assume ℓ_1 is assigned to false, and $\ell_2; \ell_3$ are assigned to true. Assume there is a coalition cycle C and $m_i(\ell_1)w_{i-1}(\ell_1) \in C$. It follows that $m_i(\ell_1)w_{i-1}(\ell_3); m_i(\ell_2)w_{i-1}(\ell_1); m_{i-1}(\ell_3)w_{i-1}(\ell_3); m_i(\ell_2)w_i(\ell_2) \in C$. Since there is no other path between $m_{i-1}(\ell_3)$ and $w_i(\ell_2)$, we obtain a contradiction to C being a cycle. The same argument is applied if there are two literals assigned to false, or if there are two literals in total. Consider a positive clause with three literals $\ell_1; \ell_2; \ell_3$. Since all men in a gadget that corresponds to true assignment $m_9; m_{10}; \dots; m_{l+8}$ are matched to their top choice, there could not be a coalition cycle that includes edges between the gadgets.

Since all women are matched, we conclude that $M^0(\cdot)$ is trade-in-free and maximal. From Lemma 1.3.4 it follows that $M^0(\cdot)$ is Pareto-optimal. \square

Lemma 4.5.5. *Let \mathcal{I} be a satisfying instance. Then $\Delta(G) = 4p_2 + 6p_3$, where p_2 and p_3 are the number of clauses with two and three literals respectively.*

Proof. Consider the Pareto-Optimal matching M^0 and stable matching M corresponding to a satisfying assignment \cdot , as defined above. Now fix a gadget H . By construction, $\Delta_H(M; M^0) = 2$. Hence, $\Delta_C(M; M^0) = 4$ for each clause C with two literals and $\Delta_C(M; M^0) = 6$ for each clause C with three literals. Therefore:

$$\Delta(G) = \Delta(M) + \Delta(M; M^0) = 4p_2 + 6p_3:$$

\square

Lemma 4.5.6. *Let p_2 and p_3 be the number of clauses with two and three literals respectively. If any consistent assignment to \mathcal{I} has at least $k \neq 0$ unsatisfied clauses, then $\Delta(G) \geq k(l+8)$.*

Proof. Based on Corollary 4.5.1 there is a bijection between stable matchings and consistent assignments. Consider any stable matching M and its corresponding as-

segment σ . Since σ leaves k unsatisfied clauses, then there are k clauses that has all literals assigned to false. Consider an unsatisfied positive clause C with three literals $\ell_1; \ell_2; \ell_3$. Edges $m_i(\ell_1)w_{i-1}(\ell_1); m_i(\ell_2)w_{i-1}(\ell_2); m_i(\ell_3)w_{i-1}(\ell_3)$, for $i = 10; 11; \dots; l+8$, form $l+8 - 10 + 1 = l - 1$ disjoint coalition cycles. From Lemma 4.2.4, this implies $\Delta_C(M) \geq l - 1$.

Consider now an unsatisfied negative clause with three literals $\ell_1; \ell_2; \ell_3$. Edges $m_i(\ell_1)w_i(\ell_1); m_i(\ell_2)w_i(\ell_2); m_i(\ell_3)w_i(\ell_3)$, for $i = 9; 11; \dots; l+8$, form $l+8 - 10 + 1 = l - 1$ disjoint coalition cycles, and again using Lemma 4.2.4 we can conclude $\Delta_C(M) \geq l - 1$. A similar argument holds if C has 2 literals instead of 3. Hence, $\Delta_C(M) \geq l - 1$ for any stable matching M and any clause C that is not satisfied in the consistent assignment corresponding to M . Since any consistent assignment leaves at least k clauses unsatisfied, the thesis follows. \square

Using the hardness of approximation for MONOTONE 3-SAT, we can deduce the proof of Theorem 4.1.3.

Proof. Let \mathcal{I} be an instance of the MONOTONE 3-SAT with n variables and $m = O(n)$ clauses, and consider the construction from this section, taking $\ell = \Theta(n)$ with $\epsilon = \frac{1}{8}$. Hence, $|A_j| = \Theta(n^{\ell+1})$.

If \mathcal{I} is satisfiable, then from Lemma 4.5.5 we have:

$$\Delta(G) \leq 4p_2 + 6p_3 < c^\ell |A_j|$$

for an appropriate constant $c^\ell > 0$.

Conversely, if every consistent assignment leaves at least $k = (\frac{1}{8} + \epsilon)m$ clauses of unsatisfied, we deduce:

$$\Delta(G) \geq k(l+8) > c |A_j|$$

for an appropriate constant $c > 0$.

Now suppose that an algorithm A as in the hypothesis of the theorem holds. We construct an algorithm A^ℓ contradicting the hardness result from Theorem 1.4.2.

Indeed, suppose algorithm A accepts. This implies $\Delta(G) < c^j A_j$, hence, we can deduce that there exists an assignment that satisfies at least $(\frac{7}{8} + \epsilon)m$ clauses from and we can make A^ℓ accept. Conversely, if A rejects, then $\Delta(G) \leq c^d A_j$ hence is not satisfiable, and we can make A^ℓ reject. \square

Computing $\Delta(M)$ for a stable matching M is NP-Hard

We showed in Section 4.4 that $\Delta(M)$ for a Pareto-optimal matching M can be computed in polynomial time. We show here that the “symmetric” problem, that is, computing $\Delta(M)$ for a stable matching M is NP-Hard.

Theorem 4.5.1. *Given a marriage instance $(G; <)$ and a stable matching M , computing $\Delta(M)$ is NP-hard.*

In this section we prove Theorem 4.5.1 by reduction from the Maximum independent set problem.

Consider the Maximum independent set problem for a graph G . We construct a corresponding marriage instance in the following way. For each vertex i of G there is an instance $(G_i; <)$ that we will call a block. One block i consists of three men $m_i^0; m_i^1; m_i^2$ and three women $w_i^0; w_i^1; w_i^2$, Fig. 4.8. Their preferences lists are the following:

$$\begin{array}{ll}
 m_i^0 : & w_i^2 \\
 m_i^1 : & w_i^2; w_i^1 \\
 m_i^2 : & w_i^0; w_i^2; w_i^1 \\
 w_i^0 : & m_i^2 \\
 w_i^1 : & m_i^2; m_i^1 \\
 w_i^2 : & m_i^2; m_i^1; m_i^0
 \end{array}$$




Figure 4.8: One block i consists of three men $m_i^0; m_i^1; m_i^2$ and three women $w_i^0; w_i^1; w_i^2$ with their preferences labeled on the edges.

If two vertices i and j from G are adjacent, then we add the following extra edges to the marriage instance: $(m_i^1; w_j^2)$ and $(m_j^1; w_i^2)$. Correspondingly, for each man m_i^1 and m_j^1 add w_j^2 and w_i^2 to the beginning of his preference list respectively. For each women w_i^2 and w_j^2 add m_i^1 and m_j^1 to the end of her preferences list respectively.

Lemma 4.5.7. *There is exactly one stable matching $M_0 = f(m_i^1; w_j^2); (m_j^2; w_i^0) : i = 1 :: jG(V)jg$ in the instance I .*

Proof. Notice that edges $(m_i^1; w_j^2) \not\geq M$ for any $i \neq j$, otherwise $(m_j^0; w_j^2) = (+; +)$. We conclude that $(m_i^2; w_i^0) \geq M$, since w_i^0 is the top choice for m_i^2 . Edge $m_i^0 w_i^2 \not\geq M$, otherwise $m_i^1 w_i^2 = (+; +)$. Since w_i^2 has to be matched, then $(m_i^1; w_i^2) \geq M$. \square

If there are no extra edges, or all vertices in graph G disconnected, then $M_0 = f(m_i^1; w_i^2); (m_i^2; w_i^0) : i = 1 :: jG(V)jg$ is also Pareto-optimal, since there is no coalition cycle. We call a block with edges $(m_i^1; w_i^2); (m_i^2; w_i^0)$ a *stable block*.

Lemma 4.5.8. *If two blocks are connected by extra edges, then at most one of them can be a stable block in a Pareto-optimal matching M_0 .*

Proof. Assume by contradiction there is a Pareto-optimal matching M that contains two stable blocks i and j that are connected. Since w_j^2 and w_i^2 are the top choices for m_i^1 and m_j^1 correspondingly, then there is a coalition cycle $m_i^2; w_j^2; m_j^1; w_i^2$. This contradicts to Pareto-optimality of M . \square

Let S be an independent set. Let I_S be the set of blocks that corresponds to S , and $I_{V \setminus S}$ be the set of blocks that corresponds to $V \setminus S$. Assign $(m_i^1; w_i^2); (m_i^2; w_i^0) \succeq M_1$ for each block from I_S , and $(m_i^0; w_i^2); (m_i^2; w_i^2); (m_i^2; w_i^0) \succeq M_1$ for each block from $I_{V \setminus S}$. Notice that M_1 is a Pareto-optimal matching.

