TWO PROBLEMS IN DISCRETE MATHEMATICS

By

JOHN CHIARELLI

A dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Mathematics

Written under the direction of

Michael Saks

And approved by

New Brunswick, New Jersey

October, 2020

ABSTRACT OF THE DISSERTATION

Two Problems in Discrete Mathematics

By JOHN CHIARELLI

Dissertation Director: Michael Saks

This thesis centers around two projects that I have undertaken in the subject of discrete mathematics. The primary project pertains to the stable matching problem, and puts particular focus on a relaxation of stability that we call S-stability. The secondary project looks at boolean functions as polynomials, and seeks to understand and use a complexity measure called the maxonomial hitting set size.

The stable matching problem is a well-known problem in discrete mathematics, with many practical applications for the algorithms derived from it. Our investigations into the stable matching problem center around the operation $\psi : E(G(I)) \to E(G(I))$; we show that for sufficiently large k, ψ_I^k maps everything to a set of edges that we call the hub, and give algorithms for evaluating $\psi_I(S)$ for specific values of S. Subsequently, we extend results on the lattice structure of stable matchings to S-stability and consider the polytope of fractional matchings for these same weaker notions of stability. We also reflect on graphs represented by instances with every edge in the hub.

Given a boolean function $f : \{0,1\}^n \to \{0,1\}$, it is well-known that it can be represented as a unique multilinear polynomial. We improve a result by Nisan and Szegedy on the maximum number of relevant variables in a low degree boolean polynomial using the maxonomial hitting set size, and look at the largest possible maxonomial hitting set size for a degree d boolean function.

Acknowledgements

I would like to thank Dr. Michael Saks for being my mentor, teaching me about the way research in mathematics works, and helping me learn how to ask questions that I want to find the answers to. During my work with Dr. Saks, I was funded in part by the Simons Foundation under award 332622, and I would like to thank them for their aid. Last but not least, I would also like to thank Dr. Doron Zeilberger for instilling in me an appreciation for experimental techniques, without which I would never have found many of my results.

Table of Contents

Chapter 1

Introduction

This thesis centers around two major projects that I have undertaken in the subject of discrete mathematics. The first project, discussed in Chapters 3-7, pertains to the stable matching problem, and puts particular focus on a relaxation of the notion of stability that we call S-stability. The second project, which appears in Chapters 8 and 9, looks at the properties of boolean functions as polynomials.

The stable matching problem is a well-known problem in discrete mathematics; a full background on the topic appears in Chapter 2. Our investigations into the stable matching problem center around the operation $\psi : E(G(I)) \to E(G(I))$, which we establish a framework for in Chapter 3 and define in Chapter 4; we show that for sufficiently large k, ψ_I^k maps everything to a set of edges that we call the hub, and give algorithms for evaluating $\psi_I(S)$ for specific values of S. In later chapters, we extend results on the lattice structure of stable matchings to the discussed weaker notions of stability (Chapter 5) and consider the polytope of fractional matchings for these same weaker notions (Chapter 6). We also reflect on graphs represented by instances with every edge in the hub (Chapter 7).

It is well known that any boolean function $f : \{0,1\}^n \to \{0,1\}$ can be represented as a unique multilinear polynomial. In Chapter 8, we consider a sensitivity measure that we call the maxonomial hitting set size, and apply it in order to improve a result by Nisan and Szegedy on the maximum number of variables in a low degree boolean polynomial (Theorem 8.1).[1] In Chapter 9, we focus on expanding our understanding of the largest possible maxonomial hitting set size for a degree d boolean function.

---

[1] This section was previously published on ArXiv in 2018, and was published in Combinatorica earlier this year [CHS20].

Chapter 2

Background on the Stable Matching Problem

In this chapter, we review the stable matching problem, which will be the central focus for most of this thesis. We focus particular attention on the structure of the lattice of stable matchings. The results of this section are not original to our work, though some of the notation is our invention. We refer the reader to the excellent book by Dan Gusfield and Robert W. Irving [GI89].

2.1 Stable Matchings and the Domination Ordering

In the stable matching problem, an $n \times n$ **instance** I contains n mentors $V_m = \{m_1, \ldots, m_n\}$ and n students $V_d = \{d_1, \ldots, d_n\}$; in addition, each individual is associated with a **preference list** - an ordered list of a subset of the individuals of the opposite type.[1] (We can think of this list as representing the partners that individual would accept, listed in order of desirability.) In general, we assume that m_i is on d_j's preference list iff d_j is on m_i's preference list. (If one individual won't accept another one as a partner, whether the second individual would accept the first is a moot point.) An individual v **prefers** v' to v'' if v' appears no later than v'' in v's preference list; in addition, every v prefers every v' on its preference list to v. (We also say that v **strictly prefers** v' to v'' if v prefers v' to v'' and $v' \neq v''$.) The **graph of the instance** $G(I)$ is the bipartite graph with $V(G(I)) = V_m \cup V_d$ such that $(m_i, d_j) \in G(I)$ iff m_i and d_j are in each other's preference list. (Since $G(I)$ is bipartite with parts equal to V_m and V_d, every edge e can be described as (m_e, d_e), where $m_e \in V_m$ and $d_e \in V_d$.) An instance I is **complete** if $G(I)$ is the complete bipartite graph between $V_m(I)$ and $V_d(I)$.

[1] Historically, the two types of individual were commonly referred to as men and women respectively. We feel that this terminology reflects an outdated notion of gender dynamics and elect to use a more socially conscientious notation.

A **matching** M is a subgraph of $G(I)$ where every vertex has degree at most 1 - in this case, every vertex in M with degree 1 has a **partner**, the vertex it is adjacent to in M. We can also describe a matching via the function $p_M : V(G(I)) \to V(G(I))$, where $p_M(v) = v$ if v has degree 0 in M, and is v's partner in M otherwise. A matching M is perfect if it is 1-regular - i.e. $p_M(v) \neq v$ for all $v \in V(G(I))$.

An edge $e \in E(G(I))$ **destabilizes** M if m_e prefers d_e to $p_M(m_e)$ and d_e prefers m_e to $p_M(d_e)$; a matching M is **stable** if no $e \in E(G(I))$ destabilizes M. While it is not immediately obvious that a stable matching exists over an arbitrary instance, David Gale and Lloyd Shapley showed that complete instances always have at least one perfect stable matching, and discovered an algorithm - the Gale-Shapley algorithm - that could find one such stable matching.

Algorithm 2.1. *Given a stable matching instance I, we construct a matching as follows:*

1. *Set $p_M(v) = v$ for all $v \in V(G(I))$. At any point, each $m \in V_m(I)$ can be in either of two states, frustrated or not frustrated; we initially set every $m \in V_m(I)$ to be not frustrated.*

2. *While there exists some $m \in V_m(I)$ such that m is not frustrated and $p_M(m) = m$, do the following*

   (a) *Select any such m. If m has proposed to every member of their preference list, m becomes frustrated; otherwise, m proposes to the first elements of their preference list that they has not yet proposed to.*

   (b) *When m proposes to d, if d prefers m to $p_M(d)$, then $p_M(m)$ becomes d and $p_M(d)$ becomes m. If $p_M(d)$ was previously some other $m' \in V_m(I)$, then $p_M(m')$ becomes m'.*

Theorem 2.2. *For any instance I, any execution of Algorithm 2.1 outputs the same stable matching M. ([GS62])*

(As we note in Theorem 2.7, Algorithm 2.1 gives the mentors much greater influence over the output matching than the students, giving each mentor their best possible partner in a stable matching and each student their worst. We can find an algorithm that produces the matching that gives each student their best possible partner and each mentor their worst by switching the roles of the mentors and students. Later work, such as in [GIL87], found alternate algorithms that were type-neutral - that is to say, switching the roles of the mentors and the students would not change the output of the algorithm.)

Theorem 2.3. *If I is a complete $n \times n$ instance, then the matching M created by Algorithm 2.1 is a perfect stable matching. ([GS62],Theorem 1).*

A stable matching is not necessarily perfect; however, as shown by Gale and Marilda Sotomayor ([GS85]), a vertex v is unmatched in a stable matching over an instance I iff it is unmatched in every stable matching over that instance. We refer to an instance I as **satisfactory** if every stable matching over I is perfect.

Theorem 2.4. *Every stable matching covers the same vertices. ([GS85], Theorem 1)*

Corollary 2.5. *An instance I is satisfactory if there exists a perfect stable matching over I.*

Theorem 2.6. *For all $n \in \mathbb{N}$, every complete $n \times n$ instance is satisfactory.*

We note that a given instance can be considered to have other, "smaller" instances within it. A **restriction** $I[S]$ of the instance I to $S \subseteq E(G(I))$ is the instance on the same vertex set such that the preference list of every vertex in $I[S]$ is the order-preserving sublist of its vertex list in I where, for any $v_1, v_2 \in V(G(I))$, v_1 appears on v_2's preference list iff $(v_1, v_2) \in S$. A particularly noteworthy type of restriction is a **truncation** - a restriction created by iteratively selecting a vertex and removing the final element of that vertex's preference list. We can construct any truncation of I by taking a subset $V \subseteq V(G(I))$ and selecting, for each $v \in V$, a minimum acceptable partner $a(v)$ from v's preference list, and for all v' such that v strictly prefers $a(v)$ to v', we remove v and v' from each other's preference lists. We write this truncation as $I_{(T_d,T_m)}$, where $T_m = \{(m, a(m)) : m \in V_m(I) \cap V\}$ and $T_d = \{(a(d), d) : d \in V_m(I) \cap V\}$.

2.2 The Lattice of Stable Matchings

For an instance I, the set $\mathcal{L}_s = \mathcal{L}_s(I)$ of all stable matchings over I has a natural partial order given by $M \preceq M'$ iff every mentor m prefers $p_M(m)$ to $p'_M(m)$, and every student d prefers $p_{M'}(d)$ to $p_M(d)$. We say that M **dominates** M' if $M \preceq M'$, and refer to $\preceq$ as the **domination ordering**. We refer to a stable matching over I as **mentor-optimal** if it dominates every other stable matching, and **student-optimal** if every other stable matching dominates it. (It is trivial to see that there can be at most one mentor-optimal and one student-optimal stable matching.)

Theorem 2.7. *Given an instance I, Algorithm 2.1 generates the unique mentor-optimal stable matching over I. ([GS62], Theorem 2)*

In particular, Theorem 2.7 implies that for any instance I, there exists a unique mentor-optimal stable matching (and similarly a unique student-optimal stable matching) over I. We may also contemplate the idea of "combining" two stable matchings to create another stable matching such that some number of vertices are given their preferred partners among the two input matchings. This intuition prompted the following theorems, which Donald E. Knuth ([Knu76]) attributed to John H. Conway.

Theorem 2.8. *Let M_1 and M_2 be two stable matchings over I. Then, the following hold:*

- *There exists a unique matching $M_1 \wedge_d M_2$ such that each student is matched with their preferred partner among their partners in M_1 and M_2, and $M_1 \wedge_d M_2$ is stable.*

- *There exists a unique matching $M_1 \wedge_m M_2$ such that each mentor is matched with their less preferred partner among their partners in M_1 and M_2, and $M_1 \wedge_m M_2 = M_1 \wedge_d M_2$.*

- *There exists a unique matching $M_1 \vee_m M_2$ such that each mentor is matched with their preferred partner among their partners in M_1 and M_2, and $M_1 \vee_m M_2$ is stable.*

- *There exists a unique matching $M_1 \vee_d M_2$ such that each student is matched with their less preferred partner among their partners in M_1 and M_2, and $M_1 \vee_d M_2 = M_1 \vee_m M_2$.*

([Knu76], p. 87-88)

We refer to the matchings $M_1 \wedge_d M_2$ and $M_1 \vee_m M_2$ as $M_1 \wedge M_2$ and $M_1 \vee M_2$ respectively. It is not trivial that the matchings $M_1 \wedge M_2$ and $M_1 \vee M_2$ are stable, or that they even exist. The proof of this depends on both M_1 and M_2 being stable, and on the fact thatthe operations of $\wedge$ and $\vee$ have their domains limited to pairs of stable matchings. (In Chapter 3, we will look at how we could naturally expand the domains of these operations.) As observed in [Knu76], the domination ordering forms a lattice with meet and join operations given by Theorem 2.8. Furthermore, it is easy to show that $\wedge$ and $\vee$ distribute over one another, and therefore:

Theorem 2.9. *Given two stable matchings M, M' over I, M dominates M' iff every $d \in V_d(I)$ prefers $p_{M'}(d)$ to $p_M(d)$. In addition, the poset $\mathcal{L}_s$ of the stable matchings with the domination ordering forms a distributive lattice. ([Knu76], p. 87-92)*

Furthermore, Charles Blair ([Bla84], Theorem 1), answering a question posed by Knuth ([Knu76], p. 92) showed that for every distributive lattice $\mathcal{L}$, there exists an instance I such that the resulting lattice $(\mathcal{L}_s, \preceq)$ is isomorphic to $\mathcal{L}$. An algorithm to create such an instance with relatively few vertices is given by Dan Gusfield, Robert W. Irving, Paul Leather, and Michael Saks ([GILS87], Section 2.2). (We present it as Algorithm 5.11.)

2.3 Rotations Over the Stable Matchings

We want to better understand the structure of the distributive lattice $\mathcal{L}_s(I)$ associated with an instance I. As a first step, we review the well-known Birkhoff Representation Theorem, which allows us to express the elements of a distributive lattice in terms of its join-irreducible elements.

A distributive lattice $\mathcal{L}$ has a least element and greatest element, which we represent as $\hat{0}_{\mathcal{L}}$ and $\hat{1}_{\mathcal{L}}$ respectively. (In cases where $\mathcal{L}$ is implied, we shorten these to $\hat{0}$ and $\hat{1}$.)

We say that an element $l \in \mathcal{L}$ is **join-irreducible** if for any subset of elements $L \subseteq \mathcal{L}$ such that $\vee_{j \in L} j = l$, $l \subseteq L$. Since the join of the empty set is $\hat{0}$, $\hat{0}$ is not join-irreducible, despite the fact that it cannot be expressed as the join of any number of elements $\prec \hat{0}$; in fact, $\hat{0}$ is the unique $l \in \mathcal{L}$ that is not join-irreducible such that if $|L| = 2$ and $\vee_{j \in L} j = l$, then $l \in L$. Similarly, l is **meet-irreducible** if for any subset of elements $L \subseteq \mathcal{L}$ such that $\wedge_{j \in L} j = l$, $l \subseteq L$; $\hat{1}$ is the unique $l \in \mathcal{L}$ that is not meet-irreducible such that if $|L| = 2$ and $\wedge_{j \in L} j = l$, then $l \in L$.

Theorem 2.10. *Given a distributive lattice $\mathcal{L}$ with partial order $\preceq$, let J be the poset of the join-irreducible elements of $\mathcal{L}$. Then, there exists an isomorphism κ from $\mathcal{L}$ to the downsets of J, such that for all $l \in \mathcal{L}$, κ maps l to $\hat{0} \cup \{j \in J : j \preceq l\}$. ([Bir37], Theorem 5)*

Thus, a distributive lattice is completely determined by its poset of join irreducible elements. We want to apply this to better understand the lattice $\mathcal{L}_s(I)$ of stable matchings. To do this, we want to provide an explicit way to describe the poset of join irreducibles. The key to this is the concept of a **rotation**.

Let I be an instance, and $\mathcal{C}_s(I)$ be the set of pairs (M, M') of stable matchings where M' covers M in $\mathcal{L}_s(I)$. We define a rotation over I to be a pair $\rho = (\rho_m, \rho_d)$ with $\rho_m, \rho_d \subseteq E(G(I))$ such that there is a pair $(M, M') \in \mathcal{C}_s(I)$ such that $\rho_m = M - M'$ and $\rho_d = M' - M$. Note that for a rotation ρ, ρ_m and ρ_d are matchings in $G(I)$ that cover the same vertices.