Lemma 4.5.9. *M_1 achieves $\Delta(M_0) = 3 \cdot |V \setminus S|$ if and only if S is the largest independent set.*

Proof. Assume that S is the largest independent set, but there is a Pareto-optimal matching M_2 such that $\Delta(M_0; M_2) < \Delta(M_0; M_1) = 3 \cdot |V \setminus S|$. This would imply that M_2 has more stable blocks and this would correspond to a larger stable set S' , contradiction.

Assume M_1 achieves $\Delta(M_0) = 3 \cdot |V \setminus S|$, but S' is the largest independent set such that $|S'| > |S|$, then S' corresponds to a Pareto-optimal matching M_2 that has more stable blocks than M_1 and would achieve smaller $\Delta(M_0)$ correspondingly. \square

4.6 Proofs of Theorem 4.1.5 and Theorem 4.1.6

In this section, we prove Theorem 4.1.6, which in particular implies Theorem 4.1.5. For the reader's convenience, the theorem is restated below.

Theorem 4.1.6. *There exists a universal constant c such that the following holds. For each $\epsilon > 0$, unless $P=NP$, there is no polynomial-time algorithm that, given a preference system $(G(A \cup B; E); <)$, accepts if $\Gamma = 0$, rejects if $\Gamma \geq c|A|$, and is free to accept or reject if none of these cases apply.*

Because of Lemma 4.2.5, it suffices to show the statement with Δ replaced by Δ_{ϵ} .

For a fixed ϵ , the instance of the marriage problem is constructed from gadgets corresponding to literals and some edges that connect them. The following is the description of a gadget that corresponds to one literal. Positive and negative literals

have the same associated gadgets H , but the true and false assignments correspond to different popular matchings in those gadgets. Nodes of the gadget that correspond to literal ℓ_i are labeled as $a(\ell_i)$, $b(\ell_i)$; $c(\ell_i)$ and $d(\ell_i)$. Sometime we omit the dependency on ℓ_i if it is clear to which gadget the nodes belong.

One gadget consists of two men b ; d and two women a ; c . Fig. 4.9 and Fig. 4.10 provide an example of the construction.




Figure 4.9: One gadget.

If ℓ_i and ℓ_j are literals that correspond to a variable and its negation respectively, then the corresponding gadgets are connected by extra edges $c(\ell_i)b(\ell_j)$ (consistency edges). For woman $c(\ell_i)$, add man $b(\ell_j)$ to the beginning of her preference list, and for a man $b(\ell_j)$, add woman $c(\ell_i)$ to the end of his preference list. If c only has two adjacent edges, then d is the top choice, and b is the second (last) choice in the preference list of c .

For a clause with exactly two literals ℓ_1 ; ℓ_2 , add connectivity edges $d(\ell_1)a(\ell_2)$ and $d(\ell_2)a(\ell_1)$. For men $d(\ell_1)$; $d(\ell_2)$ add women $a(\ell_1)$; $a(\ell_2)$ to the top of their preferences list respectively. For women $a(\ell_1)$; $a(\ell_2)$ add men $d(\ell_1)$; $d(\ell_2)$ to the end of their preferences list respectively. If the clause has three positive literals ℓ_1 ; ℓ_2 ; ℓ_3 , add edges $d(\ell_1)a(\ell_2)$, $d(\ell_2)a(\ell_3)$ and $d(\ell_3)a(\ell_1)$. The preference lists are modified similarly to the case with two literals in a clause. We define graph $G(\cdot)$ as follows. First, construct disjoint graphs $G(\ell)$ for each literal ℓ , then add edges that connect literals in the same clause, and edges that connect literals that correspond to a variable and its negation,

as described above.

Lemma 4.6.1. $G(\varphi)$ is bipartite.

Proof. We assign each vertex from $V := V(G(\varphi))$ to exactly one of two sets A or B , and prove that if $u, v \in A$ (or $u, v \in B$) then $uv \notin E := E(G(\varphi))$. The proof can be followed in Fig. 4.10, where nodes are colored according to the set of the bipartition they belong to. For each literal ℓ assign $b; d$ to A and $a; c$ to B . One easily checks that edges in all gadgets connect nodes in different sets of the bipartition. Now consider $b(\ell)c(\ell^\theta)$, where by construction ℓ (resp. ℓ^θ) is a negative (resp. positive) literal. Hence $b(\ell) \in A$, $c(\ell^\theta) \in B$, as required. \square

To a consistent assignment φ , we associate a matching $M = M(\varphi)$ defined as follows: $T(H(\ell)) \in M$ if literal ℓ is set to true, $F(H(\ell)) \in M$ otherwise, where (dependency on ℓ is omitted):

$$\begin{aligned} T(H) &:= fab; cdg, \\ F(H) &:= fad; bcg. \end{aligned}$$

Claim 4.6.1. An assignment to \mathcal{L} is consistent if and only if it corresponds to a popular matching.

Proof. Assume we are given a consistent assignment that corresponds to a matching M . By the construction all nodes are matched and in any gadget for a literal ℓ either $a(\ell)b(\ell); c(\ell)d(\ell) \in M$ or $a(\ell)d(\ell); b(\ell)c(\ell) \in M$. This implies that $d(\ell)a(\ell^\theta) = (+; \ell)$ for all $\ell; \ell^\theta$. Similarly, $c(\ell)a(\ell^\theta) = (+; \ell)$ for all $\ell; \ell^\theta$.

There is no M -alternating path from an unmatched vertex to a $(+; +)$ edge, since there are no unmatched vertices in M . Note that the only $(+; +)$ edges are all and only those of the form $a(\ell)d(\ell)$ for all ℓ such that ℓ is set to true. Consider any maximal M -alternating path P which contains a certain $(+; +)$ edge $a(\ell)d(\ell)$. P must also contain the subpath $b(\ell)a(\ell)d(\ell)c(\ell)$.




Figure 4.10: A monotone 3-SAT instance $\Phi = (x_1 \wedge x_2 \wedge x_3) \wedge (\neg x_1 \wedge \neg x_2 \wedge x_4) \wedge (x_2 \wedge x_4 \wedge x_5)$ with the corresponding preference system $(G(\Phi); <)$.

Suppose first that φ is a positive literal. For positive literals, $b(\varphi)$ has degree one in G_M , implying that there is no M -alternating cycle in G_M containing $a(\varphi)d(\varphi)$. The only way to continue the M -alternating path P is by extending it through a consistency edge incident to $c(\varphi)$. Suppose therefore that P contains the consistency edge $c(\varphi)b(\varphi_1)$, and that φ_1 is in the clause formed by literals $\varphi_1; \varphi_2; \varphi_3$ in this order. Since φ is set to true and M is consistent, φ_1 is set to false. This implies that P has as subpath $b(\varphi)a(\varphi)d(\varphi)c(\varphi)b(\varphi_1)c(\varphi_1)d(\varphi_1)a(\varphi_1)d(\varphi_3)$. Now we have two possibilities:

1. φ_3 is set to true;
2. φ_3 is set to false, φ_2 is set to true;

3. ℓ_3 is set to false, ℓ_2 is set to false;

In the first case, P continues along $d(\ell_3)c(\ell_3)$ where we cannot continue since $c(\ell_3)$ has degree one in G_M . In the second case there are two possibilities, either P continues along $d(\ell_3)a(\ell_3)b(\ell_3)c(\ell_3)$, or along $d(\ell_3)a(\ell_3)d(\ell_2)c(\ell_2)$. Similar to the first case, the path stops afterwards, without traversing a second $(+;+)$ edge. In the third case there is one additional possibility $d(\ell_3)a(\ell_3)d(\ell_2)a(\ell_2)b(\ell_2)c(\ell_2)$, since $d(\ell_1) \not\subset P$ and $c(\ell_2)$ has degree one in G_M .

Now suppose that ℓ is a negative literal. Here, $c(\ell)$ has degree one in G_M which implies that there is no M -alternating cycle in G_M containing $a(\ell)d(\ell)$. Similar to above, the only way to continue the path P is by extending it through a consistency edge leaving at $b(\ell)$. We may suppose that P contains the consistency edge to the gadget $H(\ell_1)$ and that ℓ_1 is in the clause containing literals $\ell_1; \ell_2; \ell_3$ in this order. Since ℓ is a negative literal which is set to true, we know that $\ell_1; \ell_2; \ell_3$ are positive literals and ℓ_1 is set to false, since M is consistent. This implies that P has as subpath $c(\ell)d(\ell)a(\ell)b(\ell)c(\ell_1)b(\ell_1)a(\ell_1)d(\ell_1)a(\ell_2)$. Again we have two possibilities:

1. ℓ_2 is set to true;
2. ℓ_2 is set to false.

In the first case, P continues along $a(\ell_2)b(\ell_2)$ where we cannot continue since $b(\ell_2)$ has degree one in G_M . In the second case there are two possibilities, either P continues along $a(\ell_2)d(\ell_2)c(\ell_2)b(\ell_2)$, $a(\ell_2)d(\ell_2)a(\ell_3)b(\ell_3)$ or $a(\ell_2)d(\ell_2)a(\ell_3)d(\ell_3)c(\ell_3)b(\ell_3)$ where the path stops since $a(\ell_1) \not\subset P$.

So there is no M -alternating path containing two $(+;+)$ edges and no M -alternating cycle containing a $(+;+)$ edge in G_M . We conclude that M is popular.

Assume we are given a popular matching M . We will show that the corresponding assignment is consistent. Assume that the assignment is not consistent, meaning there are two literals ℓ and ℓ' in a positive and negative gadgets correspondingly that are

both assigned to true. In this case there is a path $a(\cdot)d(\cdot)c(\cdot)b(\cdot^\theta)a(\cdot^\theta)d(\cdot^\theta)$ with two $a(\cdot)d(\cdot); a(\cdot^\theta)d(\cdot^\theta) = (+; +)$. This contradicts popularity. \square

Claim 4.6.2. *All feasible assignments to \mathcal{C} correspond to Pareto-optimal matchings.*

Proof. Assume we are given a feasible assignment that corresponds to a matching M . We need to show that there is no coalition cycle in M to get Pareto-optimality. We can easily check that there is no coalition cycle inside of a gadget. Since b prefer its possible partners in their own gadget to the one at the other end of any consistency edge, the consistency edge is never used in a coalition cycle of M . Assume there is a coalition cycle. It does not contain consistency edges, hence it is contained in the union of gadgets of literals that belong to exactly one clause, and since there are no coalition cycles inside a gadget, it must use all connectivity edges in that clause. But if \cdot is set to true, it is not possible to obtain an M -alternating path in G_M containing both connectivity edges incident to the $H(\cdot)$, since $b(\cdot)c(\cdot) = (\cdot; \cdot)$. Since there is a true literal in every clause, we conclude that there are no coalition cycles in G_M , the thesis follows. \square

Claim 4.6.3. *Let M be a popular matching that corresponds to a consistent assignment to \mathcal{C} with k unsatisfied clauses. Then $\min_{M^0 \text{ Pareto opt}} jM - M^0_j = 4k$.*

Proof. Let M^0 be any Pareto-optimal matching. Let c be a clause that is not satisfied by \cdot , and consider the matching M_c (resp. M_c^0) obtained by restricting matching M (resp. M^0) to the subgraph of G corresponding to gadgets of c . Let $\cdot_1; \cdot_2$ (and possibly \cdot_3) be the literals from c . Next we prove that $jM_c - M_c^0_j = 4$.

Assume all literals in the clause C correspond to all negated variables. Similar proof holds for the case when all literals correspond to positive variables. Notice, that if $a(\cdot_i)d(\cdot_i) \not\geq M_c$ for all $i = 1; 2; 3$, then there is a coalition cycle $a(\cdot_1)d(\cdot_1)a(\cdot_2)d(\cdot_2)a(\cdot_3)d(\cdot_3)$. We conclude that there is at least one edge

$a(i)d(i) \not\subseteq M_c^0$ for some i . Assume without loss of generality that $i = 1$. Next, we will consider all possible Pareto-optimal matchings by analyzing the following cases: $c(1)d(1); d(1)a(2) \subseteq M_c^0$, $c(1)d(1) \subseteq M_c^0$ and $d(1)a(2) \subseteq M_c^0$.

If $c(1)d(1); d(1)a(2) \subseteq M_c^0$, then $b(1)c(1) \subseteq M_c^0$, otherwise node $d(1)$ prefers unmatched node $c(1)$, which is a contradiction to Pareto-optimality. Similarly, we argue that $a(1)d(3) \subseteq M_c^0$, and either $a(3)b(3) \subseteq M_c^0$ or $a(3)d(2) \subseteq M_c^0$. In both cases $jM_c - M_c^j > 4$, since $a(1)d(1); a(3)d(3)$, $b(3)c(3) \subseteq M_c^0$ and $a(1)d(3)$, $a(3)b(3) \subseteq M_c^0$.

If $c(1)d(1) \subseteq M_c^0$, then either $b(1)a(1) \subseteq M_c^0$ or $a(1)d(3) \subseteq M_c^0$, since otherwise node $b(1)$ prefers unmatched node $a(1)$, which is a contradiction to Pareto-optimality. In the first case $jM_c - M_c^j = 4$, since $a(1)d(1); b(1)c(1) \subseteq M_c^0$ and $a(1)b(1); c(1)d(1) \subseteq M_c^0$. In the second case $jM_c - M_c^j > 4$, since $a(1)d(1); b(1)c(1)$, $a(3)d(3) \subseteq M_c^0$ and $c(1)d(1); a(1)d(3) \subseteq M_c^0$.

If $d(1)a(2) \subseteq M_c^0$, then, similarly to the argument above, either $a(1)b(1) \subseteq M_c^0$ or $b(1)c(1)$, $a(1)d(3) \subseteq M_c^0$. $jM_c - M_c^j = 4$, since $a(1)d(1); b(1)c(1)$, $a(2)d(2) \subseteq M_c^0$ and $a(1)b(1)$, $d(1)a(2) \subseteq M_c^0$, or $a(1)d(1)$, $a(3)d(3) \subseteq M_c^0$ and $d(1)a(2)$, $a(1)d(3) \subseteq M_c^0$.

In all cases the arguments hold even if there are only two literals per clause.

This implies that for all unsatisfied clauses there must be at least four edges changed from either being matched to unmatched or vice versa, in order to obtain a Pareto-optimal matching. Notice that $(M_c; M_c^0) = 4$, when the corresponding assignment of the literal changes parity, so the corresponding clause becomes satisfied. If for each unsatisfied clause of \mathcal{C} we change parity for exactly one literal, the obtained assignment is feasible, since all clauses are satisfied (but not necessarily consistent). According to Claim 4.6.2, this feasible assignment corresponds to a Pareto-optimal matching M^0 , and $jM - M^j = 4k$. \square

Claim 4.6.4. Let $M; M^0$ achieve the minimum in $\min_{M \text{ popular}, M^0 \text{ Pareto opt}} jM - M^0j$. Then every consistent assignment has at least $k = \frac{\min_{M^0 \text{ Pareto opt}} jM - M^0j}{4}$ unsatisfied clauses.

Proof. By contradiction. Assume that there is a consistent assignment such that the number of unsatisfied clauses is less than k . By Claim 4.6.1, this assignment corresponds to a popular matching M . Change the parity for exactly one literal in each unsatisfied clause as to obtain a feasible (possibly not consistent) assignment, and let M^0 be the corresponding matching. According to Claim 4.6.2 M^0 is Pareto-optimal. Then $\Delta(M; M^0) < 4k = \frac{\min_{M^0 \text{ Pareto opt}} jM - M^0j}{4} = \min_{M \text{ popular}, M^0 \text{ Pareto opt}} jM - M^0j$, a contradiction. \square

Take $c = \frac{1}{12}$, and suppose that an algorithm A^0 as in the hypothesis of the theorem exists. We show that we can provide an algorithm A^0 as in the hypothesis of Lemma 1.4.2, a contradiction (assuming $P \neq NP$). Feed $(G(A; B; E); <)$ to A^0 . If A^0 accepts, then $\min_{M \text{ popular}, M^0 \text{ Pareto opt}} jM - M^0j < (c + \frac{1}{4})jAj$. By Claim 4.6.3, there is a consistent assignment that satisfies

$$m \frac{\min_{M \text{ popular}, M^0 \text{ Pareto opt}} jM - M^0j}{4} < m \frac{c}{4}jAj \quad (1 - \frac{3}{2}c + \frac{3}{2})m > (1 - \frac{3}{2}c + \frac{1}{2})m = (\frac{7}{8} + \frac{1}{2}c)m$$

clauses. We can then make A^0 accept. Suppose conversely that A^0 rejects. Then $\min_{M \text{ popular}, M^0 \text{ Pareto opt}} jM - M^0j > 0$, hence $(1 - \frac{3}{2}c + \frac{1}{2})m > 0$ is not satisfiable (by Claim 4.6.4), and we can make A^0 reject.