Theorem 2.11. *Let ρ be a rotation over I. Then, there exists an $r \in \mathbb{N}$, a sequence $\{m_1, \dots, m_r\} \subseteq V_m(I)$, and a sequence $\{d_1, \dots, d_r\} \subseteq V_d(I)$ such that $\rho = (\{(m_1, d_1), (m_2, d_2), \dots (m_r, d_r)\}, \{(m_1, d_2), \dots, (m_{r-1}, d_r), (m_r, d_1)\})$. ([GI89], Theorem 2.5.3)*

For any $i \in [r]$, m_i prefers $p_{\rho_m}(m_i) = p_M(m_i)$ to $p_{\rho_d}(m_i) = p_{M'}(m_i)$ and d_i prefers $p_{\rho_d}(d_i) = p_{M'}(d_i)$ to $p_{\rho_m}(d_i) = p_M(d_i)$, so $\rho_m \prec \rho_d$. We say that a vertex $v \in V(G(I))$ is in a rotation ρ if there exists some vertex $v' \in V(G(I))$ such that $(v, v') \in \rho_m$. (We note by Theorem 2.11 that this occurs iff there exists some vertex $v'' \in V(G(I))$ such that $(v, v'') \in \rho_d$.)

If we have a stable matching M over I, we can consider the truncation $I_{(M,\emptyset)}$, created by deleting from I all edges (m, d) such that d strictly prefers $p_M(d)$ to m. We note that the edges deleted this way include all edges such that m strictly prefers d to $p_M(m)$ - otherwise, (m, d) would destabilize M. Therefore, M matches each mentor with their top choice in $I_{(M,\emptyset)}$ and each student with their bottom choice. We say that M **exposes** a pair of matchings (ρ_m, ρ_d) over I that cover the same vertices if $\rho_m \subseteq M$ and, for each mentor $m \in \rho$, ρ_d matches m with their second choice in $I_{(M,\emptyset)}$. In particular, we see that every pair of matchings exposed by some stable matching is a rotation.

Proposition 2.12. *If M exposes (ρ_m, ρ_d) over I, then (ρ_m, ρ_d) is a rotation over I and $M \cup \rho_d - \rho_m$ is a stable matching over I that covers M in $\mathcal{L}_s(I)$. ([GI89], Theorem 2.5.1)*

The following lemmas show the converse, that every rotation is exposed by some stable matching.

Lemma 2.13. *If ρ is a rotation over I, then there exists a stable matching M such that M exposes ρ. ([GI89], Theorem 2.5.3)*

Lemma 2.14. *Given an instance I, let $\{M_0, M_1, \ldots, M_k\}$ be any maximal chain in $\mathcal{L}_s(I)$. (Note that this implies that M_0 is the mentor-optimal stable matching over I, and M_k is the student-optimal stable matching over I.) Then, $\{(M_{i-1} - M_i, M_i - M_{i-1}) : i \in [k]\}$ is the set of all rotations over I. ([GI89], Theorem 2.5.4)*

Corollary 2.15. *Given an instance I, let $M_0, M_1, \ldots, M_r$ be any maximal chain in $\mathcal{L}_s(I)$ such that M_0 and M_r are the mentor-optimal and student-optimal stable matchings over I respectively. Then, the set of all edges that appear in some stable matching over I is $\cup_{i=0}^{r} M_i$. ([Gus87], Theorem 2)*

Proposition 2.16. *Let M, M' be a pair of stable matchings and ρ be a rotation over I such that $\rho = (\{e \in E(G(I)) : e \in M, e \notin M'\}, \{e \in E(G(I)) : e \notin M, e \in M'\}$. Then, M' covers M. ([GI89], Theorem 2.4.2)*

The above theorems show us that we can consider a more compact form of describing the lattice of stable matchings - namely, through its rotations. For a given instance I, we define the **rotation poset** of I to be the poset on the set of rotations of I with the partial order that $\rho \leq \rho'$ iff for every stable matching M over I and $m \in V_m(I)$, m either prefers $p_{\rho_d}(m)$ to $p_M(m)$ or prefers $p_M(m)$ to $p_{\rho_m}(m)$. We represent the rotation poset of I by $\Pi(I)$. The following theorem gives a more explicit description of the order relation of $\Pi(i)$.

Theorem 2.17. *For a given stable matching instance I, let $\mathcal{R}$ be the digraph such that $V(\mathcal{R})$ is the set of all rotations over I, and (ρ, ρ') is an edge in $\mathcal{R}$ iff at least one of the following holds:*

- $\rho_d \cap \rho'_m \neq \emptyset$.
- *There exists a mentor $m_0 \in \rho'$ and a student $d_0 \in \rho$ such that (m_0, d_0) does not appear in any stable matching over I and, in I, m_0 prefers $p_{\rho'_m}(m_0)$ to d_0 to $p_{\rho'_d}(m_0)$ and d_0 prefers $p_{\rho_d}(d_0)$ to m_0 to $p_{\rho_m}(d_0)$.*

Then, $\Pi(I)$ is the transitive closure of $\mathcal{R}$. ([Gus87], Theorem 4 and Lemma 6)

Theorem 2.18. *Let ν be the map from the downsets of $\Pi(I)$ to the stable matchings over I such that for any downset $D \in \Pi(I)$, $\nu(D) = M_0 \cup (\cup_{\rho \in D} \rho_d) - (\cup_{\rho \in D} \rho_m)$. Then, ν is an isomorphism, and for any stable matching M, $\nu^{-1}(M)$ is the set of all rotations ρ such that, for all $(m, d) \in \rho$, m strictly prefers d to $p_M(m)$. ([IL86], Theorem 5.1)*

One major advantage of representing the lattice of stable matchings through the rotation poset is its compact nature. The entire lattice of stable matchings over an $n \times n$ instance could potentially be superpolynomial in terms of n. The lattice of stable matchings over an $n \times n$ instance can have size exponential in n ([IL86], Corollary 2.1). However, the number of rotations is at most $O(n^2)$ (since the elements of $\{\rho_d : \rho \in \Pi(I)\}$ are disjoint by Lemma 2.14, and each ρ_d contains at least two edges in $G(I)$). Therefore, the rotation poset provides a compact representation of the lattice $\mathcal{L}_s(I)$. The more compact representation of $\mathcal{L}_s(I)$ afforded by $\Pi(I)$ allows us to perform certain computational tasks far more efficiently - as seen in the following theorems, as well as in Chapter 5.

Theorem 2.19. *Given an $n \times n$ instance I, we can construct the rotation poset of I in $O(n^2)$ time. ([Gus87], Theorem 5)*

Corollary 2.20. *Given an $n \times n$ instance I, there exists an algorithm that determines, in $O(n^2)$ time, the set of all edges in $G(I)$ that appear in a stable matching over I. ([Gus87], Theorem 3)*

2.4 The vNM-Stable Matchings

A significant portion of this thesis is dedicated to weakenings of stability. A weakening that is of particular interest here is von Neumman-Morgenstern stability, or vNM-stability, which was studied in [Ehl07], [Wak08], and [Wak10].

Given a stable matching instance I, a set of matchings $\mathcal{M}$ is **vNM-stable over** I if it satisfies the following conditions:

- For all $M_1, M_2 \in \mathcal{M}$ and $v \in V_m \cup V_d$, at least one of v and $p_{M_1}(v)$ prefers their partner in M_2 to the other.
- For all $M \notin \mathcal{M}$, there exists an $M' \in \mathcal{M}$ and $v \in V_m \cup V_d$ such that v and $p_{M'}(v)$ strictly prefer each other to their respective partners in M.

The first major result on vNM-stable sets we need is attributed to Lars Ehlers.

Theorem 2.21. *If $\mathcal{M}$ is a vNM-stable set of matchings over an instance I, then $(\mathcal{M}, \preceq)$ is a distributive lattice, and every stable matching over I appears in $\mathcal{M}$. ([Ehl07], Theorem 2))*

It is not clear from the definition that every instance has a vNM-stable set of matchings, and if it does, whether it is unique. This was established by Jun Wako, who showed:

Theorem 2.22. *For any instance I, there exists a unique vNM-stable set of matchings over I. ([Wak08], Theorem 5.1)*

This allows us to talk about *the* vNM-stable set of matchings of an instance I. We say that a matching is vNM-stable over I if it belongs to the vNM-stable set of matchings.

It is clear from the second part of the definition that a vNM-stable set of matchings must contain all stable matchings, and so vNM stability is a weakening of stability. If I is an instance where every edge is in a stable matching, then stability and vNM-stability coincide.

The proof that appears in Wako is based on a construction that follows this outline:

1. Initially, let C_0 be the set of all stable matchings over I and set $n = 0$.

2. Let UD^n be the set of all matchings that are not destabilized by any edge that appears in a matching in C_0.

3. If $C_n \subsetneq UD^n$, find C_{n+1}, the set of all stable matchings over $I[\cup_{M \in UD^n} M]$. Return to step 2 with $n := n + 1$. If $C_n = UD^n$, then C_n gives the unique vNM-stable set of matchings.

Wako was able to make the final assertion in the above construction by the following lemma.

Lemma 2.23. *If $C_n = C_{n+1}$ for any $n \in \mathbb{N}$, then $C_n = UD^n$. ([Wak08], Lemma 5.1)*

It is an interesting question as to how many iterations are required to find the vNM-stable set of an $n \times n$ instance I in the algorithm provided by [Wak08] - we will show in Chapter 4 that at most $2n - 3$ iterations are needed. Later, Wako discovered an algorithm that would construct, given an instance I, a compact representation of the vNM-stable set of matchings.

Theorem 2.24. *For any instance I, there exists an algorithm that, in $O(n^2)$ time, outputs an instance I' such that the set of stable matchings over I' is the vNM-stable set of I. ([Wak10], Theorem 6.1+6.2)*

This algorithm does not use the iterative technique in [Wak08] described above. In Chapter 4, we reformulate the Wako algorithm in a slightly modified form that allows us to reveal a few additional conclusions, and provide an alternate polynomial-time algorthm to construct a compact representation of the vNM-stable set of an $n \times n$ instance.

Finally, we note that Theorem 2.24 implies that there exists a mentor-optimal vNM-stable matching, and that there exists an algorithm to compute it in $O(n^2)$ time. In one of the appendices, we will show an algorithm that outputs a specific matching M_0 over the $n \times n$ instance I in $O(n^3)$ time, and prove that M_0 is the mentor-optimal vNM-stable matching over I. (The algorithm was originally found by Mircea Digulescu in [Dig16], but the proof that it creates the mentor-optimal vNM-stable matching is our own.) This construction has the following direct consequence.

Theorem 2.25. *Let M_0 be the mentor-optimal vNM-stable matching over the $n \times n'$ instance I. Then, we may label the vertices of $V_m(I)$ as $\{m_1, \ldots, m_n\}$ and the vertices of $V_d(I)$ as $\{d_1, \ldots, d_{n'}\}$ such that $M_0 = \{(m_1, d_1), (m_2, d_2), \ldots, (m_k, d_k)\}$ for some $k \in \mathbb{N}$, and for all $i \in [k]$ and $j > i$, m_i prefers d_i to d_j.*

The proof of this statement appears in Appendix C. This property has an obvious analogue for the student-optimal vNM-stable matching.

Corollary 2.26. *Let M_1 be the student-optimal vNM-stable matching over the $n \times n'$ instance I. Then, we may label the vertices of $V_m(I)$ as $\{m_1, \ldots, m_n\}$ and the vertices of $V_d(I)$ as $\{d_1, \ldots, d_{n'}\}$ such that $M_1 = \{(m_1, d_1), (m_2, d_2), \ldots, (m_k, d_k)\}$ for some $k \in \mathbb{N}$, and for all $i \in [k]$ and $j > i$, d_i prefers m_i to m_j.*

2.5 Overview of the Pertinent Sections

The overarching focus of Chapters 3 through 7 of this thesis is on a relaxation of stability that we refer to as S-stability. In Chapter 3, we generalize the notion of join and meet on stable matchings, find the conditions on sets of matchings where such notions can be applied, and use them to introduce the notion of S-stability. In Chapter 4, we consider the operation $\psi_I : E(G(I)) \to E(G(I))$ for an instance I, and use it to replicate the results of [Wak08]. We also consider how the operation of $\psi_{I[S]}$ compares to ψ_I for restrictions of the form $I[S]$ - most notable in Theorem 4.29. Lastly, we show that, for any S, $\psi_I^k(S)$ is the unique hub (as defined for Theorem 4.1) for sufficiently large k, and use it to construct an alternate algorithm to represent the vNM-stable matchings (Theorem 4.54).

The next three chapters look at ways to apply the concepts of S-stability to other questions that stem from the structures of the stable matchings. In Chapter 5, we extend the representation of distributive lattices as seen in [GILS87] to the vNM-stable matchings, and discuss the necessary and sufficient conditions that a lattice-sublattice pair must uphold to respectively represent the vNM-stable matchings and stable matchings of some instance (Theorem 5.1). In Chapter 6, we look at the concept of a fractional S-stable matching, consider the necessary and sufficient constraints on the polytope of fractional S-stable matchings for important values of S (Theorem 6.3), and attempt a classification of this polytope for general S. In Chapter 7, we look at representing a graph as the union of a stable matching for some instace, and talk about the discoveries and interesting examples we have found. (Chapters 8 and 9 discuss work on an unrelated problem about boolean functions.)

The appendices pertain to results that originated in previous papers and were rediscovered by us; we present their proofs in our own notation. In Appendix A, we look at a result that [GI89] presents with the skeleton of a proof; in particular, we clarifiy some ambiguous phrasing from the book and present our proof of the result. Appendix B features a proof of Lemma 4.10, which follows the same logic as Wako uses in their proof of Lemma 2.23. Appendix C gives the algorithm for the mentor-optimal vNM-stable matching that originated in [Dig16], and shows how it lets us show Theorem 2.25 and replicate the results in Theorem 2.24.

Chapter 3

An Expanded Notion of Join and Meet

We recall that from Theorem 2.9 ([Knu76]), the set of stable matchings $\mathcal{L}_s$ of an instance I form a distributive lattice under the domination ordering, and that any two stable matchings have a join and a meet that are also stable matchings. We will be interested in relaxations of the stability condition, and in this context it is natural to ask under what conditions do two (not necessarily stable) matchings have a meet and join.

3.1 Join and Meet on Assignments

Recall that, given two stable matchings M_1 and M_2, $M_1 \vee M_2$ is the stable matching such that each student is partnered with their preferred partner among M_1 and M_2, and each mentor is partnered with their preferred partner among M_1 and M_2. It is not obvious that these outputs should be stable, or even matchings; this fact is heavily dependent on M_1 and M_2 being stable. If we wish to extend the notion of $\vee$ and $\wedge$ to operate over a larger domain than just all pairs of stable matchings, we need to extend our domain beyond just matchings.

We may think of a matching (not necessarily stable, or even complete) as a subgraph of $G(I)$ with maximum degree 1. We define an arbitrary subgraph $A \subseteq G(I)$ to be a **mentor-assignment** if every mentor in A has degree at most 1; similarly, we define a subgraph $B \subseteq G(I)$ to be a **student-assignment** if every student in B has degree at most 1. Let $\mathcal{A}$ be the family of all mentor-assignments, $\mathcal{B}$ the family of all student-assignments, and $\mathcal{C}$ the family of all matchings. (Trivially, $\mathcal{C} = \mathcal{A} \cap \mathcal{B}$.) For a mentor-assignment A and mentor m, $p_A(m) = m$ if m has degree 0 in A, and equals the (singular) student adjacent to m in A otherwise; similarly, for a student-assignment B and student d, $p_B(d) = d$ if d has degree 0 in B, and equals the (singular) mentor adjacent to d in B otherwise.

We can order $\mathcal{A}$ via the ordering $\preceq_m$, where $A_1 \preceq_m A_2$ iff for all $m \in V_m(I)$, m prefers $p_{A_1}(m)$ to $p_{A_2}(m)$. This ordering is a product of chains (where each chain corresponds to some $m \in V_m(I)$ and consists of m's ordered preference list of students), and so $\mathcal{A}$ is a distibutive lattice with join and meet defined as follows:

- $A_1 \wedge_m A_2$ consists of all edges of the form (m, d), where m is any mentor and d is their most preferred partner among $p_{A_1}(m)$ and $p_{A_2}(m)$.
- $A_1 \vee_m A_2$ consists of all edges of the form (m, d), where m is any mentor and d is their least preferred partner among $p_{A_1}(m)$ and $p_{A_2}(m)$.

It is trivial to see that, for any two mentor-assignments A_1 and A_2, $A_1 \wedge_m A_2$ and $A_1 \vee_m A_2$ are preserved when the instance I is replaced with any restriction $I[S]$ such that $A_1 \cup A_2 \subseteq S$.

We can similarly order $\mathcal{B}$ via the ordering $\preceq_d$, where $B_1 \preceq_d B_2$ iff for all $d \in V_d(I)$, d prefers $p_{B_1}(d)$ to $p_{B_2}(d)$. This ordering is a product of chains (where each chain corresponds to some $d \in V_m(I)$ having its partner increase in desirability), and so $\mathcal{B}$ is a distibutive lattice with join and meet defined as follows:

- $B_1 \wedge_d B_2$ consists of all edges of the form (m, d), where d is any student and m is their least preferred partner among $p_{B_1}(d)$ and $p_{B_2}(d)$.
- $B_1 \vee_d B_2$ consists of all edges of the form (m, d), where d is any student and m is their most preferred partner among $p_{B_1}(d)$ and $p_{B_2}(d)$.

It is trivial to see that, for any two student-assignments B_1 and B_2, $B_1 \wedge_d B_2$ and $B_1 \vee_m B_2$ are preserved when the instance I is replaced with any restriction $I[S]$ such that $B_1 \cup B_2 \subseteq S$. [1]

Since a subgraph is a matching iff it is both a mentor-assignment and a student-assignment, for any two matchings M_1 and M_2, we can find $M_1 \wedge_m M_2$, $M_1 \vee_m M_2$,

[1] The reader might find it curious that $\wedge_m$ matches vertices with their preferred partners, and $\wedge_d$ does not (and vice versa for $\vee_m$ and $\vee_d$). We use this initially counterintuitive notation because, in the domains we focus on most closely, $\wedge_m$ and $\wedge_d$ will be equal as operations (and similarly for $\vee_m$ and $\vee_d$).

$M_1 \wedge_d M_2$, and $M_1 \vee_m M_2$; furthermore, if M_1 and M_2 are stable matchings, then $M_1 \wedge_m M_2$ and $M_1 \wedge_d M_2$ both equal $M_1 \wedge M_2$ in the lattice of stable matchings, while $M_1 \vee_m M_2$ and $M_1 \vee_d M_2$ both equal $M_1 \vee M_2$ in the lattice of stable matchings. However, if M_1 and/or M_2 are not stable, the resulting assignments are not necessarily matchings. (As an example, consider the instance I such that $V_m(I) = \{m_1, m_2\}$, $V_d(I) = \{d_1, d_2\}$, m_1 and m_2 each have $[d_1, d_2]$ as their respective preference list, and d_1 and d_2 each have $[m_1, m_2]$ as their respective preference list. If $M_1 = \{(m_1, d_1), (m_2, d_2)\}$ and $M_2 = \{(m_1, d_2), (m_2, d_1)\}$, then it is trivial to see that none of $M_1 \vee_m M_2$, $M_1 \wedge_m M_2$, $M_1 \vee_d M_2$, and $M_1 \wedge_d M_2$ are matchings.)

3.2 Costable Matchings

For sets $A \subseteq V_m(I)$ and $B \subseteq V_d(I)$ such that $|A| = |B|$, let $\mathcal{M}(A, B)$ the the set of all perfect matchings between M' and D'. By Theorem 2.4, for any instance I, there exist A, B such that $\mathcal{L}_s(I) \subseteq \mathcal{M}(A, B)$. We know that $\mathcal{L}_s(I)$ is closed under $\vee$ and $\wedge$. The generalizations of $\vee$ and $\wedge$ to mentor-assignments and student-assignments allows us to extend these operations to non-stable matchings. We consider the following questions: given two matchings M_1 and M_2, under what conditions is $M_1 \vee_m M_2$ (resp. $M_1 \vee_d M_2$, $M_1 \wedge_m M_2$, and $M_1 \wedge_d M_2$) a matching? Under what conditions does $M_1 \vee_m M_2 = M_1 \vee_d M_2$ (resp. $M_1 \wedge_m M_2 = M_1 \wedge_d M_2$)?

The answers to these questions lead to the concept of co-stability which we now define. Given a stable matching instance I, we recall that a matching M on the instance is **destabilized** by $e \in E(G(I))$ if m_e prefers d_e to $p_M(m_e)$ and d_e prefers m_e to $p_M(d_e)$; if $S \subseteq E(G(I))$ and M is not destabilized by any $e \in S$, we say that M is **S-stable**. (We generally denote the set of all S-stable matchings as $\mathcal{M}_S$.)

Theorem 3.1. *Let $M, M' \subseteq S$ be two matchings that are also S-stable. Then, all of $M \wedge_m M'$, $M \wedge_d M'$, $M \vee_m M'$, and $M \vee_d M'$ are S-stable matchings. In addition, $M \wedge_m M' = M \wedge_d M'$ and $M \vee_m M' = M \vee_d M'$.*

Proof. Consider the restriction $I[S]$. Both M and M' are stable matchings over $I[S]$, so the following hold over $I[S]$:

- $M \wedge_m M' = M \wedge_d M' = M \wedge M'$.
- $M \vee_m M' = M \vee_d M' = M \vee M'$.

Furthermore, by the properties of stable matchings, $M_0 \equiv M \wedge M'$ and $M_1 \equiv M \vee M'$ are stable matchings over $I[S]$; they are also matchings over I, and retain the property of being S-stable. □

In such a case, we may define $M \wedge M' \equiv M \wedge_m M'$ and $M \vee M' \equiv M \vee_m M'$; it is trivial to see that this agrees with our previous definition of $\wedge$ and $\vee$ in the context of stable matchings. We define two matchings M, M' to be **costable** if M is M'-stable and M' is M-stable. (Note that if M is S-stable, it is also T-stable for any $T \subseteq S$.)

Corollary 3.2. *Let M, M' be two costable matchings. Then, all of $M \wedge_m M'$, $M \wedge_d M'$, $M \vee_m M'$, and $M \vee_d M'$ are matchings. In addition, $M \wedge_m M' = M \wedge_d M'$ and $M \vee_m M' = M \vee_d M'$.*

Proof. Let $S = M \cup M'$; since M and M' are costable (and no matching can be destabilized by an edge in that matching), M and M' are both $\subseteq S$ and S-stable. By Theorem 3.1, we are done. □

In particular, Corollary 3.2 implies that we may naturally extend the operations $\wedge$ and $\vee$ to accept any pair of costable matchings as input. There are some additional observations that we can make on costable matchings.

Proposition 3.3. *Let M and M' be any pair of costable matchings. Then, M and M' cover the same set of vertices.*

Proof. We prove this by contradiction. Assume, for the sake of contradiction, that there exists a vertex v that only one of the matchings covers; WLOG, we may assume v is a mentor and is covered by M but not M'. We may construct a pair of sequences $\{m_0, m_1, m_2, \ldots\}$ and $\{d_1, d_2, \ldots\}$ inductively by setting $m_0 \equiv v$ and, for all positive $i \in \mathbb{N}$, $d_i = p_M(m_{i-1})$ and $m_i = p_{M'}(d_i)$.

Lemma 3.4. *For all positive $i \in \mathbb{N}$, $d_i \in V_d(I)$ prefers $p_{M'}(d_i)$ to $p_M(d_i)$, and $m_i \in V_m(I)$ prefers $p_M(m_i)$ to $p_{M'}(m_i)$.*

Proof. We prove this result by induction on i. For our base case, we note that $m_0 = v$ is a mentor that prefers their partner in M to that in M' - as they has a partner in M but not in M'.

For our inductive step, assume for any positive $i \in \mathbb{N}$ that m_{i-1} is a mentor which is paired under M, and prefers their partner in M to that in M'. By definition, $d_i = p_M(m_{i-1})$; by our inductive assumptions, m_{i-1} is a mentor that is paired under M, so d_i is a student. Given that m_{i-1} prefers their partner in M to that in M', d_i must prefer their partner in M' to that in M - otherwise, M' would be destabilized by $(m_{i-1}, d_i) \in M$, contradicting costability. Since d_i strictly prefers $p_{M'}(d_i)$ to m_{i-1} to d_i, $m_i = p_{M'}(d_i)$ is a mentor; in addition, m_i must prefer their partner in M to that in M' - otherwise, M would be destabilized by $(m_i, d_i) \in M'$, contradicting costability. (This also tells us that m_i is paired under M, since M prefers $p_M(m_i)$ to d_i to m_i.) By induction, we see that for all positive $i \in \mathbb{N}$, d_i is a student that prefers their partner in M' to that in M, and m_i is a mentor that prefers their partner in M to that in M'. (In particular, d_i and m_i have partners in both M and M'.) $\square$

We may define $L \equiv \{m_0, d_1, m_1, d_2, m_2, \ldots\}$, the sequence such that, for all $i \in \mathbb{N}$, $L_{2i} = m_i$ and $L_{2i+1} = d_{i+1}$. Since this sequence is an infinite sequence in a finite domain (namely, $V(I)$), there must be a minimum k such that $L_k = L_j$ for some $j \leq k$. L_k has a partner in M', and $L_0 = v$ doesn't, so the resulting j cannot equal 0. However, if $j \geq 1$, then the fact that $L_k = L_j$ means they have the same type, so j and k are either both even or both odd.

- If k is even, then $L_k = p_{M'}(L_{k-1})$ and $L_j = p_{M'}(L_{j-1})$. Since being paired in M' is a symmetric property, $L_{k-1} = p_{M'}(L_k) = p_{M'}(L_j) = L_{j-1}$. This contradicts the minimality of k such that L_k is not a new term in L, so k cannot be even.
- If k is odd, then $L_k = p_M(L_{k-1})$ and $L_j = p_M(L_{j-1})$. Since being paired in M is a symmetric property, $L_{k-1} = p_M(L_k) = p_M(L_j) = L_{j-1}$. This contradicts the minimality of k such that L_k is not a new term in L, so k cannot be odd.

However, k must be even or odd. This creates a contradiction, so no such v can exist, and M and M' cover the same set of vertices. $\square$

The following corollary is not used in this section, but will be referenced in future ones:

Corollary 3.5. *If $S \subseteq G(I)$ is a set of edges such that $S \supseteq M_0$ for some S-stable matching M_0, then every S-stable matching covers the same set of vertices as M_0.*

Proof. Let M be an arbitrary S-stable matching; then, M is also M_0-stable. In addition M_0 is M-stable (by virtue of being stable), so M and M_0 are costable. By Proposition 3.3, M and M_0 cover the same set of vertices. □

Proposition 3.6. *Let M and M' be any pair of costable matchings over an instance I, $V_m' \subseteq V_m$ be the set of all mentors m that strictly prefer $p_M(m)$ to $p_{M'}(m)$, and $V_d' \subseteq V_d$ be the set of all students d that strictly prefer $p_{M'}(d)$ to $p_M(d)$. Then, $|V_m'| = |V_d'|$, and for all $m \in V_m$, $m \in V_m'$ iff $p_M(m) \in V_d'$ iff $p_{M'}(m) \in V_d'$.*

Proof. We first note, by Proposition 3.3, that every vertex v that is unpaired in M is also unpaired in M', and so $p_M(v) = p_{M'}(v) = v$; as a result, $v \notin V_m'$ or V_d'.

Let $V_m^* \subseteq V_m$ and $V_d^* \subseteq V_d$ respectively represent the mentors and students that are paired under M; by the definition of p_M and Proposition 3.3, p_M and $p_{M'}$ are bijections between V_m^* and V_d^*. For any $m \in V_M^*$, if $m \in V_m'$ and $p_M(m) \notin V_d'$, then m and $p_M(m)$ strictly prefer each other to their respective partners in M', so M' is destabilized by $(m, p_M(m)) \in M$; this contradicts the fact that M' is M-stable, so we have a contradiction and see that if $m \in V_m'$, $p_M(m) \in V_d'$. As p_M is a bijection between $V_m^* \supseteq V_m'$ and $V_d^* \supseteq V_d'$, $|V_m'| \leq |V_d'|$.

Similarly, for any $m \in V_M^*$, if $m \notin V_m'$ and $p_{M'}(m) \in V_d'$, then m and $p_{M'}(m)$ strictly prefer each other to their respective partners in M, so M is destabilized by $(m, p_{M'}(m)) \in M'$; this contradicts the fact that M is M'-stable, so we have a contradiction and see that if $p_{M'}(m) \in V_d'$, $m \in V_m'$. As $p_{M'}$ is a bijection between V_m^* and V_d^*, $|V_m'| \geq |V_d'|$. However, this means that $|V_m'| = |V_d'|$; consequentially every element of V_d' can be expressed as $p_M(m)$ for some $m \in V_m'$, or as $p_{M'}(m)$ for some $m \in V_m'$. □

The natural converse of Proposition 3.6 also holds.

Proposition 3.7. *Let I be a stable matching instance, and M, M' be two matchings such that both p_M and $p_{M'}$ are bijections between V'_m and V'_d (as defined in Proposition 3.6). Then, M and M' are costable.*

Proof. Consider an arbitrary $e \in M$. If m_e prefers d_e to $p_{M'}(m_e)$, then $m_e \in V'_m$; this implies that $d_e = p_M(m_e) \in V'_d$, so d_e prefers $p_{M'}(d_e)$ to m_e. As a result, for every $e \in M$, either m_e or d_e prefers their partner in M' to the other, and so M' is M-stable.

Similarly, consider an arbitrary $e \in M'$. If d_e prefers m_e to $p_M(d_e)$, then $d_e \in V'_d$; this implies that $m_e = p_{M'}(d_e) \in V'_m$, so m_e prefers $p_M(m_e)$ to d_e. As a result, for every $e \in M'$, either m_e or d_e prefers their partner in M to the other, and so M is M'-stable. By the definition of costability, M and M' are costable. □

We can now show that the notion of $\vee$ and $\wedge$ extends to all pairs of costable matchings in the following way.

Theorem 3.8. *Let M, M' be two costable matchings. Then, $M \vee_m M'$ and $M \vee_d M'$ are both matchings. Furthermore, if M and M' are costable, then $M \vee_m M' = M \vee_d M'$ and $M \wedge_m M' = M \wedge_d M'$.*

Proof. Let V'_m, V'_d, V^*_m, and V^*_d be defined as in Proposition 3.6. If M and M' are costable, then $M \vee_m M'$ partners each $m \in V_m(I)$ to $p_{M'}(I)$ if $m \in V'_m$ and $p_M(I)$ otherwise. However, we note by Proposition 3.6 that $p_{M'}$ is a bijection between V'_m and V'_d, and p_M is a bijection between $V^*_m - V'_m$ and $V^*_d - V'_d$, so $M \vee_m M'$ is a matching. By the same reasoning, $M \wedge_m M'$, $M \vee_d M'$, and $M \wedge_d M'$ are all matchings.