*Combinatorial Scheduling for Adaptive Machine Learning
in Cybersecurity*

5.1 Introduction

Cyber-security is an important problem in today's world. Enormous amounts of computations across different industries rely on secure operation of networked computing devices. These systems contain information that can be used for malicious intentions such as identity theft, damage of reputation and business operations, loss of valuable resources, loss of intellectual property and even compromise national security. With the advances of cryptography algorithms the cyber-attacks become more elaborate as well. These attacks exploit the smallest flaws in the system and are guided by intelligent and adaptive adversaries with resources. They can last for long periods of time while remaining undetected and progress slowly by learning the system and its design.

Corporations and governments spend tremendous amounts of resources to avoid data breaches and system infiltration to protect their information, as it's compromise can lead to serious and costly consequences. The development of modern cryptography algorithms reflects the need of design that incorporates the possibility of data being partially compromised. Moving target defense (MTD) strategies are one of the techniques that attempt to mitigate complicated attacks on a system.

In [51] authors introduce game *FlipIt* to model attacks that compromise a system or take control of a resource completely, but are not necessarily revealed immediately.

In this model there are two players, the attacker and the defender. They are both interested in gaining control of a single resource that in real-life could be a password, design of a system, cryptographic key etc. Players can take a move at any given time that allows them to immediately gain control of the resource. Attacker and defender are not aware of their opponents moves and who controls the resource except at the time when they make a move. Each move has a cost. Players are interested in maximizing their utility which is the amount of time they controlled the resource minus the total costs associated with the moves. This motivates the game’s second name, *game of stealthy takeover*. There are multiple papers ([39], [38], [44]) that explore FlipIt game and its extensions.

In Fig. 5.1 it is shown a schematic diagram of an instance of FlipIt game.




Figure 5.1: An example of FlipIt game. Defender controls the resource in the beginning, then the attacker makes a move and the defender loses the control. Flags indicate the moves of the players. Red and blue colors show time when the defender and attacker control the resource respectively.

Moving target defense (MTD) strategies allow to mitigate sophisticated attacks by incorporating randomness or making a static part of a system dynamic. For example, certain parts of the system might be required to restart to a known, not compromised state, such as change of a password. The additional cost required to obtain needed information from such systems might prevent the adversary to start the attack. In [42] authors provide survey on MTD strategies. In [30] authors develop a strategic game, PLADD, which is an abstract model for MTD. PLADD stands for Probabilistic Learning Attacker, Dynamic Defender. This model incorporates key characteristics that are shared among different categories of MTD strategies [30]:

1. There exists some information that provides competitive advantage to the attacker when its in his/her possession.
2. The defender can take this information from the attacker at least temporary.

The inspiration for PLADD comes from FlipIt game, but there are several important differences [30]:

1. The attacker does not gain the control of the resource immediately after the attack. There is a random time-to-success wait period.
2. Time-to-success incorporates the ability of the attacker to learn. Its distribution changes over time.
3. The defender can make two types of moves that are accomplished without delays. The first move "take" allows to gain the control of the resource. The second move "morph" allows to gain the control of the resource and take information away from the attacker (stop the ongoing attack). This allows the attacker to detect both defender's moves when the attacker controls the resource (the attacker knows when they lose control) and to detect a morph move even if the defender controls the resource, since morph move restarts the system. After the morph the first attacker move comes from f_{base} distribution, and the following attacks are drawn from f_{learned} distribution until the next morph move. A take move does not restart the system, so it does not reset the distribution to f_{base} . At any time only one attack can be in progress. Fig. 5.2 shows an instance of PLADD game.

The attacker's and defender's goal is to maximize their utility which is total time in control minus total cost associated with moves. The defender has two costs C_{take} and C_{morph} associated with the corresponding moves. The attacker has a fixed cost to start an attack and a variable cost during the attack. The variable cost is proportional the duration of the attack.

The analysis of PLADD consists of deriving strategies of when and what type of moves the defender should make.

In [30] authors analyzed PLADD, derived closed-form solutions in multiple game variants and developed simulation of the game that verified their analytical solution. They also formulated a stochastic programming model and proposed another approach for analyzing PLADD, a new combinatorial scheduling problem. In this work we analyze the proposed latter approach.




Figure 5.2: This diagram illustrates an example of PLADD game. Small flags indicate take moves and large flags indicate morph moves. The defender starts the game with a morph move. At the same time the attacker starts the attack that does not succeed immediately but have time-to-success wait time drawn from f_{base} distribution. During this wait time the defender makes two take moves that are wasted and do not have any influence on the system. After that the attacker gains the control of the resource until the defender places a take and regains the control of the system. The attacker knows that he lost control and immediately starts a new attack. This time the wait-to-success time is drawn from $f_{learned}$ distribution, since the attacker already acquired some knowledge about the system from the previous attack. The defender places a take that does not influence the game. The next defender’s move, morph, cancels the attack and reset the distribution of wait-to-success time to base. The game resets. The attacker lost acquired knowledge and starts a new attack from f_{base} distribution.

5.2 The PLADD game and a new scheduling problem

In this section we present a new scheduling problem for PLADD analysis [30]. Since a morph move completely resets the system by setting the time-to-success distribution to f_{base} (all prior attacker’s knowledge about the system is deleted), we will analyse

the game between two successive morph moves with fixed number of takes. Let T be the game's time horizon that equals to the time between two morph moves. Assume there are m scenarios, and each scenario is a list of t_{k+1} draws such that the first draw is from f_{base} distribution, and the consecutive k draws are from $f_{learned}$ distribution. The goal is to schedule k takes withing T time such that the attacker control time is minimized over all scenarios. This approach allows to find a strategy for a defender that minimizes on average the expected attacker control time.

Each scenario can be modelled as a machine or a server. Each such machine has a list of jobs with fixed length that must be processed in order. At time $t_0 = 0$ all machines start their first jobs. If there is a start time (a take in a PLADD scenario), a machine can only start the next job if it finished its current job (the machine was idle), otherwise the machine continues processing its current job. The problem is to find a schedule of start times, or synchronization points, $t_i; i = 1; \dots; k$, that minimizes total idle time on all machines, where $k + 1$ is the length of the longest list on a machine. Without loss of generality, we assume that the total processing time of all jobs on any machine is at most T . In order to achieve this we can truncate the last job or drop jobs from the end of the list. This implies that on any machine either the total processing time of all jobs equals to T , or number of jobs is $k + 1$, or both. Fig. 5.3 provides an example of a scheduling problem with three machines.

If there is one machine, then the optimal solution is a schedule where each job starts immediately the previous job finishes, or all start times are placed at the end of jobs. For multiple machines there is an optimal solution such that all synchronization points are scheduled at the end of some jobs that do not necessary belong to one machine, [30]. This statement follow from the observation that if there is a schedule such that all machines idle t time units before a start time, then scheduling this synchronization point and all following it t time units earlier might allow to schedule more job at the end, or provide additional processing time for jobs that were not

finished.




Figure 5.3: This diagram illustrates an example of PLADD game. There are three servers. Different jobs on one machine are colored differently for visualization purpose. On all servers j th jobs are the same color. For example, first and second jobs on all servers have red and blue colors respectively. All job lengths are measured in units of time that is needed to process them. The first server has 9 jobs of lengths 1, 4, 1, 1, 4, 1, 1, 4, 1. The second server has 9 jobs of lengths 2, 3, 1, 2, 3, 1, 2, 3, 1. The third server has 6 jobs, each of length 3. There are 9 takes (vertical black lines) and $T = 18$ (vertical red line). The first take t_0 is scheduled at time $t = 0$. The rest of the takes are scheduled at times 1, 2, 5, 6, 8, 12, 13, 15. The idle time is represented by white color between jobs on one server. Jobs after time T shown in dark color to illustrate that they were not processed. The diagram shows the optimal schedule for this instance with the total idle time of 6 units.