Furthermore, for any edge $e \in M - M'$, $e \in M \vee_m M' \Leftrightarrow m_e \in V'_m \Leftrightarrow d_e = p_M(m_e) \in V'_d$ (since M and M' are costable) $\Leftrightarrow e \in M \vee_d M'$; similarly, for any edge $e \in M'$, $e \in M \vee_m M' \Leftrightarrow m_e \notin V'_m \Leftrightarrow d_e = p_M(m_e) \notin V'_d$ (since M and M' are costable) $\Leftrightarrow e \in M \vee_d M'$. Since $M \vee_m M', M \vee_d M' \subseteq M \cup M'$, this is sufficient to show that $M \vee_m M' = M \vee_d M'$. By similar reasoning, we may show $M \wedge_m M' = M \wedge_d M'$. □

3.3 Rotations Over the S-Stable Matchings

A noteworthy example of a set of costable matchings is the set of all S-stable matchings, given that S is a set of edges such that every S-stable matching is $\subseteq S$; we refer to such an S as **stable-closed** over I. We take particular note of Theorem 3.1 in this context.

Proposition 3.9. *Let I be any instance, and S be stable-closed over I. Then, the set of all S-stable matchings over I is the set of all stable matchings over $I[S]$.*

Proof. By the definition of S-stability and $I[S]$, it is trivial to see that any matching M over $I[S]$ is stable over $I[S]$ iff it is S-stable over I. Every matching over $I[S]$ is also a matching over I; in addition, since S is stable-closed, every S-stable matching is also a matching over $I[S]$. Therefore, the proposition holds. □

Theorem 3.10. *Suppose that $S \subseteq E(G(I))$ is stable-closed. Then, the collection of S-stable matchings forms a distributive lattice $\mathcal{L}'_S$, where $M_1 \leq M_2$ iff M_1 dominates M_2, and the operations $\vee$ and $\wedge$ in Theorem 3.1 are the join and meet operations on $\mathcal{L}'_S$ respectively.*

Proof. By Proposition 3.9, the collection of S-stable matchings is the collection of stable matchings over $I[S]$; as a result, by Theorem 2.9, $\mathcal{L}_s(I[S])$ a distributive lattice under the ordering where $M_1 \leq M_2$ iff M_1 dominates M_2, with $\vee$ and $\wedge$ as the join and meet operators respectively. Since the operations of domination, $\vee$, and $\wedge$ are defined only by local properties, it is trivial to see that these properties extend to the poset $\mathcal{L}'_S$ under the same ordering. □

For a given instance, the structure of $\mathcal{L}'_S$ may change for different stable-closed S. However, all of them contain the lattice of stable matchings $\mathcal{L}_{G(I)}$ as a sublattice (since the stable matchings are closed under $\vee$ and $\wedge$). In fact, this sublattice also preserves the covering property, which we will spend the remainder of this section showing.

Proposition 3.11. *Let $S \subseteq E(G(I))$ be any stable-closed set of edges over I. Then, any rotation ρ over I is also a rotation over $I[S]$.*

Proof. Take any such ρ. By Lemma 2.13, there exists a stable matching M_0 over I that exposes ρ. Let $M_1 = (M - \rho_m \cup \rho_d)$; by the definition of an exposed rotation, M_1 is a stable matching. M_0 and M_1 are obviously S-stable as well, and so appear as stable matchings over $I[S]$.

Now, every mentor's preference list in the truncation $I[S]_{(M_0,\emptyset)}$ is a subset of their preference list in $I_{(M_0,\emptyset)}$. In addition, every edge in M_0 or M_1 is still in $I[S]_{(M_0,\emptyset)}$ (since every student weakly prefers their partner in either of M_0 and M_1 to their partner in M_0), so for all mentors $m \in \rho$, $p_{\rho_m}(m)$ and $p_{\rho_d}(m)$ continue to be m's first and second choice respectively in $I[S]_{M_0}$. Consequentially, ρ is a rotation over $I[S]$ exposed by M_0, and so is a rotation over $I[S]$. $\square$

Given a lattice $\mathcal{L}_1$ and a sublattice $\mathcal{L}_0$, we say that $\mathcal{L}_0$ is a **cover-preserving** sublattice of $\mathcal{L}_1$ if for all $l, l' \in \mathcal{L}_0$ such that l' covers l in $\mathcal{L}_0$, l' covers l in $\mathcal{L}_1$. (Note that if $l, l' \in \mathcal{L}_0$ and l' covers l in $\mathcal{L}_1$, l' covers l in $\mathcal{L}_0$ trivially.)

Theorem 3.12. *Let I be any instance and $S \subseteq E(G(I))$ be stable-closed over I. Then, $\mathcal{L}_s$ is a cover-preserving distributive sublattice of $\mathcal{L}'_S$.*

Proof. By virtue of being closed under $\vee$ and $\wedge$, $\mathcal{L}_s$ is a distributive sublattice of $\mathcal{L}'_S$. Now, take any $M, M' \in \mathcal{L}_s$ such that M' covers M in $\mathcal{L}_s$; $(\{e \in E(G(I)) : e \in M, e \notin M'\}, \{e \in E(G(I)) : e \notin M, e \in M'\}$ is therefore a rotation over I, and by Proposition 3.11, is also a rotation over $I[S]$. By Proposition 2.16, M' covers M in $\mathcal{L}'_S$. However, since our choice of M and M' is arbitrary, $\mathcal{L}_s$ must be a cover-preserving sublattice of $\mathcal{L}'_S$. $\square$

Chapter 4

The ψ Operation and the Pull of the Hub

In [Wak08], Jun Wako presented an algorithm that could find the set of vNM-stable matchings for a given $n \times n$ instance I (see Theorem 2.22). However, the algorithm had a prohibitively long runtime in the form that it was presented.

Associated to each stable matching instance I, we define a mapping $\psi_I : 2^{E(G(I))} \rightarrow 2^{E(G(I))}$. As we'll see, the result that every instance has a unique vNM-stable set (Theorem 2.22) is equivalent to the statement that ψ_I has a unique fixed point. We also consider how the operation of $\psi_{I[S]}$ compares to ψ_I for restrictions of the form $I[S]$, most notably in Theorem 4.29. Our foremost conclusions show that if I is an $n \times n$ instance, for $k \geq max(n, 2n-3)$, ψ_I^k maps everything to the unique fixed point of ψ_I (Theorem 4.1 and Theorem 4.52); we use it to construct an alternate algorithm that produces the fixed point of ψ_I (where I is an $n \times n$ instance) in $O(n^3)$ time (Theorem 4.54).

4.1 Preliminaries on the ψ Operation

Associated to every stable matching instance I is a function $\psi_I : 2^E \rightarrow 2^E$, where $E = E(G(I))$. For any $S \subseteq E(G(I))$, we define $\psi_I(S) = \cup_{M \in \mathcal{M}_S} M$, the union of all S-stable matchings. (In cases where the instance I is implied, $\psi_I(S)$ is shortened to $\psi(S)$.) We are especially interested in the fixed points of ψ_I - we define a subset of the edges $S \subseteq E$ to be a **hub** if $\psi(S) = S$.

Theorem 4.1. *1. There exists a set ψ_I^∞ and an integer r so that for all $s \geq r$, $\psi_I^s(\emptyset) = \psi_I^\infty$. (In particular, we note that ψ_I^∞ is a hub.)*

2. Let $\xi(I)$ be the minimum r such that $\psi_I^r(\emptyset) = \psi_I^\infty$. Then, for all $S \subseteq E(G(I))$,

$\psi_I^{\xi(I)}(S) = \psi_I^\infty$.

3. ψ_I^∞ is the unique hub of I.

In cases where I is implied, ψ_I^∞ is shortened to ψ^∞. We define a matching to be **hub-stable** over I if it is ψ_I^∞-stable. In particular, we note that a matching is hub-stable over I iff it is vNM-stable over I, and so item 1 also follows from Theorem 2.22. In this section, we will present a proof of Theorem 4.1 that follows a similar path as the proof for Theorem 2.22 in [Wak08]; however, we discovered the result independently of Wako, and only found their result after. The strategy that we will use to prove item 1 is to consider the sequences $Q = \{\emptyset, \psi^2(\emptyset), \ldots\}$ and $Q' = \{\psi(\emptyset), \psi^3(\emptyset), \ldots\}$, then show that these sequences converge to the same set of edges. By focusing on $\psi : 2^{E(G(I))} \to 2^{E(G(I))}$, as opposed to a function that maps sets of matchings over I to sets of matchings over I, we are then able to use our arguments to show items 2 and 3. (While item 3 also follows from Theorem 2.22, item 2 does not.)

Before we show these results, we note some elementary properties of ψ.

Proposition 4.2. *For any instance I, $\psi(\emptyset) = E$.*

Proof. Since the range of ψ is 2^E, every possible output of ψ is $\subseteq E$, including $\psi(\emptyset)$. For any matching in E, the property of being $\emptyset$-stable is vacuous; therefore, for any edge $e \in E$, the subgraph with edge set $\{e\}$ is a $\emptyset$-stable matching. This shows that $\psi(\emptyset) \supseteq E$, and thus $\psi(\emptyset) = E$. $\square$

Proposition 4.3. *For any instance I, ψ is weakly order-reversing - i.e. if $S_1, S_2 \subseteq E$ and $S_1 \subseteq S_2$, then $\psi(S_1) \supseteq \psi(S_2)$.*

Proof. Suppose $S_1 \subseteq S_2 \subseteq E$. Every matching that is S_2-stable is also S_1-stable (since each such matching is stable with respect to every edge in S_2 - which includes every edge in S_1). Take any edge $e \in \psi(S_2)$; since it appears in a matching which is S_2-stable, it appears in a matching which is S_1-stable (the exact same matching), and so $e \in \psi(S_1)$. The edge e is arbitrary, so $\psi(S_1) \supseteq \psi(S_2)$. $\square$

Corollary 4.4. *For any instance I, ψ^2 is weakly order-preserving - i.e. if $S_1, S_2 \subseteq E$ and $S_1 \subseteq S_2$, then $\psi^2(S_1) \subseteq \psi(S_2)$.*

The above properties are sufficient for us to begin making observations on the sequences Q and Q' described above.

Lemma 4.5. *For any instance I, the sequence $Q = \{\emptyset, \psi^2(\emptyset), \ldots, \psi^{2n}(\emptyset), \ldots\}$ is an increasing sequence that converges to a set of edges $S_\emptyset \subseteq E$ in a finite number of steps (i.e. there exists an $n' \in \mathbb{N}$ such that $Q_n = S_\emptyset$ for all $n \geq n'$).*

Proof. We will prove that Q is increasing by induction on the elements of Q. For our base case, $Q_0 = \emptyset$ is a subset of every element in the range of ψ; since $Q_1 = \psi(\psi(\emptyset))$ is in this range, $Q_1 \subseteq Q_2$. By induction via Corollary 4.4, we see that $Q_i \subseteq Q_{i+1}$ for every positive integer i, and so Q is increasing.

However, every element of Q is in 2^E, a finite set; since Q is also increasing, it must converge to an element of 2^E in a finite number of steps - i.e. there exists an $n' \in \mathbb{N}$ such that $Q_i = S_\emptyset$ for all $i \geq n'$. □

We define n_0 to be the minimum such n' from Lemma 4.5.

Corollary 4.6. *For any instance I, the sequence $Q' = \{E, \psi^2(E), \ldots, \psi^{2n}(E), \ldots\}$ is a decreasing sequence that converges to a set of edges $S_E \subseteq E$ in at most n_0 steps (i.e. $Q_n = S_E$ for all $n \geq n_0$).*

Proof. Since $E = \psi(\emptyset)$, for all $n \in \mathbb{N}$, $Q'_n = \psi^{2n}(E) = \psi^{2n+1}(\emptyset) = \psi(\psi^{2n}(\emptyset)) = \psi(Q_n)$, so $Q' = \psi(Q)$. Since ψ is weakly order-reversing and Q is increasing, Q' is decreasing, and converges to $S_E = \psi(S_\emptyset) \subseteq E$ in at most n_0 steps. □

We will show that $S_\emptyset = S_E$; this implies that the sequence $\{\emptyset, \psi(\emptyset), \ldots, \psi^n(\emptyset), \ldots\}$ converges to a hub. To see why this sequence converges to a hub, we need to identify some important properties about $S_\emptyset$ and S_E.

Proposition 4.7. *$S_E = \psi(S_\emptyset)$ and $S_\emptyset = \psi(S_E)$.*

Proof. As noted in the proof of Corollary 4.6, $Q'_n = \psi(Q_n)$ for all $n \in \mathbb{N}$, so $S_E = \psi(S_\emptyset)$. Now, set $n_0 \in \mathbb{N}$ such that $Q_{n_0} = S_\emptyset$. As Q is an increasing sequence that converges to $S_\emptyset$, every subsequent term of Q equals $S_\emptyset$ - including Q_{n_0+1} - and so $S_\emptyset = Q_{n_0+1} = \psi^2(Q_{n_0}) = \psi(\psi(S_\emptyset))$. However, $\psi(S_\emptyset) = S_E$, so by substitution, $S_\emptyset = \psi(S_E)$. □

Proposition 4.8. $S_\emptyset \subseteq S_E$.

Proof. By the definitions of Q and Q', $Q_0 = \emptyset \subseteq E = Q'_0$. The function ψ^2 is weakly order-preserving, so $(\psi^2)^{n_0} = \psi^{2n_0}$ is weakly order preserving as well, and $\psi^{2n_0}(\emptyset) \subseteq \psi^{2n_0}(E)$. However, $\psi^{2n_0}(\emptyset) = Q_{n_0}$ and $\psi^{2n_0}(E) = Q'_{n_0}$. By the definition of n_0, $Q_{n_0} = S_\emptyset$ and $Q'_{n_0} = S_E$, so by substitution, $S_\emptyset \subseteq S_E$. □

As an aside, all of these propositions allow us to show that the elements of $\{\psi^{2k}(\emptyset) : k \in \mathbb{N}\}$ form a chain, with the order $\emptyset \subseteq \psi^2(\emptyset) \subseteq \psi^4(\emptyset) \subseteq \ldots \subseteq \psi^3(\emptyset) \subseteq \psi(\emptyset)$. [1]

Theorem 4.9. *Let $i, j \in \mathbb{N}$ with $i \leq j$. Then, $\psi^i(\emptyset) \subseteq \psi^j(\emptyset)$ if i is even, and $\psi^j(\emptyset) \subseteq \psi^i(\emptyset)$ if i is odd.*

Proof. Suppose that i is even, so $\frac{i}{2} \in \mathbb{N}$. If j is also even, then $\frac{j}{2} \in \mathbb{N}$, so $\psi^{2(\frac{i}{2})}(\emptyset) \subseteq \psi^{2(\frac{j}{2})}(\emptyset)$ by Lemma 4.5. Otherwise, $\psi^i(\emptyset) \in Q$ and $\psi^j(\emptyset) \in Q'$, so $\psi^i(\emptyset) \subseteq S_\emptyset \subseteq S_E \subseteq \psi^j(\emptyset)$, by Lemma 4.5, Proposition 4.8, and Corollary 4.6 respectively.

Now, suppose that i is odd, so $\frac{i-1}{2} \in \mathbb{N}$. If j is also odd, then $\frac{j-1}{2} \in \mathbb{N}$, so $\psi^{2(\frac{j-1}{2})+1}(\emptyset) \subseteq \psi^{2(\frac{i-1}{2})+1}(\emptyset)$ by Corollary 4.6. Otherwise, $\psi^i(\emptyset) \in Q'$ and $\psi^j(\emptyset) \in Q$, so $\psi^j(\emptyset) \subseteq S_\emptyset \subseteq S_E \subseteq \psi^i(\emptyset)$, by Lemma 4.5, Proposition 4.8, and Corollary 4.6 respectively. □

Given these propositions, we now consider the following lemma:

Lemma 4.10. *Let $J, K \subseteq E$. If $J \subseteq K$, $\psi(J) = K$, and $\psi(K) = J$, then $J = K$. ([Wak08], Lemma 5.1)*

---

[1] We do not immediately use this theorem in this section, but we will use it a number of times in the following sections, and also find it useful for the purposes of visualizing the chain.

We note that this lemma is equivalent to Lemma 2.23. We rediscovered it independently of Wako, and include our own phrasing of the proof in Appendix B. Now, since $S_\emptyset$ and S_E satisfy the hypotheses of the lemma, we obtain:

Corollary 4.11. $S_\emptyset = S_E$.