5.3 Integer Programming Formulation

In this section we present an integer programming formulation (IP) for the problem. The time scale is discretized as the set $\{0, \dots, T\}$. Any job that does not start in $\{0, \dots, T-1\}$ is permitted to start at time T without additional take and without satisfying precedence constraints, since we are only concerned about the idle time before time T and are not interested about the schedule after this time horizon. All n_m jobs for each machine $m \in \{1, \dots, M\}$ are indexed in order they must be processed. The length of job j on machine m is $\ell_{m,j}$ for $j \in \{1, \dots, n_m\}$ and $m \in \{1, \dots, M\}$. Let $L(m; i; j) = \sum_{k \in [i; j]} \ell_{m,k}$ be the total sum of lengths of all jobs with indices i, \dots, j on machine m . $L_m = L(m; 1; n_m)$ is the total length of all the jobs on machine m . Let $\Delta_{j,m}(t) = \{s \in \mathbb{Z}^+ : s + \ell_{m,j} \leq t\}$ be a set of start times s for job j on a machine m such that if j starts at time s , then it still runs at time t . The earliest possible start

time for job j on machine m is $ET(m;j) := L(m;1;j - 1)$. The binary variables for the formulation are the following:

$x_{m;j;t} = 1$ indicates that job j starts at time t on the machine m and 0 otherwise,
 $y_t = 1$ indicates that jobs are allowed to start at time t , i.e. a take places at time t ,

$z_{m;t} = 1$ indicates that the machine m is idle at time t .

The variables $x_{m;j;t}$ are defined for $j \in [1; n_m]$, $m \in [1; M]$ and $t \in [ET(m;j); T]$. The variables y_t and $z_{m;t}$ are defined for $t \in [0; T]$. We are given k takes (not including the first take at start time $t = 0$). The objective is to minimize idle time:

$$\min \sum_{m \in [1; M]} \sum_{t \in [0; T]} z_{m;t}$$

We require that every job is scheduled:

$$\sum_{t \in [ET(m;j); T]} x_{m;j;t} = 1; \forall m; j: \tag{5.1}$$

The first jobs on all machines start at time $t = 0$:

$$x_{m;1;0} = 1; \forall m: \tag{5.2}$$

We also require that at most one job is in process at any time $[0; T)$. Recall that at time T multiple jobs are allowed to be scheduled, and time T is not included in measuring idle time. If all $x_{m;j;s}$ in the summation below are zeros, then $z_{m;t}$ must be 1, indicating that the server is idle.

$$z_{m;t} + \sum_j \sum_{s \in m_j(t)} x_{m;j;s} = 1; \forall t < T \text{ and } \forall m: \tag{5.3}$$

We impose precedence constraints to ensure that for a job i , all jobs with a greater index j must start after all jobs between it and i have completed. For an IP formulation the precedence constraints only for $j = i + 1$ are needed. However, in Section 5.4 we provide an example that illustrates that including the constraints for $j > i + 1$ result in a stronger LP relaxation.

$$x_{m;j;t} \sum_{s \in [t+L(m;j);T]} x_{m;j+1;s} \delta_{m;j;t} \leq ET(m;j) \quad t < T: \quad (5.4)$$

Precedence constraints (5.4) with constraints for $j > i + 1$:

$$x_{m;i;t} \sum_{s \in [t+L(m;i;j-1);T]} x_{m;j;s} \delta_{m;i;j;t} \leq ET(m;j) \quad t < T; j > i: \quad (5.5)$$

The following constraints are complementary to (5.4) and (5.5) respectively. They are not necessary for an IP formulation. They ensure that for a job i , each job with a smaller index j must end before i starts and leave enough time for processing jobs between j and i . Similarly to above, for an IP formulation the precedence constraints only for for $j = i - 1$ are needed.

$$x_{m;j+1;t} \sum_{s \in [ET(m;j);t-L(m;j)]} x_{m;j;s} \delta_{m;j;t} \leq ET(m;j) \quad t < T: \quad (5.6)$$

Precedence constraints (5.6) with constraints for $j < i + 1$:

$$x_{m;i;t} \sum_{s \in [ET(m;j);t-L(m;j;i-1)]} x_{m;j;s} \delta_{m;i;j;t} \leq ET(m;j) \quad t < T; j < i: \quad (5.7)$$

Constraints (5.4) - (5.7) can be replaced by the following constraints (5.8) - (5.11). We replace the left-hand side with an appropriate sum. This results in a stronger LP relaxation as shown in Section 5.4.

$$\sum_{r \in [t; T]} x_{m; i; r} \quad \sum_{s \in [t + l_{m; j}; T]} x_{m; j+1; s} \quad \delta_{m; j; t} : ET(m; j) \quad t < T: \quad (5.8)$$

$$\sum_{r \in [t; T]} x_{m; i; r} \quad \sum_{s \in [t + L(m; j - 1); T]} x_{m; j; t} \quad \delta_{m; i; j; r} : ET(m; j) \quad t < T; j > i: \quad (5.9)$$

$$\sum_{r \in [ET(j+1); s]} x_{m; j+1; r} \quad \sum_{s \in [ET(m; j); t - l(m; j)]} x_{m; j; s} \quad \delta_{m; j; t} : ET(m; j) \quad t < T: \quad (5.10)$$

$$\sum_{r \in [ET(i); s]} x_{m; i; r} \quad \sum_{s \in [ET(m; j); t - L(m; j; i - 1)]} x_{m; j; s} \quad \delta_{m; i; j; t} : ET(m; j) \quad t < T; j < i: \quad (5.11)$$

We ensure that jobs may only start if a take is placed. Number of takes is budgeted:

$$\sum_j x_{m; j; t} \quad y_t \quad \delta_{m; t} : ET(m; j) \quad t < T: \quad (5.12)$$

$$\sum_{t \in [0; T]} y_t \quad k + 1: \quad (5.13)$$

At time $t = 0$ all machines start their first jobs. This implies the first take is always placed at $t = 0$.

$$y_0 = 1: \quad (5.14)$$

5.4 Linear Programming Relaxation

In this section we explain the meaning of variables in linear programming relaxation (LP) and give some examples that illustrate how redundant constraints for IP

strengthen LP. The solutions to IP and LP were found using Python implementation with Gurobi Optimizer. For each example we provide input data and the output for variables in IP or LP formulations. The input data consists of list *length*, time horizon T and number of takes k (without the take at time $t = 0$). List *length* contains lists that correspond to servers and contain lengths of jobs in order. The output data contains variables z , x and y . We will slightly change notation of variables to resemble that of Python. Variables $z[m][t]$ are presented in a table where the rows correspond to servers and columns correspond to time line $t = [0; \dots; T)$. Variable $x[m][j][t]$ are presented for each server separately in a table, where each row corresponds to a job in order $0; \dots; n_m - 1$ and columns correspond to time line $t = [0; \dots; T)$. Variables $y[t]$ are given in a vector, each entry correspond to a value of y at time $t = [0; \dots; T)$. All number are rounded to two decimals.

Example 5.4.1. *The IP formulation includes constraints (5.1) - (5.3), (5.9), (5.11), (5.12) - (5.14).*

Input:

length = [[1, 4, 1, 1, 4], [2, 3, 2, 2, 2], [3, 3, 3, 2]]

T = 11

K = 4

Output for IP

$z[m][t]$

[0][1][1][0][0][0][0][1][1][0][0]

[0][0][1][0][0][0][0][0][1][0][0]

[0][0][0][0][0][0][0][0][0][0][0]

$x[1][j][t]$

[1][0][0][0][0][0][0][0][0][0][0]

[0][0][0][1][0][0][0][0][0][0][0]
[0][0][0][0][0][0][0][0][0][1][0]
[0][0][0][0][0][0][0][0][0][0][1]
[0][0][0][0][0][0][0][0][0][0][0]

$x[2][j][t]$

[1][0][0][0][0][0][0][0][0][0][0]
[0][0][0][1][0][0][0][0][0][0][0]
[0][0][0][0][0][0][1][0][0][0][0]
[0][0][0][0][0][0][0][0][0][1][0]
[0][0][0][0][0][0][0][0][0][0][1]
[0][0][0][0][0][0][0][0][0][0][0]

$x[3][j][t]$

[1][0][0][0][0][0][0][0][0][0][0]
[0][0][0][1][0][0][0][0][0][0][0]
[0][0][0][0][0][0][1][0][0][0][0]
[0][0][0][0][0][0][0][0][0][1][0]

$y[t]$

[1][0][0][1][0][0][1][0][0][1][1]

Output for LP

$z[m][t]$

[0.00][0.67][0.33][0.00][0.00][0.00][0.00][0.67][0.33][0.00][0.00]
[0.00][0.00][0.67][0.00][0.00][0.00][0.00][0.00][0.33][0.00][0.00]
[0.00][0.00][0.00][0.33][0.33][0.00][0.00][0.00][0.00][0.00][0.00]

$x[1][j][t]$

[1.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00]
 [0.00][0.33][0.33][0.33][0.00][0.00][0.00][0.00][0.00][0.00][0.00]
 [0.00][0.00][0.00][0.00][0.00][0.33][0.33][0.00][0.00][0.33][0.00]
 [0.00][0.00][0.00][0.00][0.00][0.00][0.33][0.00][0.33][0.00][0.33]
 [0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.33][0.00][0.33][0.00]

x[2][j][t]

[1.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00]
 [0.00][0.00][0.33][0.67][0.00][0.00][0.00][0.00][0.00][0.00][0.00]
 [0.00][0.00][0.00][0.00][0.00][0.33][0.67][0.00][0.00][0.00][0.00]
 [0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.33][0.33][0.33][0.00]
 [0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.33][0.33]

x[3][j][t]

[1.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00]
 [0.00][0.00][0.00][0.67][0.00][0.33][0.00][0.00][0.00][0.00][0.00]
 [0.00][0.00][0.00][0.00][0.00][0.00][0.67][0.00][0.33][0.00][0.00]
 [0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.00][0.67][0.00]

y[t]

[1.00][0.33][0.33][0.67][0.0][0.33][0.67][0.33][0.33][0.67][0.33]

Fig. 5.4 and Fig. 5.6 illustrate the optimal solutions for IP and LP respectively. Fig. 5.5 also give another optimal solution for IP.

The total idle time for IP is 6 time units, for LP the total idle time is $3\frac{2}{3}$ time units. Notice that in LP takes can have fractional values. A fractional take of value p at time t means that p fraction of a new job can start at time t if the server has capacity. For example, in Fig. 5.6, at time $t = 0$ all servers started rst jobs, take




Figure 5.4: Optimal schedule for example 5.4.1. At time $t = 0$ all jobs are scheduled, the next take is placed at time $t = 3$, etc.




Figure 5.5: Optimal schedule for example 5.4.1. At time $t = 0$ all jobs are scheduled, the next take is placed at time $t = 1$, etc.




Figure 5.6: Optimal solution for LP for example 5.4.1. Takes are fractional.

at $t = 0$ is 1. At time $t = 1$ the take is $\frac{1}{3}$. Since the first server is idle at $t = 1$ then it starts $\frac{1}{3}$ of the second job (blue color), and continue processing $\frac{1}{3}$ of it during its total length of 4 time units. At time $t = 3$ the take is $\frac{2}{3}$, so all servers can not start more than $\frac{2}{3}$ fraction of some jobs. The first server has only time left to start $\frac{1}{3}$ of the next job, so it starts $\frac{1}{3}$ of the second job (blue color), while servers 2 and 3 have start processing $\frac{2}{3}$ of the second jobs, even though server 3 has $\frac{1}{3}$ of idle time left.

Example 5.4.2.

Input:

length = [[1, 4, 1, 1, 4], [2, 3, 2, 2, 2], [3, 3, 3, 2]]

T = 11

K = 4

Optimal schedule has idle time of 6 units.

The LP formulation includes constraints (5.1) - (5.3), (5.9), (5.11), (5.12) - (5.14).

Optimal has idle time of $3\frac{2}{3}$ units.

The LP formulation includes constraints (5.1) - (5.3), (5.5), (5.7), (5.12) - (5.14).

Optimal has idle time of 3.52 units.

We conclude that using constraints (5.9) and (5.11) strengthen LP formulation.

Example 5.4.3.

Input:

length = [[1, 3, 1, 1, 5], [3, 3, 3, 2]]

T = 11

K = 4

Optimal schedule has idle time of 3 units.

The LP formulation includes constraints (5.1) - (5.3), (5.5) and (5.7), (5.12) - (5.14). Optimal solution has idle time of $2\frac{1}{3}$ units.

The LP formulation includes constraints (5.1) - (5.3), (5.4) and (5.6), (5.12) - (5.14). Optimal solution has idle time of 2.01 units.

We conclude that using constraints (5.5) and (5.7) strengthen LP formulation.

5.5 Integrality Gap

In this section we construct an example for which the integrality gap is an unbounded function of number of jobs on a server. We employ the notations of the IP formulation

with constraints (5.1) - (5.3), (5.9), (5.11) - (5.14).

Consider the following PLADD scheduling problem with two servers. First server has $3n$ jobs, each of length 2 units. Second server has $2n$ jobs, each of length 3 units. Global time T is $6n$ units, and there are $3n$ takes including the take at time 0. Notice that the total length of jobs on any server equals to T .

Lemma 5.5.1. *Optimal LP solution for the proposed problem has idle time of at most 1 unit for $n \geq 2$.*

Proof. Consider the following LP solution:

$$y_0 = 1; y_1 = 0; \text{ and } y_i = 1=2 \text{ for } i = 2; \dots; T - 1;$$

$$x_{0,0,0} = 1; x_{0,i,2i} = 1=2; x_{0,i,2i+1} = 1=2 \text{ for } i = 1; \dots; 3n - 1;$$

$$x_{1,0,0} = 1; x_{1,i,3i} = 1=2; x_{1,i,3i+1} = 1=2 \text{ for } i = 1; \dots; 2n - 1;$$

$$z_{0,2} = 1=2 \text{ and } z_{0,3} = 1=2;$$

All other variables equal to 0. Recall that $x_{i,j,t}$ indicates a fraction of job j that starts on server i at time t . All constraints of LP are satisfied, hence the solution is feasible.

Fig. 5.7 provides the schedule for $n = 2$.




Figure 5.7: For each server blocks of the same color represent the same job. Each server is divided into two parts for visualization of fractional takes. Vertical black lines represent takes, and vertical red line represents the global time T

Total idle time is 1 unit. □

Lemma 5.5.2. *Optimal IP solution for the proposed problem has idle time of n units.*

Proof. Any schedule can be represented as a union of the blocks in Fig. 5.8, since blocks of types $x; y; z$ and t represent all possible combinations of placing takes and not repeating the existing combination. Blocks $x; y; z; t$ can be repeated many times in any order in a schedule, and blocks $x_1; x_2; y_1; y_2; y_3; z_1; z_2; z_3; z_4; t_1; t_2; t_3; t_4$ represent all possible combinations for the last block, if it is not $x; y; z; t$. Note that block y_2 can be represented by x_1 , since we only concern about the schedule before the global time T (red vertical line); similarly $y_3; z_3; z_4; t_3; t_4$ are represented by $x_2; y_1; x_1; y_1; x_1$ respectively.

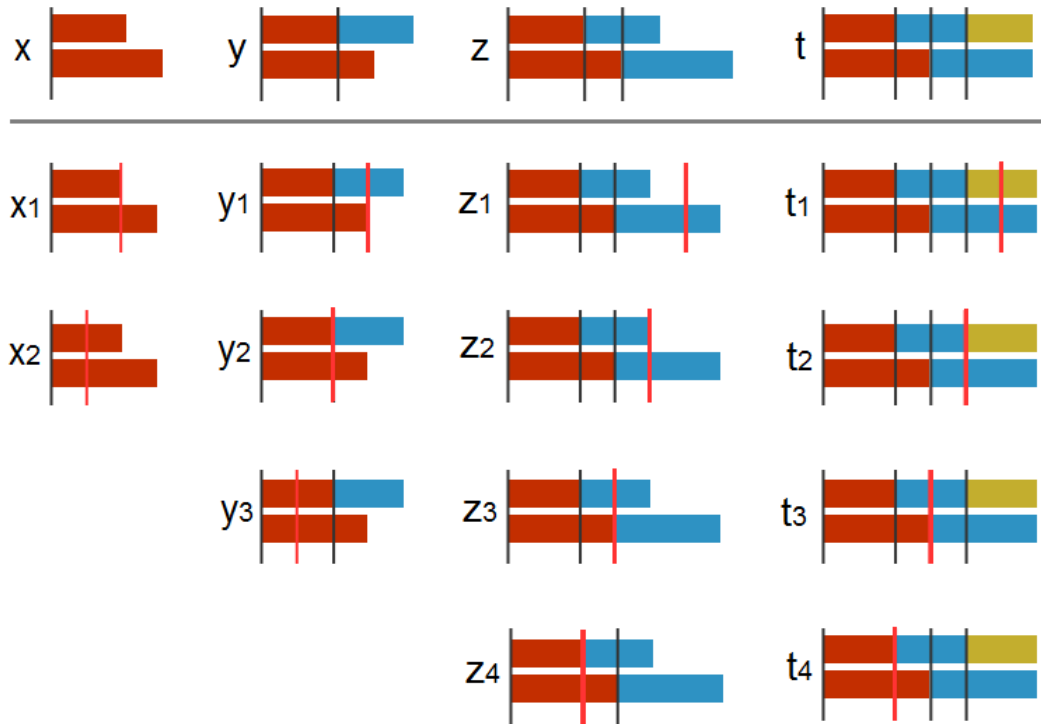


Figure 5.8: Any schedule can be represented as a union of the blocks above. Red vertical line represent global time T

The following is the IP formulation for the problem with $n = 2k$. Each variable represents number of blocks of certain type from Fig. 5.8. For example, variable x equals to the number of blocks of type x in a schedule.