As a result, we see that item 1 of Theorem 4.1 holds, and $\psi_I^\infty = S_\emptyset$. The following theorem and corollary show that items 2 and 3 hold as well.

Theorem 4.12. *For any stable matching instance I and $S \subseteq E(G(I))$, $\psi_I^{\xi(I)}(S) = \psi_I^\infty$.*

Proof. By the prior corollary, $S_\emptyset$ is a hub, as $\psi(S_\emptyset) = S_E = S_\emptyset$. Now, let S be any set of edges for the instance. Since $S \subseteq E$, $\emptyset \subseteq S \subseteq \psi(\emptyset) = E$; by the order-reversing property of ψ, this implies that $\psi(\emptyset) \supseteq \psi(S) \supseteq \psi^2(\emptyset)$. By repeating this process a total of $2n$ times for any $n \in \mathbb{N}$, we see that $\psi^{2n}(\emptyset) \subseteq \psi^{2n}(S) \subseteq \psi^{2n+1}(\emptyset)$ for all $n \in \mathbb{N}$. For any sufficiently large value of n, $\psi^{2n}(\emptyset) = S_\emptyset$ and $\psi^{2n+1}(\emptyset) = S_E$, so the above relation becomes $S_\emptyset \subseteq \psi^{2n}(S) \subseteq S_E$. However, since $S_\emptyset = S_E$, $S_\emptyset \subseteq \psi^{2n}(S) \subseteq S_\emptyset$, which can only occur if $\psi^{2n}(S) = S_\emptyset$; $S_\emptyset$ is a hub, so this implies that for all $r \geq 2n$, $\psi^r(S) = \psi^{r-2n}(\psi^{2n}(S)) = \psi^{r-2n}(S_\emptyset) = S_\emptyset$. □

Corollary 4.13. *ψ_I^∞ is the unique hub of I.*

In this way, we see that the above $S_\emptyset$ is the unique hub ψ_I^∞. Furthermore, since every hub-stable matching is $\subseteq \psi^\infty$, Proposition 3.11 and Theorem 3.12 have the following trivial corollaries.

Proposition 4.14. *The collection of hub-stable matchings forms a distributive lattice $\mathcal{L}_h$, where $M_1 \leq M_2$ iff M_1 dominates M_2, and the operations $\vee_m$ and $\wedge_m$ in Theorem 3.1 are the join and meet operations on $\mathcal{L}_h$ respectively. Furthermore, the collection of hub-stable matchings is the collection of stable matchings on $I[\psi_I^\infty]$, the instance created by restricting I to ψ_I^∞.*

Theorem 4.15. *Over any given instance I, $\mathcal{L}_s$ is a cover-preserving sublattice of $\mathcal{L}_h$.*

One final conjecture that we may contemplate is that every S such that $\psi(S) \subseteq S$ contains ψ^∞ as a subset. However, this is not the case - if we take any instance I such that $\psi^\infty \neq \psi(E)$, and let e be any edge in $\psi^\infty - \psi(E)$, then $E - \{e\} \supseteq \psi(E - \{e\})$, but $E - \{e\}$ does not contain ψ^∞ as a subset.

4.2 Preliminaries on Satisfactory Instances

We recall that an instance I is satisfactory if there exists a perfect stable matching M_c over I. As noted in [GS85], this is equivalent to every stable matching being perfect. As we have previously noted, matchings (including stable matchings and hub-stable matchings) do not have to be perfect matchings; however, for a given instance I, all of the stable matchings will cover the same vertices. Note that this is not the case for the S-stable matchings in general; as an example, when $S = \emptyset$ and $E(G(I)) \neq \emptyset$, every matching over I is S-stable, so if $e \in E(G(I))$, then $\emptyset$ and $\{e\}$ are S-stable matchings that cover different vertices. However, we will show in Theorem 4.18 that for particularly important values of S, the S-stable matchings do all cover the same vertices.

It is straightforward to see that for any complete $n \times n$ instance, every $e \in \psi_I(S)$ appears in a perfect S-stable matching (any non-perfect matching can be made perfect by arbitrarily matching unpaired vertices, with no vertex becoming less happy). In this section, we will show that satisfactory instances have the same property when $S = \psi_I^k(\emptyset)$ for some $k \in \mathbb{N}$.

We recall that a restriction $I[S]$ of the instance I to $S \subseteq E(G(I))$ is the instance on the same vertex set such that the preference list of every vertex in $I[S]$ is the order-preserving sublist of its vertex list in I where, for any $v_1, v_2 \in V(G(I))$, v_1 appears on v_2's preference list iff $(v_1, v_2) \in S$.

Proposition 4.16. *Suppose, over a given instance I, that there exists a perfect hub-stable matching $M_{c'}$. Then, every hub-stable matching over I is perfect.*

Proof. Consider the restriction $I[\psi_I^\infty]$; by Proposition 4.14, since $M_{c'}$ is stable over I', I' is satisfactory. As a result, every stable matching over I' is perfect. However, every hub-stable matching over I is a stable matching over I', so every hub-stable matching over I is perfect. $\square$

Proposition 4.17. *Let I be an instance. Then, there exists a perfect hub-stable matching over I iff I is satisfactory.*

Proof. Suppose that I is satisfactory, so there exists a perfect stable matching M_c over I. Then, M_c is also hub-stable over I.

Conversely, suppose that there exists a perfect hub-stable matching over I. By Proposition 4.16, every hub-stable matching over I is perfect. Since every stable matching is also hub-stable, I is satisfactory by Corollary 2.5. $\square$

We can identify further properties of satisfactory instances with the following theorem. For any $k \in \mathbb{N}$, we define a matching to be **k-stable** over I if it is $\psi_I^k(\emptyset)$-stable. (In particular, we note that every matching is 0-stable, and the 1-stable matchings over I are the stable matchings over I.)

Theorem 4.18. *For all $k \geq 1$, every k-stable matching covers the same set of vertices.*

Proof. For all $k \geq 1$, $\psi^k(\emptyset) \supseteq \psi^2(\emptyset)$, by Theorem 4.9. $\psi^2(\emptyset)$ is the union of all stable matchings, and so $\supseteq M_0$ for some stable matching. By Corollary 3.5, this implies that every $\psi^k(\emptyset)$-stable matching covers the same set of vertices as M_0. $\square$

Corollary 4.19. *Let I be a satisfactory instance. Then, for all $k \geq 1$, every k-stable matching is a perfect matching.*

Proof. Consider any stable matching M_0. Since I is satisfactory, M_0 must be a perfect matching. In addition, M_0 is $\psi_I^k(\emptyset)$-stable (by virtue of being stable), so every k-stable matching must cover the same set of vertices as M_0 - i.e. every vertex in I. $\square$

4.2.1 Instances with Unique Top Choices

An interesting special case of satisfactory instance is one where every vertex has a distinct top choice. (In such an instance, the mentor-optimal stable matching has every mentor paired with their top choice, and the student-optimal stable matching has every student paired with their top choice.) We will show that, for these instances, the hub is in fact the set of all edge that appear in a stable matching. The following theorem was implicitly used in [Wak10] for the construction of the vNM-stable set of matchings (Lemma 6.2).

Theorem 4.20. *Let I be an instance such that the mentor-optimal and student-optimal stable matchings are the mentor-optimal and student-optimal hub-stable matchings respectively. Then, $\psi_I^2(\emptyset)$ is the unique hub of I. (In other words, the hub of I is the union of all stable matchings over I.)*

Proof. By Theorem 4.15, $\mathcal{L}_s$ is a cover-preserving sublattice of $\mathcal{L}_h$. There exists a maximal chain C in $\mathcal{L}_s$; by the fact that $\mathcal{L}_s$ is a cover-preserving sublattice of $\mathcal{L}_h$ with the same greatest and least elements, C is a maximal chain in $\mathcal{L}_h$ as well. By Corollary 2.15, the edges that appear in a stable matching over I are exactly the edges that appear in an element of C. Similarly, since $I[\psi_I^\infty]$ is the instance created by restricting I to ψ_I^∞, the edges that appear in a stable matching over $I[\psi_I^\infty]$ (i.e. a hub-stable matching over I) are exactly the edges that appear in an element of C; however, by the definition of a hub, these are also the edges in ψ_I^∞, and so $\psi_I^\infty = \cup_{M \in C} M = \cup_{M \in \mathcal{L}_s} M$. □

Corollary 4.21. *Let I be an instance such that the mentor-optimal stable matching has every mentor partnered with their top choice, and the student-optimal stable matching has every student partnered with their top choice. Then, $\psi_I^2(\emptyset)$ is the unique hub of I. (In other words, the hub of I is the union of all stable matchings over I.)*

Proof. In such an instance, the mentor-optimal and student-optimal stable matchings are trivially the mentor-optimal and student-optimal hub-stable matchings as well (since no vertex of the relevant type can find a better partner). By Theorem 4.20, we are done. □

4.3 Making Arbitrary Instances Complete

Over the rest of this chapter, we will look at a number of algorithms that act on stable matching instances; many of these algorithms require the input instance to be satisfactory (which implies that for all $k \geq 1$, all k-stable matchings are perfect). However, these results can ultimately all be extended to nonsatisfactory instances.

In this section, we will discuss how, given a arbitrary instance I, we can construct a complete instance I' that preserves the operation of ψ, in the sense that for any $S \subseteq E(G(I'))$, $\psi_I(S \cap E(G(I))) = \psi_{I'}(S) \cap E(G(I))$.

Conside any instance I with $V_m(I) = \{m_1, m_2, \ldots, m_{n_1}\}$ and $V_d(I) = \{d_1, d_2, \ldots, d_{n_2}\}$ (where n_1 and n_2 are not necessarily equal). Each vertex v has a preference list P_v of vertices of the opposite type; since this instance is not necessarily complete, P_v need not contain every vertex of the opposite type. We define I' from I as follows:

- $V(I')$ has its set of mentors as $\{m_1, m_2, \ldots, m_n\}$ and its set of students as $\{d_1, d_2, \ldots, d_n\}$, where $n = max(n_1, n_2)$.
- For every m_i such that $i \leq n_1$, m_i's preference list P'_{m_i} consists of P_{m_i}, followed by every student not in P_{m_i} in order of increasing index. For every m_i such that $i > n_1$, its preference list is every student listed in order of increasing index.
- For every d_i such that $i \leq n_2$, d_i's preference list P'_{d_i} consists of P_{d_i}, followed by every mentor not in P_{d_i} in order of increasing index. For every d_i such that $i > n_1$, its preference list is every mentor listed in order of increasing index.

We refer to I' created this way as the **completion** of I. The key property of the completion of an instance is as follows.

Proposition 4.22. *Let I be any instance, and I' be the completion of I. Then, for any set of edges $S \subseteq G(I')$, $\psi_I(S \cap G(I)) = \psi_{I'}(S) \cap G(I)$.*

Proof. We show this equality in two parts, that $\psi_I(S \cap G(I)) \supseteq \psi_{I'}(S) \cap G(I)$, and $\psi_I(S \cap G(I)) \subseteq \psi_{I'}(S) \cap G(I)$.

For the former inequality, we may prove this by showing that, for every S-stable matching M in I', $M \cap G(I)$ is $(S \cap G(I))$-stable in I. For every edge $e \in S \cap G(I)$, this edge is also in S, and so M is e-stable in I'. This implies that at least one of m_e and d_e is partnered with someone that they prefer to the other. (Let us refer to such a vertex as v_e, and the other vertex as v'_e.) Since $p_M(v_e)$ appears earlier in P'_{v_e} than v'_e - which, by virtue of e being in $G(I)$, must appear in P_{v_e} - $p_M(v_e)$ must appear in P_{v_e}, and specifically earlier than v'_e. This implies that in I, v_e remains partnered to $p_M(v_e)$ in $M \cap G(I)$, and continues to rank their partner higher than v'_e. Consequentially, $M \cap G(I)$ is e-stable, and since this is true for all $e \in S \cap G(I)$, $M \cap G(I)$ is $(S \cap G(I))$-stable in I.

For the latter, we can show that any $(S \cap G(I))$-stable matching M_0 in I can be extended to an S-stable matching in I', thereby implying that $\psi_I(S \cap G(I)) \subseteq \psi_{I'}(S)$; since $\psi_I(S \cap G(I)) \subseteq G(I)$, this proves the desired inequality. Given M_0, we perfom the following algorithm to produce a perfect matching M:

1. Set $i = 0$.

2. If M_i is a perfect matching, set $M = M_i$ and return. Otherwise, define $M_{i+1} = M_i \cup \{(m_{a(i)}, d_{b(i)})\}$, where $a(i)$ and $b(i)$ are defined such that $m_{a(i)}$ and $d_{b(i)}$ are, respectively, the lowest-index mentor and student that are unmatched in M_i.

3. Set $i = i + 1$ and go to step 2.

Since the iteration in step 2 preserves the property of being a matching and adds a new edge, this algorithm terminates with a perfect matching in at most n cycles. In addition, the indices a and b strictly increase as i increases. We now show that M is S-stable in I'.

Consider any edge $(m, d) \in S$ that is not in M_0.

- If $(m, d) \in G(I)$, then either $p_{M_0}(m) \neq m$ and m prefers their to d, or $p_{M_0}(d) \neq d$ and d prefers them to m. In the former case, m has the same partner in M, and since $(m, p_{M_0}(m)) \in G(I)$, m still prefers $p_{M_0}(m)$ to d in I', so M is (m, d)-stable. In the latter case, d has the same partner in M, and since $(d, p_{M_0}(d)) \in G(I)$, d still prefers $p_{M_0}(d)$ to m in I', so M is (m, d)- stable.

- If $(m, d) \notin G(I)$ and $p_{M_0}(m) \neq m$, then m has the same partner in M as in M_0. Since $(m, p_{M_0}(m)) \in G(I)$ and $(m, d) \notin G(I)$, we know that d was added to m's preference list after any d' such that $(m, d') \in G(I)$; as a result, m prefers $p_{M_0}(m)$ to d in I', and M is (m, d)-stable.

- If $(m,d) \notin G(I)$ and $p_{M_0}(d) \neq d$, then d has the same partner in M as in M_0. Since $(d, p_{M_0}(d)) \in G(I)$ and $(m,d) \notin G(I)$, d prefers $p_{M_0}(d)$ to m in I', so M is (m,d)-stable.

- If $(m,d) \notin G(I)$ and both m and d are unpaired by M_0, then $m = m_{a(i_1)}$ and $d = d_{b(i_2)}$ for some $i_1, i_2 \in \mathbb{N}$. If $i_1 < i_2$, then $p_M(m) = d_{a(i_1)}$ has a smaller index than d; since $(m,d) \notin G(I)$, this means that m prefers $p_M(m)$ to d, and M is (m,d)-stable. Otherwise (since $(m,d) \notin M$), $i_1 > i_2$, so $p_M(d) = m_{a(i_2)}$ has a smaller index than m; since $(m,d) \notin G(I)$, this means that d prefers $p_M(d)$ to m, and M is (m,d)-stable.

As a result, M is (m,d)-stable for every $(m,d) \in S - M$; M is also trivially (m,d)-stable for every $(m,d) \in M$. Therefore, M is S-stable. $\square$

We note two important consequences of Proposition 4.22.

Corollary 4.23. *Let I be any instance, and I' be any completion of I. Then, for all $k \in \mathbb{N}$, $\psi_I^k(\emptyset) = \psi_{I'}^k(\emptyset) \cap G(I)$.*

Proof. We prove this result by induction on k. For our base case, when $k = 0$, $\psi_I^0(\emptyset) = \emptyset = \emptyset \cap G(I) = \psi_{I'}(\emptyset) \cap G(I)$.