The objective is to minimize idle time:

$$\min x + y + 2z + 0t + 0x_1 + 0x_2 + 0y_1 + 1z_1 + 0z_2 + 0t_1$$

Total number of takes equals $6k$ (for example, in block x there is only one take used):

$$x + 2y + 3z + 4t + x_1 + x_2 + 2y_1 + 3z_1 + 3z_2 + 4t_1 + 3t_2 = 6k$$

Total length of all blocks equals $T = 12k$:

$$3x + 4y + 6z + 6t + x_1 + 2x_2 + 3y_1 + 5z_1 + 4z_2 + 5t_1 + 4t_2 = 12k$$

We can use at most one block from $x_1; x_2; y_1; z_1; z_2; t_1; t_2$:

$$x_1 + x_2 + y_1 + z_1 + z_2 + t_1 + t_2 \leq 1$$

All variables represent number of blocks of a certain type in a schedule:

$$x; y; z; t; x_1; x_2; y_1; z_1; z_2; t_1; t_2 \geq 0$$

$$x; y; z; t; x_1; x_2; y_1; z_1; z_2; t_1; t_2 \geq 0$$

Consider LP relaxation of IP. The following is a basic feasible solution: $x = 2k - 1 = 3; t = k - 1 = 6; x_2 = 1$, all other variables are zero, since there are exactly 12 (we add one slack variable for the inequality to transform the problem into canonical form) linearly independent active constraints. Reduced costs vector corresponding to $(x; y; z; t; x_1; x_2; y_1; z_1; z_2; t_1; t_2; s)$ is $(0; 1 = 3; 1; 0; 2 = 3; 0; 1 = 3, 1; 2 = 3; 1; 2 = 3; 1 = 3)$. We conclude that this vertex is optimal solution to the LP. Optimal value of the objective function is $2k - 1 = 3$. Notice that $Opt_{LP} = Opt_{IP}$ and the cost vector is integer, hence

optimal value for the IP is bounded below by $2k$.

The following schedule achieves the lower bound $2k$ units of idle time. Consider an instance of the problem with $n = 2k$ and create $2k$ block of type t and $2k$ blocks of type x . Blocks of type x have one unit of idle time each. Total idle time is $2k$ units. Fig. 5.9 shows optimal schedules for $n = 2; 4$.

Similar argument can be applied for odd values of n . The optimal schedule for any n has n units of idle time. □





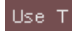
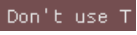


Figure 5.9: The top and bottom schedules are optimal for the problems with $n = 2$ and $n = 4$

5.6 Visualization Software

The Schedule Simulator (see Fig. 5.10) allows the visualization and calculation of optimal configurations for small instances. A user can interact with the schedule by dragging and dropping synchronization points. It provides tools for creating custom scenarios where a user can add/delete jobs from the end or servers from the bottom using buttons  and  correspondingly. In order to change a job length a user needs to select the job and enter its length in .

Buttons  and  switch between problems with and without global time T correspondingly.

 and  buttons store a snapshot of a previously considered example.

 and  corresponds to cancel/recover previous versions.

There is an option between Brute and Greedy algorithms (heuristics). The first one considers all possible cases and chooses the one with the smallest total idle time.

Save and **Load** provides a tool to save and load scenarios.

Export CSV saves the idle times of all finite cases explored by the optimization algorithms.

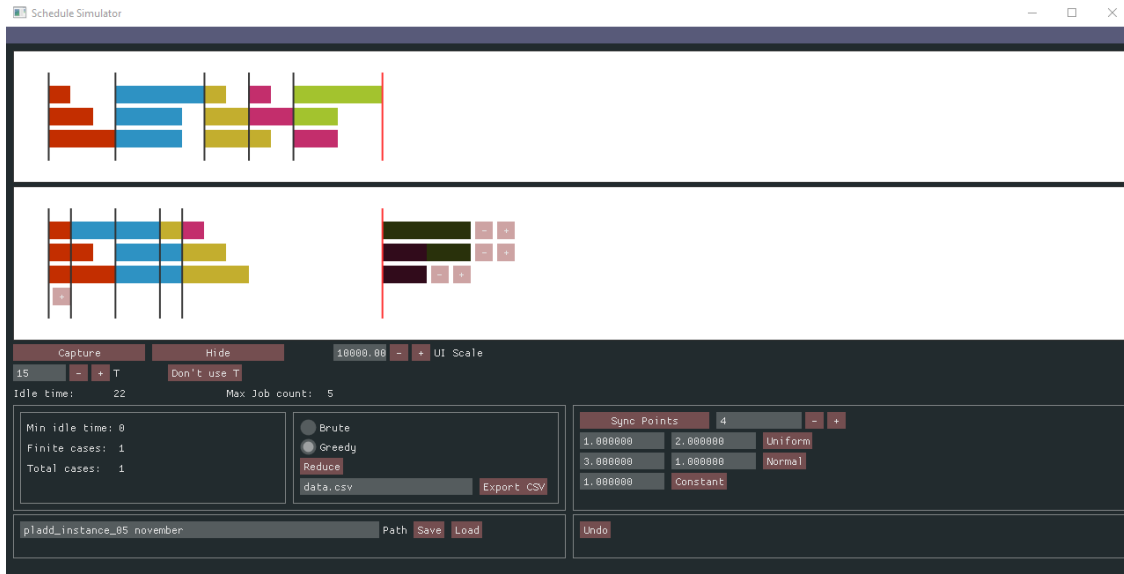


Figure 5.10: A screenshot of Schedule Simulator application

The source code and the built application can be found on GitHub at <https://github.com/VladenaPowers/PLADD>.

Chapter 6

Conclusions

In this thesis we studied several problems arising in popular matchings and scheduling, proving positive and negative complexity results, investigating the effectiveness of current tools for attacking those problems, as well as providing new ones.

In Chapter 2, we settle the complexity of many open problems around popular matchings in bipartite graphs [12, 13, 41, 28, 33]. In particular, we developed a framework that allows to show NP-hardness and inapproximability results for problems on popular and stable matchings, as well as providing new, and we believe simpler, proofs for recently established problems in the literature. This framework is based on a reduction from a variant of the classical 3-SAT.

In Chapter 4, we develop this framework further. First, we introduce a certain concept of distance between the sets of stable (resp. popular) and Pareto-optimal matchings. We show that hardness of approximation results for (a variant of) 3-SAT can be employed to show that this distance is, in general, very hard to approximate. We complement those results showing special cases in which it can be computed exactly, by exploiting the lattice structure of stable matchings.

We believe that future work in this area could expand our framework further, and apply it to more problems on popular matchings, and more generally to matchings under preferences.

In Chapter 3, we show that some problems around popular matchings are solvable in polynomial time if we restrict the input graph to have bounded treewidth. The latter is a classical assumption, often turning intractable problems into tractable ones.

However, for popular matchings, classical techniques do not seem to apply, and we have to develop a suitable generalization of the concept of popularity, that we called *local popularity*. The algorithm we propose is, for the moment, of theoretical interest only. However, we believe that this result could be the starting point for applying the concept of popular matchings to real-world problems whose underlying graphs has small treewidth. Indeed, stable matchings are found in many real-world applications, and replacing them with popular matching (or subclasses of those) could increase the cardinality of the output matching, while preserving a certain notion of global stability.

Last, in Chapter 5, we analyze a combinatorial formulation for PLADD game introduced in [30]. This game represents an abstract model for moving target defense strategies that allows to mitigate sophisticated cyber attacks by incorporating randomness or changing static parts of the system into dynamic.