Now, for our inductive step, assume, for some arbitrary $k \in \mathbb{N}$, that $\psi_I^k(\emptyset) = \psi_{I'}^k(\emptyset) \cap G(I)$; we look to show that $\psi_I^{k+1}(\emptyset) = \psi_{I'}^{k+1}(\emptyset) \cap G(I)$. Let $S \equiv \psi_{I'}^k(\emptyset)$; by our inductive assumption, $\psi_I^k(\emptyset) = S \cap G(I)$. As a result, $\psi_I^{k+1}(\emptyset) = \psi_I(\psi_I^k(\emptyset)) = \psi_I(S \cap G(I)) = \psi_{I'}(S) \cap G(I)$ by Proposition 4.22. However, $\psi_{I'}(S) = \psi_{I'}(\psi_{I'}^k(\emptyset)) = \psi_{I'}^{k+1}(\emptyset)$, so $\psi_I^{k+1}(\emptyset) = \psi_{I'}^{k+1}(\emptyset) \cap G(I)$.

By induction, $\psi_I^k(\emptyset) = \psi_{I'}^k(\emptyset) \cap G(I)$ for all $k \in \mathbb{N}$. $\square$

Corollary 4.24. *Let I be any instance, and I' be any completion of I. Then, $\psi_I^\infty = \psi_{I'}^\infty \cap G(I)$.*

Proof. Let $S \equiv \psi_{I'}^\infty$. By Proposition 4.22, $\psi_I(S \cap G(I)) = \psi_{I'}(S) \cap G(I)$. Since S is a hub of I', $\psi_{I'}(S) = S$, so $\psi_I(S \cap G(I)) = S \cap G(I)$, and $S\ capG(I)$ is a hub of I. By Theorem 4.12, this means that $\psi_I^\infty = S \cap G(I) = \psi_{I'}^\infty \cap G(I)$. $\square$

4.4 The Behavior of ψ on Restrictions

In this section, we investigate the relationship of the operators ψ_I and $\psi_{I'}$, where I' is a restriction on I. As we recall, a restriction $I[S]$ on I is an instance on the same set of mentors and students such that $G(I[S]) = S \subseteq G(I)$ and, for all $v, v_1, v_2 \in V(G(I))$ such that $(v, v_1), (v, v_2) \in G(I')$, v's preference ordering of v_1 and v_2 is the same in I and $I[S]$.

Proposition 4.25. *Let I be any instance, and I' be any restriction of I. Then, for all $S \subseteq G(I')$, $\psi_{I'}(S) \subseteq \psi_I(S)$.*

Proof. By its definition, $\psi_{I'}(S)$ is the union of every S-stable matching M over I'. For every such M, M is also a matching over I, and is S-stable there as well; consequentially, the set of S-stable matchings over I contains every such M, and so $\psi_I(S) \supseteq \psi_{I'}(S)$. $\square$

Note that $\psi_{I'}(S)$ is not necessarily equal to $\psi_I(S)$ under such conditions - for example, if I is any instance, I' is any restriction of I with $G(I') \neq G(I)$, and $S = \emptyset$. In this way, we see how properties of $\psi_{I'}$ are modified in ψ_I. Furthermore, we note that the domain of $\psi_{I'}$ is a subset of the domain of ψ_I - namely, if S contains any edge in $G(I) - G(I')$, then $\psi_I(S)$ is defined, but $\psi_{I'}(S)$ is not. For an instance I' and a set of edges S, we may consider $\psi_{I'}(S) \equiv \psi_{I'}(S \cap G(I'))$.

Proposition 4.26. *Let I be any instance, and I' be any restriction of I. Then, for all $S \subseteq G(I)$ such that $\psi_I(S) \subseteq G(I')$, $\psi_{I'}(S) \supseteq \psi_I(S)$.*

Proof. By its definition, $\psi_I(S)$ is the union of every S-stable matching M over I. For every such M, M is also a matching over I' (since $M \subseteq \psi_I(S) \subseteq G(I')$), and is S-stable there as well; consequentially, the set of S-stable matchings over I' contains every such M, and so $\psi_I(S) \subseteq \psi_{I'}(S)$. $\square$

The two above propositions give us the following result:

Theorem 4.27. *Let I be any instance, and I' be any restriction of I. Then, for all $S \subseteq G(I')$ such that $\psi_I(S) \subseteq G(I')$, $\psi_{I'}(S) = \psi_I(S)$.*

Corollary 4.28. *Let I be any instance, and I' be any restriction of I such that $\psi_I^\infty \subseteq G(I')$. Then, $\psi_{I'}^\infty = \psi_I^\infty$.*

Proof. Let $S = \psi_I^\infty$. Since $S \subseteq G(I')$, and $\psi_I(S) = S \subseteq G(I')$, $\psi_{I'}(S) = \psi_I(S) = S$ by Theorem 4.27. However, by the definition of the hub, this implies that S is a hub of I' - and by Theorem 4.1, is the unique hub of I'. $\square$

We may hope that this preservation of the operation of ψ on restrictions holds for general S; however, as seen below, it is possible to find I, I', and S such that $\psi_I(S)$ and $\psi_{I'}(S)$ differ dramatically. The reason why, from an intuitive perspective, is because the most "appealing" edges in $G(I)$, which have a very large impact on what matchings aren't stable, are not present in $G(I')$. However, truncations, which we recall are restrictions created by iteratively removing a vertex from the bottom of another vertex's preference list, specifically avoid removing these "appealing" edges, and so we can conclude much stronger results on how $\psi_I(S)$ influences $\psi_{I'}(S)$ when I' is a truncation of I.

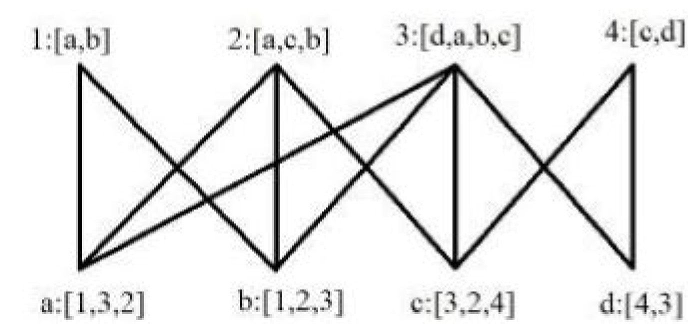

$\Psi_I^2(\phi)$

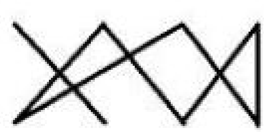

$\Psi_{I[H]}^2(\phi)$

For the remainder of this section, we will focus on truncations of the form $I_{(M_1,M_2)}$, where M_1 and M_2 are matchings. We refer to such truncations as **subinstances**. In the following theorem, we note that, for any $S \subseteq E(G(I))$, when the mentors truncate their preference lists to a matching that is $\subseteq S$ and S-stable, the behavior of S under the ψ operation is preserved on the new subinstance.

Theorem 4.29. *Let M be any matching over an instance I, and $I' = I_{(\emptyset,M)}$. Then, for any $S \subseteq G(I)$ such that $M \subseteq S$ is S-stable and $\psi_I(S)$-stable, $\psi_{I'}(S \cap G(I')) = \psi_I(S) \cap G(I')$.*

Proof. We prove this result by showing, for any edge $e \in G(I')$, if e appears in one of $\psi_I(S) \cap G(I')$ and $\psi_{I'}(S \cap G(I'))$, then it appears in the other.

Suppose that $e \in \psi_{I'}(S \cap G(I'))$. To show that $e \in \psi_I(S)$, let M' be an $S \cap G(I')$-stable matching over I' that contains e; we will claim that M' is also $\psi_I(S)$-stable. To do this, we first note that M' is also a $(S \cap G(I'))$-stable matching over I (as the relative preference orderings through edges in M' and $S \cap G(I')$ is preserved). Furthermore, if $m \in V_m(I)$, then m weakly prefers $p_{M'}(m)$ to $p_M(m)$, and prefers $p_M(m)$ to any student d such that $(m, d) \notin G(I')$ (by the definition of I'). Consequentially, M' is also $(G(I) - G(I'))$-stable; as a result, M' is $\{e_S\}$-stable for every $e_S \in S$, and so is S-stable over I. Since $e \in M'$, $e \in \psi_I(S)$, so $e \in \psi_I(S) \cap G(I')$.

Now, suppose that $e \in \psi_I(S) \cap G(I')$; then, there exists an S-stable matching M^* over I that contains e. Since $M \subseteq S$, M^* is also M-stable; in addition, since $M^* \subseteq \psi_I(S)$, M is M^*-stable. As a result, M and M^* are costable, and so their meet $M' \equiv M \vee M^*$ is a matching by Theorem 3.8. We note the following properties of M':

- M' consists only of edges (m, d) where m weakly prefers d to $p_M(m)$, implying $M' \subseteq G(I')$. As a result, we see that M' also exists as a matching over I'.

- If $e = (m_e, d_e)$, then m_e weakly prefers $p_{M^*}(m_e) = d_e$ to $p_M(m_e)$ (as $e \in G(I')$), so $e \in M'$.

- Since M and M^* are both S-stable, M' is as well, and so is $(S \cap G(I'))$-stable. This property is preserved over I'.

As a result, M' is an $(S \cap G(I'))$-stable matching over I' that contains e, thereby proving that $e \in \psi_{I'}(S \cap G(I'))$. $\square$

We note that this preservation of the behavior of S under ψ also occurs when the students truncated their preference lists to such a matching, or when the mentors truncate their preference lists to one matching and the students to another.

Corollary 4.30. *Let M be any matching over an instance I, and $I' = I_{(M,\emptyset)}$ be the instance created by restricting I to edges (m, d) such that d weakly prefers m to $p_M(d)$. Then, for any $S \subseteq G(I)$ such that $M \subseteq S$ is S-stable and $\psi_I(S)$-stable, $\psi_{I'}(S \cap G(I')) = \psi_I(S) \cap G(I')$.*

Corollary 4.31. *Let $M_1 \preceq M_2$ be any two costable matchings over an instance I, and $I' = I_{(M_1,M_2)}$. Then, for any $S \subseteq G(I)$ such that M_1 and M_2 are both subsets of S, S-stable and $\psi_I(S)$-stable, $\psi_{I'}(S \cap G(I')) = \psi_I(S) \cap G(I')$.*

Proof. By Theorem 4.29, the instance $I'' = I_{(\emptyset,M_2)}$ preserves the operation of ψ on S - i.e. $\psi_{I''}(S \cap G(I'')) = \psi_I(S) \cap G(I'')$. In I'', M_1 remains both a subset of $S \cap G(I'')$ and $S \cap G(I'')$-stable over I'', and since $\psi_{I''}(S \cap G(I'') = \psi_I(S) \cap G(I'') \subseteq \psi_I(S)$, it is $\psi_{I''}(S \cap G(I''))$-stable as well. Since $I' = I''_{(M_1,\emptyset)}$, by Corollary 4.30, $\psi_{I'}(S \cap G(I')) = \psi_I(S) \cap G(I')$. $\square$

There are two particular consequences of note for this theorem - each consequence is respectively presented here in the form of a theorem and two corollaries.

Theorem 4.32. *Let M be any stable matching over an instance I, and $I' = I_{(\emptyset,M)}$. Then, for any $k \in \mathbb{N}$, $\psi^k_{I'}(\emptyset) = \psi^k_I(\emptyset) \cap G(I')$.*

Proof. We prove this result by induction on k. For our base case, when $k = 1$, $\psi_{I'}(\emptyset) = G(I')$, while $\psi_I(\emptyset) \cap G(I') = G(I) \cap G(I') = G(I')$.

For our inductive step, assume the statement is true for $k = k_0$ for some $k_0 \in \mathbb{N}$; we will prove that it is true for $k = k_0 + 1$. Let $S \equiv \psi^{k_0}_I(\emptyset)$; by our inductive assumption, $\psi^{k_0}_{I'}(\emptyset) = S \cap G(I')$. As a result, $\psi^{k_0+1}_{I'}(\emptyset) = \psi_{I'}(\psi^{k_0}_{I'}(\emptyset)) = \psi_{I'}(S \cap G(I'))$; similarly, $\psi^{k_0+1}_I(\emptyset) \cap G(I') = \psi_I(\psi^{k_0}_I(\emptyset)) \cap G(I') = \psi_I(S) \cap G(I')$. Furthermore, since $M \subseteq \psi^2_I(\emptyset) \subseteq$

$\psi_I^{k_0}(\emptyset)$ (given that $k_0 \geq 1$), $M \subseteq S$; it is also trivial to see that, since M is stable, it is also S-stable and $\psi(S)$-stable. By Theorem 4.29, $\psi_{I'}(S \cap G(I')) = \psi_I(S) \cap G(I')$, so by substitution, $\psi_{I'}^{k_0+1}(\emptyset) = \psi_I^{k_0+1}(\emptyset) \cap G(I')$.

By induction, we see that $\psi_{I'}^k(\emptyset) = \psi_I^k(\emptyset) \cap G(I')$ for all $k \in \mathbb{N}$. $\square$

Corollary 4.33. *Let M be any stable matching over an instance I, and $I' = I_{(M,\emptyset)}$. Then, for any $k \in \mathbb{N}$, $\psi_{I'}^k(\emptyset) = \psi_I^k(\emptyset) \cap G(I')$.*

Corollary 4.34. *Let M_1, M_2 be two stable matchings over an instance I such that M_1 dominates M_2, and $I' = I_{(M_1,M_2)}$. Then, for any $k \in \mathbb{N}$, $\psi_{I'}^k(\emptyset) = \psi_I^k(\emptyset) \cap G(I')$.*

Proof. Let $I'' = I_{(\emptyset,M_2)}$, so $\psi_{I''}^k(\emptyset) = \psi_I^k(\emptyset) \cap G(I'')$ by Theorem 4.32. We note that $I' = I''_{(M_1,\emptyset)}$, so by Corollary 4.33, $\psi_{I'}^k(\emptyset) = \psi_{I''}^k(\emptyset) \cap G(I') = \psi_I^k(\emptyset) \cap G(I'') \cap G(I') = \psi_I^k(\emptyset) \cap G(I')$. $\square$

Theorem 4.35. *Let M be any hub-stable matching over an instance I, and $I' = I_{(\emptyset,M)}$. Then, $\psi_{I'}^\infty = \psi_I^\infty \cap G(I')$.*

Proof. Let $S = \psi_I^\infty$ Since M is hub-stable, it is both a subset of S and S-stable; in addition, $\psi(S) = S$, so M is also $\psi(S)$-stable. By Theorem 4.29, $\psi_{I'}(S \cap G(I')) = \psi_I(S) \cap G(I') = S \cap G(I')$. Therefore, $S \cap G(I')$ is the unique hub over I', and so $\psi_{I'}^\infty = S \cap G(I') = \psi_I^\infty \cap G(I')$. $\square$

Corollary 4.36. *Let M be any hub-stable matching over an instance I, and $I' = I_{(M,\emptyset)}$. Then, $\psi_{I'}^\infty(\emptyset) = \psi_I^\infty(\emptyset) \cap G(I')$.*

Corollary 4.37. *Let M_1, M_2 be two stable matchings over an instance I such that M_1 dominates M_2, and $I' = I_{(M_1,M_2)}$. Then, $\psi_{I'}^\infty(\emptyset) = \psi_I^\infty(\emptyset) \cap G(I')$.*

Proof. Let $I'' = I_{(\emptyset,M_2)}$, so $\psi_{I''}^\infty(\emptyset) = \psi_I^\infty(\emptyset) \cap G(I'')$ by Theorem 4.35. We note that $I' = I''_{(M_1,\emptyset)}$, so by Corollary 4.36, $\psi_{I'}^\infty(\emptyset) = \psi_{I''}^\infty(\emptyset) \cap G(I') = \psi_I^\infty(\emptyset) \cap G(I'') \cap G(I') = \psi_I^\infty(\emptyset) \cap G(I')$. $\square$

4.5 Computing Important $\psi(S)$

We consider the computational problem: Given an $n \times n$ instance I and $S \subseteq E(G(I))$, find $\psi(S)$. The definition of $\psi(S)$ as the union of all S-stable matchings gives a natural algorithm: generate all S-stable matchings and find their union. This naive algorithm has a worst-case running time that is exponential in n, since the number of S-stable matchings can be exponential in n. We do not know a polynomial time algorithm for computing $\psi(S)$ for general S. In this section, we provide polynomial time algorithms that compute $\psi(S)$ when S meets certain natural conditions.