The problem results in finding a schedule that minimizes the total idle time for a given number of servers and ordered lists of jobs on each of them. We propose a natural integer programming formulation where variables represent jobs and takes, and we count minimize total idle time. We study multiple instances that allowed us to identify and strengthen certain parts of the formulation. The linear relaxation of this formulation introduces fractional values for the variables, for which we provided intuition and meaning. We prove that the integrality gap of all those formulations is unbounded. The software *The Schedule Simulator* was developed to find an optimal schedule for small instances by enumerating all possible scenarios and visualizing any schedule for analysis. It is available on GitHub at <https://github.com/VladenaPowers/PLADD>. It also allows to build custom instances in an interactive manner and experiment with different heuristics.

Future work on this problem could potentially include a development of a heuristic that would result in an approximation algorithm with a constant approximation ratio,

analysis of the integrality gap for a maximization version of the problem, or different integer programming formulations. This new scheduling problem could also lead to new insights in other areas of research, such as adaptive machine learning.

Bibliography

- [1] A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3): 689–701, 1998.
- [2] A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American economic review*, 93(3):729–747, 2003.
- [3] A. Abdulkadiroğlu, P. A. Pathak, and A. E. Roth. Strategy-proofness versus efficiency in matching with indifference: Redesigning the nyc high school match. *American Economic Review*, 99(5):1954–78, 2009.
- [4] D. J. Abraham, K. Cechlárová, D. F. Manlove, and K. Mehlhorn. Pareto optimality in house allocation problems. In *International Symposium on Algorithms and Computation*, pages 3–15. Springer, 2004.
- [5] D. J. Abraham, A. Blum, and T. Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 295–304, 2007.
- [6] G. Aggarwal, S. Muthukrishnan, D. Pál, and M. Pál. General auction mechanism for search advertising. In *Proceedings of the 18th international conference on World wide web*, pages 241–250, 2009.
- [7] P. Biró, R. W. Irving, and D. F. Manlove. Popular matchings in the marriage and roommates problems. In *International Conference on Algorithms and Complexity*, pages 97–108. Springer, 2010.
- [8] L. Bodin and A. Panken. High tech for a higher authority: The placement of graduating rabbis from hebrew union college—jewish institute of religion. *Interfaces*, 33(3):1–11, 2003.
- [9] H. L. Bodlaender. A tourist guide through treewidth. *Acta cybernetica*, 11(1-2): 1, 1994.
- [10] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on computing*, 25(6):1305–1317, 1996.
- [11] H. K. Büning and T. Lettmann. *Propositional logic: deduction and algorithms*, volume 48. Cambridge University Press, 1999.

- [12] Á. Cseh. Popular matchings. *Trends in Computational Social Choice*, 105, 2017.
- [13] Á. Cseh and T. Kavitha. Popular edges and dominant matchings. *Mathematical Programming*, 172(1-2):209–229, 2018.
- [14] Á. Cseh, Y. Faenza, T. Kavitha, and V. Powers. Understanding popular matchings via stable matchings. *arXiv preprint arXiv:1811.06897*, 2018.
- [15] Y. Faenza and X. Zhang. Legal assignments and fast eadam with consent via classical theory of stable matchings. *arXiv preprint arXiv:1809.08506*, 2018.
- [16] Y. Faenza, V. Powers, and X. Zhang. Two-sided popular matchings in bipartite graphs with forbidden/forced elements and weights. *arXiv preprint arXiv:1803.01478*, 2018.
- [17] Y. Faenza, T. Kavitha, V. Powers, and X. Zhang. Popular matchings and limits to tractability. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2790–2809. SIAM, 2019.
- [18] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [19] P. Gärdenfors. Match making: assignments based on bilateral preferences. *Behavioral Science*, 20(3):166–173, 1975.
- [20] S. Gupta, P. Misra, S. Saurabh, and M. Zehavi. Popular matching in roommates setting is np-hard. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2810–2822. SIAM, 2019.
- [21] V. Guruswami and S. Khot. Hardness of max 3sat with no mixed clauses. In *20th Annual IEEE Conference on Computational Complexity (CCC'05)*, pages 154–162. IEEE, 2005.
- [22] D. Gusfield and R. W. Irving. *The stable marriage problem: structure and algorithms*. MIT press, 1989.
- [23] J. Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.
- [24] M. Hirakawa, Y. Yamauchi, S. Kijima, and M. Yamashita. On the structure of popular matchings in the stable marriage problem—who can join a popular matching. In *the 3rd International Workshop on Matching under Preferences (MATCH-UP)*, 2015.
- [25] G. J. Hitsch, A. Hortaçsu, and D. Ariely. Matching and sorting in online dating. *American Economic Review*, 100(1):130–63, 2010.
- [26] G. J. Hitsch, A. Hortaçsu, and D. Ariely. Matching and sorting in online dating. *American Economic Review*, 100(1):130–63, 2010.

- [27] C.-C. Huang and T. Kavitha. Popular matchings in the stable marriage problem. *Information and Computation*, 222:180–194, 2013.
- [28] C.-C. Huang and T. Kavitha. Popularity, mixed matchings, and self-duality. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2294–2310. SIAM, 2017.
- [29] R. W. Irving, P. Leather, and D. Gusfield. An efficient algorithm for the “optimal” stable marriage. *Journal of the ACM (JACM)*, 34(3):532–543, 1987.
- [30] S. Jones, A. Outkin, J. Gearhart, J. Hobbs, J. Siirola, C. Phillips, S. Verzi, D. Tauritz, S. Mulder, and A. Naugle. Evaluating moving target defense with pladd. *Sandia National Laboratories*, 2015.
- [31] K. Joshi and S. Kumar. Matchmaking using fuzzy analytical hierarchy process, compatibility measure and stable matching for online matrimony in india. *Journal of Multi-Criteria Decision Analysis*, 19(1-2):57–66, 2012.
- [32] T. Kavitha. A size-popularity tradeoff in the stable marriage problem. *SIAM Journal on Computing*, 43(1):52–71, 2014.
- [33] T. Kavitha. Popular half-integral matchings. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- [34] T. Kavitha. Max-size popular matchings and extensions. *arXiv preprint arXiv:1802.07440*, 2018.
- [35] T. Kavitha. The popular roommates problem. *arXiv preprint arXiv:1804.00141*, 2018.
- [36] J. Kennes, D. Monte, and N. Tumennasan. The day care assignment: A dynamic matching problem. *American Economic Journal: Microeconomics*, 6(4):362–406, 2014.
- [37] O. Kesten. School choice with consent. *The Quarterly Journal of Economics*, 125(3):1297–1348, 2010.
- [38] A. Laszka, B. Johnson, and J. Grossklags. Mitigating covert compromises. In *International Conference on Web and Internet Economics*, pages 319–332. Springer, 2013.
- [39] A. Laszka, G. Horvath, M. Felegyhazi, and L. Buttyán. Flipthem: Modeling targeted attacks with flipit for multiple resources. In *International Conference on Decision and Game Theory for Security*, pages 175–194. Springer, 2014.
- [40] B. M. Maggs and R. K. Sitaraman. Algorithmic nuggets in content delivery. *ACM SIGCOMM Computer Communication Review*, 45(3):52–66, 2015.

- [41] D. Manlove. *Algorithmics of matching under preferences*, volume 2. World Scientific, 2013.
- [42] H. Okhravi, M. Rabe, T. Mayberry, W. Leonard, T. Hobson, D. Bigelow, and W. Streilein. Survey of cyber moving target techniques. Technical report, MASSACHUSETTS INST OF TECH LEXINGTON LINCOLN LAB, 2013.
- [43] M. Ostrovsky. Stability in supply chain networks. *American Economic Review*, 98(3):897–923, 2008.
- [44] V. Pham and C. Cid. Are we compromised? modelling security assessment games. In *International Conference on Decision and Game Theory for Security*, pages 234–247. Springer, 2012.
- [45] M. J. Rosenfeld, R. J. Thomas, and S. Hausen. Disintermediating your friends: How online dating in the united states displaces other ways of meeting. *Proceedings of the National Academy of Sciences*, 116(36):17753–17758, 2019.
- [46] A. E. Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of political Economy*, 92(6):991–1016, 1984.
- [47] A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *The Quarterly journal of economics*, 119(2):457–488, 2004.
- [48] A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005.
- [49] M. A. Satterthwaite and H. Sonnenschein. Strategy-proof allocation mechanisms at differentiable points. *The Review of Economic Studies*, 48(4):587–597, 1981.
- [50] L. Shapley and H. Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.
- [51] M. Van Dijk, A. Juels, A. Oprea, and R. L. Rivest. Flipit: The game of “stealthy takeover”. *Journal of Cryptology*, 26(4):655–713, 2013.