For the rest of this paper, whenever we say that an algorithm pertaining to an $n_1 \times n_2$ instance runs in polynomial time, we mean that it runs in time that is polynomial in terms of $n \equiv max(n_1, n_2)$.

One such case is outlined by Proposition 4.2 - namely, if $S = \emptyset$, then $\psi(S) = G(I)$. Another specific value of S for which $\psi(S)$ is easily computable is $S = G(I)$ - in this case, the lattice of stable matching can be constructed in $O(n^2)$ time, as noted in Corollary 2.20, and the edges that appear in $\psi(S)$ are precisely those that appear in some rotation over I. This strategy can be extended to generate $\psi(S)$ whenever S is stable-closed. (Recall that S is stable-closed when every S-stable matching over I is $\subseteq S$ - it is trivial to see that this is equivalent to saying that $\psi_I(S) \subseteq S$.

Theorem 4.38. *If $S \subseteq E(G(I))$ is stable-closed over I, then we may construct $\psi(S)$ in $O(n^2)$ time.*

Proof. By Theorem 3.10, the set of S-stable matchings over I is precisely the set of stable matchings over $I[S]$. Over this restricted $n_1 \times n_2$ instance, we may apply Corollary 2.20 to find $\psi(S)$ in $O(n^2)$ time. □

Since we have an algorithm for computing $\psi(S)$ when $S \supseteq \psi(S)$, we may consider whether a similar algorithm exists when $S \subseteq \psi(S)$. While we don't know such an algorithm, we do have an algorithm that works if S satisfies a somewhat more restrictive condition.

Theorem 4.39. *Let $S \subseteq E(G(I))$ be a stable-closed set such that $\psi^2(S) \subseteq S$. Then, we may construct $\psi^2(S)$ in $O(n^2)$ time.*

We prove this via the following:

Lemma 4.40. *Let I be a satisfactory instance. Then, given $\psi_I^2(\emptyset)$, we may construct $\psi_I^3(\emptyset)$ in $O(n^2)$ time.*

We will hold off on proving this lemma; however, we may immediately note this consequence.

Corollary 4.41. *Let I be any instance. Then, given $\psi_I^2(\emptyset)$, we may construct $\psi_I^3(\emptyset)$ in $O(n^2)$ time.*

Proof. Let I^* be the completion of I. By Theorem 2.6, I^* is satisfactory, and so we can construct $\psi_{I^*}^3(\emptyset)$ in $O(n^2)$ time. Furthermore, by Corollary 4.23, $\psi_I^3(\emptyset) = \psi_{I^*}^3(\emptyset) \cap E(G(I))$, and so we can easily construct $\psi_I^3(\emptyset)$ in $O(n^2)$ time. □

We now prove Theorem 4.39.

Proof. Let $I' \equiv I[S]$. Since $\psi(S), \psi^2(S) \subseteq G(I')$, we see by Theorem 4.27 that $\psi_{I'}(S) = \psi_I(S)$; $S = G(I') = \psi_{I'}(\emptyset)$, so $\psi_{I'}(S) = \psi_{I'}^2(\emptyset)$. Using Corollary 4.41, we can construct $\psi_{I'}^2(S) = \psi_{I'}^3(\emptyset)$ in $O(n^5)$ time. However, since $S, \psi_I(S) \subseteq G(I')$ by the inital conditions on S, $\psi_I^2(S) = \psi_{I'}^2(S)$, so we have constructed $\psi_I^2(S)$. □

Together, Theorem 4.38 and Theorem 4.39 give us a mechanism to construct the sequence:

$$\{\emptyset, \psi(\emptyset), \psi^2(\emptyset), \ldots, \psi^k(\emptyset)\}$$

in $O(kn^2)$ time for any instance I. The first two elements are constructed trivially - $\emptyset$ is explicitly given, whereas $\psi(\emptyset) = G(I)$. The subsequent elements can be determined by an inductive argument.

Theorem 4.42. *For any non-negative $i \in \mathbb{N}$, given $\psi_I^i(\emptyset)$ (and $\psi_I^{i-1}(\emptyset)$, if $i > 0$), we may construct $\psi_I^{i+1}(\emptyset)$ in $O(n^2)$ time.*

Proof. Let $S \equiv \psi_I^i(\emptyset)$. If i is odd, then by Theorem 4.9, $\psi_I^{i+1}(\emptyset) \subseteq \psi_I^i(\emptyset)$; we can therefore use Theorem 4.38 to construct $\psi_I^{i+1}(\emptyset)$ in $O(n^2)$ time. On the other hand, if i is odd, then $\psi_I^i(\emptyset) \subseteq \psi_I^{i+1}(\emptyset) \subseteq \psi_I^{i-1}(\emptyset)$ by Theorem 4.9; by applying Theorem 4.39 with $T = \psi_I^{i-1}(\emptyset)$, we may construct $\psi_I^{i+1}(\emptyset)$ in $O(n^2)$ time. □

By induction, we see that the entire sequence is generated in $O(kn^2)$ time.

4.5.1 Proof of Lemma 4.40

In this section, we provide a proof for Lemma 4.40. Recall that, for $k \in \mathbb{N}$, a matching is k-stable if it is $\psi_I^k(\emptyset)$-stable.

Since $\psi^2(\emptyset) \subseteq \psi^3(\emptyset)$ by Theorem 4.9, in order to find $\psi^3(\emptyset)$, we only need to determine, for every $e \in G(I) - \psi^2(\emptyset)$, if $e \in \psi^3(\emptyset)$. To that end, consider the mentor-optimal and student -optimal stable matchings, M_0 and M_1 respectively. Let M be any 2-stable matching. Since M_0 and M_1 are stable, they are also M-stable; similarly, $M_0, M_1 \subseteq \psi^2(\emptyset)$, so M is M_0-stable and M_1-stable. As a result, M is costable with M_0 and M_1, so by Theorem 3.1, any combination of joins and meets of these elements will result in a $\psi^2(\emptyset)$-stable matching (since M_0 and M_1 are trivially $\psi^2(\emptyset)$-stable).

Now, consider any edge $e \in G(I) - \psi^2(\emptyset)$. By Corollary 3.5, every 2-stable matching covers the same vertices as any stable matching; therefore, any edge that covers a vertex that M_0 does not cover cannot be in a 2-stable matching, and so isn't in $\psi^3(\emptyset)$. In addition, if, for any $i \in [0, 1]$, m_e prefers d_e to $p_{M_i}(m_e)$ and d_e prefers m_e to $p_{M_i}(d_e)$, then any matching that contains e cannot be costable with M_i by Proposition 3.6; as a result, any such matching cannot be 2-stable, and so $e \notin \psi^3(\emptyset)$. Similarly, if, for any $i \in [0, 1]$, m_e prefers $p_{M_i}(m_e)$ to d_e and d_e prefers $p_{M_i}(d_e)$ to m_e, then any matching that contains e cannot be costable with M_i by Proposition 3.6; as a result, any such matching cannot be 2-stable, and so $e \notin \psi^3(\emptyset)$. As a result, every edge $e \in \psi^3(\emptyset)$ must fit in one of the following categories:

1. m_e prefers $p_{M_0}(m_e)$ to d_e to $p_{M_1}(m_e)$, and d_e prefers $p_{M_1}(d_e)$ to m_e to $p_{M_0}(d_e)$.

2. m_e prefers d_e to $p_{M_0}(m_e)$, and d_e prefers $p_{M_0}(d_e)$ to m_e.

3. m_e prefers $p_{M_1}(m_e)$ to d_e, and d_e prefers m_e to $p_{M_1}(d_e)$.

Let E be the set of all edges that fulfill the second set of conditions, and E^* be the set of all edges that fulfill the third set of conditions. For each type of edge, we look at the set of all edges in $G(I)$ of that type, and consider which appear in $\psi_I^3(\emptyset)$.

Lemma 4.43. *Let $e \in \psi_I^3(\emptyset)$ such that m_e prefers $p_{M_0}(m_e)$ to d_e to $p_{M_1}(m_e)$, and d_e prefers $p_{M_1}(d_e)$ to m_e to $p_{M_0}(d_e)$. Then, $e \in \psi_I^2(\emptyset)$.*

Proof. Every such e appears in the subinstance $I_3 \equiv I_{(M_0,M_1)}$. In this subinstance, we observe that M_0 is a stable matching where each mentor is paired with their top partner, and M_1 is a stable matching where each student is paired with their top partner; by Corollary 4.21, $\psi_{I_3}^2(\emptyset)$ is the hub of I_3, and so $\psi_{I_3}^3(\emptyset) = \psi_{I_3}^2(\emptyset)$. Since M_0 and M_1 are stable over I, this implies that $\psi_I^3(\emptyset) \cap G(I_3) = \psi_I^2(\emptyset) \cap G(I_3)$ by Corollary 4.34. Consequentially, every such $e \in \psi_I^3(\emptyset)$ also appear in $\psi_I^2(\emptyset)$. $\square$

Lemma 4.44. *$\psi_I^3(\emptyset) \cap E$ is the union of all perfect matchings over E.*

Proof. We note that $E = E(G(I_{(\emptyset,M_0)}))$; set $I' \equiv I_{(\emptyset,M_0)}$. By Theorem 4.32, $\psi_{I'}^3(\emptyset) = E \cap \psi_I^3(\emptyset)$, so any edge $e \in E$ is in $\psi_I^3(\emptyset)$ iff it is in $\psi_{I'}^3(\emptyset)$.

Since $\psi_{I'}^2(\emptyset) = E \cap \psi_I^2(\emptyset) = M_0$, any 2-stable matching over I' must be perfect by Corollary 3.5. Conversely, for any edge $e \in \psi_{I'}^2(\emptyset) = M_0$, m_e prefers their partner in such a perfect matching to d_e, their partner in M_0 (by the definition of E); consequentially, every perfect matching over E is 2-stable over I'. Thus, $\psi_{I'}^3(\emptyset) = E \cap \psi_I^3(\emptyset)$ is the union of all perfect matchings in E. We also know that E contains the perfect matching M_0 over the vertices of that are matched in any 2-stable matching over I. $\square$

Corollary 4.45. *$\psi_I^3(\emptyset) \cap E^*$ is the union of all perfect matchings over E^*.*

Applying the above three results to the classification of the three types of edges in $\psi_I^3(\emptyset)$ shows us the following.

Theorem 4.46. *$\psi_I^3(\emptyset) = \psi_I^2(\emptyset) \cup P \cup P^*$, where P and P^* are the unions of all perfect matchings over E and E^* respectively.*

As an aside, we note the following corollary (which we do not use to prove Lemma 4.40, but will use in a later section of the thesis):

Corollary 4.47. *Let $e \in \psi_I^3(\emptyset) - \psi_I^2(\emptyset)$. Then, either m_e prefers d_e to their partner in the mentor-optimal stable matching over I, or d_e prefers m_e to their partner in the student-optimal stable matching over I.*

Proof. Since $e \notin \psi_I^2(\emptyset)$, $e \in P \cup P^*$. If $e \in P$, then $e \in E$, so m_e prefers d_e to their partner in the mentor-optimal stable matching over I. Similarly, if $e \in P^*$, then $e \in E^*$, so d_e prefers m_e to their partner in the student-optimal stable matching over I. □

Consequentially, in order to construct $\psi_I^3(\emptyset)$, we need only to find $\psi_I^2(\emptyset)$, P, and P^*. $\psi_I^2(\emptyset)$ can be constructed in $O(n^2)$ time, so we are only left with the task of constructing P and P^*. However, each of P and P^* is the union of all perfect matchings over a specific subgraph of $G(I)$; this allows us to apply the following result, discovered by Tamir Tassa.

Theorem 4.48. *Let G be any bipartite graph with n vertices and k edges, such that there exists a perfect matching over G. Then, there exists an algorithm that inputs G, and outputs the union of all perfect matchings over G in $O(n + k)$ time. ([Tas12], Algorithm 2)*

We may now prove Lemma 4.40 by showing that each of $\psi_I^2(\emptyset)$, P, and P^* can be constructed in $O(n^2)$ time.

Proof. By Corollary 2.20, we can construct $\psi_I^2(\emptyset)$ in $O(n^2)$ time. We note that since I is satisfactory, the mentor-optimal stable matching M_0 is a perfect matching over E. We also note that, since I is an $n \times n$ instance, $|V(E)| \leq 2n$ and $|E| \leq n^2$. Consequentially, we see that we can find P in $O(2 * n + n^2) = O(n^2)$ time. Similarly, we can find P^* in $O(n^2)$ time (the student-optimal stable matching M_1 is also found in the process of finding $\psi_I^2(\emptyset)$, and is a perfect matching over E^*). As a result, by Theorem 4.46, we can find $\psi_I^3(\emptyset)$ in $O(n^2) + O(n^2) + O(n^2) = O(n^2)$ time. □

4.6 Analysis of the Convergence Rate of ψ

We recall that the evolution of the sequence $\{\emptyset, \psi(\emptyset), \psi^2(\emptyset), \ldots\}$ corresponds to the algorithm for finding the vNM-stable matchings for a given instance described in [Wak08]. However, it was previously unknown how many iterations are needed for the sequence to converge. For a given $n \times n$ instance I, we recall that $\xi(I)$ is the minimum $r \in \mathbb{N}$ such that $\psi_I^s(\emptyset) = \psi_I^\infty$ for all $s \geq r$. (As a consequence of Theorem 4.12, $\psi_I^r(S) = \psi_I^\infty$ for all $S \subseteq G(I)$ and $r \geq \xi(I)$.) For all $n \in \mathbb{N}$, we may also define $\Xi(n)$ to be the maximum value of $\xi(I)$ over all $n \times n$ instances I; the similar $\Xi^*(n)$ is the maximum value of $\xi(I)$ over all satisfactory $n \times n$ instances I. In this section, we determine the values of $\Xi^*(n)$ and $\Xi(n)$ for all $n \in \mathbb{N}$ (see Theorem 4.51 and Theorem 4.52 respectively).

When $n = 1$ or 2, the number of possible instances is very small, and so it can easily be confirmed by hand that $\Xi(n) = \Xi^*(n) = n$ for such values of n. However, for larger values of n, the number of instances becomes far larger than can be listed out by hand. Our previous arguments allow us to make some observations on $\xi(I)$ for a general instance I.

Proposition 4.49. *For an instance I such that $|E(G(I))| = k$ and every stable matching has q edges, $\xi(I) \leq k - q + 1$.*

Proof. By Theorem 4.9, $\psi_I^2(\emptyset) \subsetneq \psi_I^4(\emptyset) \subsetneq \ldots \subsetneq \psi_I^{\xi(I)}(\emptyset) \subsetneq \ldots \subsetneq \psi_I^3(\emptyset) \subsetneq \psi_I(\emptyset)$, so each element in the sequence has a different number of edges in it. However, each of the $\xi(I)+1$ elements has at least 0 edges and at most k, so the number of distinct sets of edges in the sequence can be at most $k+1$ by the pigeonhole principle. Consequentially, $\xi(I) \leq k$. $\square$

Corollary 4.50. *For an $n \times n$ instance I, $\xi(I) \leq n^2 - n + 1$ (i.e. $\Xi(n) \leq n^2 - n + 1$).*

That said, the above bound is far from tight. In this section, we find an exact value of $\Xi(n)$, thereby finding a tight upper bound on $\xi(I)$ for an $n \times n$ instance.

Theorem 4.51. *For all $n \geq 3$, $\Xi^*(n) = 2n - 3$.*

Theorem 4.52. *For all $n \geq 3$, $\Xi(n) = 2n - 3$.*

The proof of Theorem 4.51 will be postponed to Subsections 4.6.1 and 4.6.2, where we prove Lemma 4.55 and Lemma 4.57 respeectively. For the remainder of this section, we will show how to deduce Theorem 4.52 from Theorem 4.51. We begin with the following lemma:

Lemma 4.53. *Let I' be a completion of I. Then, $\xi(I) \leq \xi(I')$.*

Proof. Let k be the least element of $\mathbb{N}$ such that $\psi^k_{I'}(\emptyset) = \psi^\infty_{I'}$; by the definition of ξ, $\xi(I')$. By Corollary 4.24, this means that $\psi^\infty_I = \psi^\infty_{I'} \cap G(I) = \psi^k_{I'}(\emptyset) \cap G(I)$; however, by Corollary 4.23, $\psi^k_{I'}(\emptyset) \cap G(I) = \psi^k_I(\emptyset)$. Therefore, $\psi^k_I(\emptyset) = \psi^\infty_I$, so $\xi(I) \leq k = \xi(I')$. $\square$

We now can prove Theorem 4.52.

Proof. Since $2n - 3 = \Xi^*(n)$ by Theorem 4.51, this statement can be considered in two parts - namely, $\Xi(n) \geq \Xi^*(n)$, and $\Xi(n) \leq \Xi^*(n)$. To show that $\Xi(n) \geq \Xi^*(n)$, we note that $\Xi(n)$ is the maximum of $\xi(I)$ over all $n \times n$ instances I, whereas $\Xi^*(n)$ is the maximum of $\xi(I)$ over only the satisfactory $n \times n$ instances; consequentially, $\Xi(n) \geq \Xi^*(n)$.

To show that $\Xi(n) \leq \Xi^*(n)$, we consider any $n \times n$ instance I. By Lemma 4.53, there exists a complete $n \times n$ instance I' such that $\xi(I) \leq \xi(I')$. Since I' is complete - and thereby satisfactory - $\xi(I') \leq \Xi^*(n)$. $\Xi(n)$ is the maximum of $\xi(I)$ over all such I, so $\Xi(n) \leq \Xi^*(n)$. $\square$

Recall from Section 2.4 that [Wak10] gave an algorithm that, given an $n \times n$ instance I, finds the hub ψ^∞_I in $O(n^3)$. Theorem 4.52 allows us to give an alternative algorithm for this:

Theorem 4.54. *Given an $n \times n$ instance I, we may find $(\emptyset, \psi(\emptyset), \psi^2(\emptyset), \ldots, \psi^\infty)$ in $O(n^3)$ time.*

Proof. The first two terms of the sequence are trivially $\emptyset$ and $E(G(I))$. By Theorem 4.42, for $k \geq 2$, we can use $\psi^{k-2}(\emptyset)$ and $\psi^{k-1}(\emptyset)$ to construct $\psi^k(\emptyset)$ in $O(n^2)$ time; therefore, the sequence $(\emptyset, \psi(\emptyset), \psi^2(\emptyset), \ldots, \psi^{2n-3}(\emptyset))$ can be constructed in $(2n-3) * O(n^2) = O(n^3)$ time. By Theorem 4.52, the final term in the sequence is ψ^∞_I. $\square$

4.6.1 Finding a Lower Bound for Ξ^*

Since $\Xi^*(n)$ is the maximum of $\xi(I)$ over all satisfactory $n \times n$ instances I, we can show that $\Xi^*(n) \geq 2n-3$ by finding a family of satisfactory instances $\{I_n : n \in \{3,4,\ldots\}\}$ such that for each $n \in \mathbb{N}$, I_n is an $n \times n$ satisfactory instance with $\xi(I_n) = 2n-3$.

Lemma 4.55. *There exists a family of satisfactory instances $\{I_n : n \in \{3,4,\ldots\}\}$ such that for each $n \in \mathbb{N}$, I_n is an $n \times n$ instance with $\xi(I_n) = 2n-3$.*

Proof. We define each I_n as follows:

- The set of mentors is $\{m_1, m_2, \ldots, m_n\}$ and the set of students is $\{d_1, d_2, \ldots, d_n\}$.
- The preference list of m_1 is $[d_1]$.
- For all $i \in \{2,3\}$, the preference list of m_i is $[d_i, d_{i-1}, d_{i+1}]$.
- For all $i \in \{4,5,\ldots,n-1\}$, the preference list of m_i is $[d_i, d_{i-1}, d_2, d_{i+1}]$.
- The preference list of m_n is $[d_n, d_{n-1}, d_2]$.
- The preference list of d_1 is $[m_2, m_1]$.
- The preference list of d_2 is $[m_n, m_{n-1}, \ldots, m_2]$.
- For all $i \in \{3,4,\ldots,n-1\}$, the preference list of d_i is $[m_{i+1}, m_{i-1}, m_i]$.
- The preference list of d_n is $[m_{n-1}, m_n]$.

Trivially, $\psi_{I_n}(\emptyset) = G(I_n)$; by using the Gale-Shapley algorithm in [GS62], we see that the mentor-optimal and student optimal stable matchings over I_n are both $\{(m_i, d_i) : i \in [n]\}$, so this is the only stable matching over I_n and $\psi^2_{I_n}(\emptyset) = \{(m_i, d_i) : i \in [n]\}$. We can further find via induction the structure of $\psi^k_{I_n}(\emptyset)$ for all $k \geq 1$. For $k \geq 2$, we define $E_k, E'_k \subseteq E(G(I_n))$ as follows:

$$E_k = \{(m_i, d_i) : i \in [n]\} \cup \{(m_i, d_2) : i \in \{3, \ldots, k\}\} \cup \{(m_{i-1}, d_i) : i \in \{3, \ldots, k\}\};$$

$$E'_k = E_n \cup \{(m_i, d_{i-1}) : i \in \{k, \ldots, n\}\}.$$

Lemma 4.56. *For all* $k \in \{2, \ldots, n\}$, $\psi_{I_n}^{2k-3}(\emptyset) = E'_k$ *and* $\psi_{I_n}^{2k-2}(\emptyset) = E_k$. *Furthermore, the mentor-optimal* $(2k-3)$*-stable matching is* $\{(m_i, d_i) : i \in [n]\}$, *and the student-optimal* $(2k-3)$*-stable matching is:*

$$\{(m_1, d_1), (m_2, d_3), \ldots, (m_{k-1}, d_k), (m_k, d_2), (m_{k+1}, d_{k+1}), \ldots, (m_n, d_n)\}.$$

Proof. We prove this result by induction on k. For the base case, when $k = 2$, we note that $\psi_{I_n}(\emptyset) = E(G(I_n)) = E'_2$ trivially. In addition, by applying the Gale-Shapley algorithm to I_n, we see that the mentor-optimal and student-optimal 1-stable matching is $\{(m_i, d_i) : i \in [n]\}$. As a consequence, this is the only 1-stable matching over I_n, and so $\psi_{I_n}^2(\emptyset) = \{(m_i, d_i) : i \in [n]\} = E_2$.

Now, for the inductive step, assume that for some $k \in \{2, \ldots, n-1\}$, $\psi_{I_n}^{2k-2}(\emptyset) = E_k$, the mentor-optimal $(2k-3)$-stable matching is $\{(m_i, d_i) : i \in [n]\}$, and the student-optimal $(2k-3)$-stable matching M_1 is:

$$\{(m_1, d_1), (m_2, d_3), \ldots, (m_{k-1}, d_k), (m_k, d_2), (m_{k+1}, d_{k+1}), \ldots, (m_n, d_n)\}.$$

In particular, we note that by Theorem 4.9, $\psi_{I_n}^{2k-3}(\emptyset) \supseteq \psi_{I_n}^{2k-1}(\emptyset) \supseteq \psi_{I_n}^{2k-2}(\emptyset)$, so by the proofs of Theorem 4.38 and Theorem 4.39, we see that if $I' = I[\psi^{2k-3}(\emptyset)]$, then $\psi_{I_n}^{2k-2}(\emptyset) = \psi_{I'}^2(\emptyset)$ and $\psi_{I_n}^{2k-1}(\emptyset) = \psi_{I'}^3(\emptyset)$. By applying Theorem 4.46 to I', we see that $\psi^{2k-1}(\emptyset) = \psi^{2k-2}(\emptyset) \cup P \cup P^*$, where P is the union of all perfect matchings over E (the edges (m_i, d_j) where m_i prefers $p_{M_1}(m_i)$ to d_j and d_j prefers m_i to $p_{M_1}(d_j)$), and P^* is the union of all perfect matchings over E^* (the edges (m_i, d_j) where m_i prefers d_j to d_i and d_j prefers m_j to m_i). We note that $P^* = E^* = \{(m_1, d_1), \ldots, (m_n, d_n)\}$ trivially. In addition, it is straightforward to see that $E = \{(m_1, d_1)\} \cup \{(m_i, d_2) : i \in \{k, \ldots, n\}\} \cup \{(m_{i-1}, d_i) : i \in \{3, \ldots, n\}\} \cup \{(m_i, d_{i-1}) : i \in \{k+1, \ldots, n\}\}$, with the additional edge (m_2, d_1) if $k = 2$; as a result, $P = \{(m_1, d_1)\} \cup \{(m_i, d_2) : i \in \{k, \ldots, n\}\} \cup \{(m_{i-1}, d_i) : i \in \{3, \ldots, n\}\} \cup \{(m_i, d_{i-1}) : i \in \{k+1, \ldots, n\}\}$. (Any perfect matching over E must have m_1 partnered with d_1, since d_1 is m_1's only available partner.) Therefore, $\psi^{2(k+1)-3}(\emptyset) = \psi^{2k-2}(\emptyset) \cup P \cup P^* = E_k \cup \{(m_i, d_2) : i \in \{k+1, \ldots, n\}\} \cup \{(m_{i-1}, d_i) : i \in \{k+1, \ldots, n\}\} \cup \{(m_i, d_{i-1}) : i \in \{k+1, \ldots, n\}\}$ (by the inductive assumption) $= E'_{k+1}$.

By Theorem 4.38, $\psi_{I_n}^{2k}(\emptyset) = \psi_{I_n[E'_{k+1}]}^{2}(\emptyset)$. We may then apply the algorithm for finding the set of stable matchings over an instance from [GS85] in order to see that the mentor-optimal $2k-3$-stable matching is $\{(m_i, d_i) : i \in [n]\}$, the student-optimal $2k-3$-stable matching is $\{(m_1, d_1), (m_2, d_3), \ldots, (m_{k-1}, d_k), (m_k, d_2), (m_{k+1}, d_{k+1}), \ldots, (m_n, d_n)\}$, and $\psi_{I_n}^{2(k+1)-2}(\emptyset) = \psi_{I_n}^{2k}(\emptyset) = E_{k+1}$. By induction, we are done. $\square$

As seen by the above lemma, $\psi_{I_n}^{2n-4}(\emptyset) \neq \psi_{I_n}^{2n-3}(\emptyset) = \psi_{I_n}^{2n-2}(\emptyset)$. By Theorem 4.1 $\psi_{I_n}(S) = S$ iff $S = \psi_{I_n}^{\infty}$, so $\xi(I_n) = 2n-3$ by the definition of ξ. $\square$

4.6.2 The Upper Bound of Ξ^*

Since we have shown in the previous section that $\Xi^*(n) \geq 2n-3$ for all $n \geq 3$, to prove Theorem 4.51, we only need to show that the following lemma is true:

Lemma 4.57. *For all $n \geq 3$, $\Xi^*(n) \leq 2n-3$.*

We will ultimately prove Lemma 4.57 by induction on n, so we initially consider the base case for such an induction argument.

Lemma 4.58. $\Xi^*(3) = 3$.

Proof. We use a Maple program to compute $\xi(I')$ for every complete 3×3 instance I', and confirm that the maximum value of $\xi(I')$ for such instances is 3. However, every satisfactory 3×3 instance I can be extended to a completion I' with $\xi(I') \geq \xi(I)$ by Lemma 4.53; as a result of this and Theorem 2.6, we see that $\Xi^*(3)$ is the maximum of $\xi(I)'$ over all complete 3×3 instances - i.e. 3. $\square$

We now consider some lemmas that we can use to construct an inductive argument. For such purposes, we note the following results.

Proposition 4.59. *Let I be any instance, and $I' = I[\psi_I^3(\emptyset)]$. Then, for all positive $k \in \mathbb{N}$, $\psi_{I'}^{k}(\emptyset) = \psi_I^{k+2}(\emptyset)$.*

Proof. We prove this by induction on k. For our base case, when $k = 1$, $\psi_{I'}(\emptyset) = G(I') = \psi_I^3(\emptyset)$.

Now, for any $k_0 \in \mathbb{N}$, assume that $\psi_{I'}^{k_0}(\emptyset) = \psi_I^{k_0+2}(\emptyset)$; we aim to show that $\psi_{I'}^{k_0+1}(\emptyset) = \psi_I^{k_0+3}(\emptyset)$. Since $k_0 \geq 1$, $\psi_I^{k_0+3}(\emptyset) \subseteq \psi_I^3(\emptyset) = G(I')$. Meanwhile, $\psi_{I'}^{k_0+1}(\emptyset) \subseteq G(I') \subseteq G(I)$, so we only need to show that any given edge in $G(I')$ is in $\psi_{I'}^{k_0+1}(\emptyset)$ iff it is in $\psi_I^{k_0+3}(\emptyset)$.

Let $e \in G(I')$. If $e \in \psi_{I'}^{k_0+1}(\emptyset)$, there exists a $\psi_{I'}^{k_0}(\emptyset)$-stable matching M over I'. This matching remains $\psi_{I'}^{k_0}(\emptyset)$-stable over I, and so by substitution is $\psi_I^{k_0+2}(\emptyset)$-stable; by the definition of ψ_I, $e \in \psi_I^{k_0+3}(\emptyset)$. Conversely, if $e \in \psi_I^{k_0+3}(\emptyset)$, there exists a $\psi_I^{k_0+2}(\emptyset)$-stable matching M over I'; by substitution, M is $\psi_{I'}^{k_0}(\emptyset)$-stable over I. Since $M \subseteq \psi_I^{k_0+3}(\emptyset) \subseteq \psi_I^3(\emptyset) = G(I')$, it consists only of edges in I'; consequentially, M is a matching over I', and preserves the property of being $\psi_{I'}^{k_0}(\emptyset)$-stable over I'. By the definition of $\psi_{I'}$, this means that $e \in \psi_{I'}^{k_0+1}(\emptyset)$.

As a result, $\psi_{I'}^{k_0+1}(\emptyset) = \psi_I^{k_0+3}(\emptyset)$, and we have shown our inductive step. By induction, $\psi_{I'}^{k}(\emptyset) = \psi_I^{k+2}(\emptyset)$ for all positive $k \in \mathbb{N}$. □

Corollary 4.60. *Let I be any instance such that $\xi(I) \geq 3$, and $I' = I[\psi_I^3(\emptyset)]$. Then, $\xi(I) = \xi(I') + 2$.*

We also need the following lemma, which we prove in Subsection 4.6.3.

Lemma 4.61. *Let I_1 and I_2 be two instances on disjoint sets of vertices, and I be the instance with vertex set $V(I_1) \cup V(I_2)$, where each vertex from I_1 and I_2 has the same preference list as in I_1 and I_2 respectively. Then, $\xi(I) = max\{\xi(I_1), \xi(I_2)\}$.*

We now proceed to the proof of Lemma 4.57. We recall that a matching is k-stable over I if it is $\psi_I^k(\emptyset)$-stable.

Proof. We prove this result by induction on n. For our base case, when $n = 3$, the statement is equivalent to Lemma 4.58.

Now, for our inductive step, suppose that, for a given $n \geq 3$, $\Xi^*(n) \leq 2n - 3$; we need to show that $\Xi(n+1) \leq 2n-1$. Let I be an arbitrary satisfactory $(n+1) \times (n+1)$ instance, with M_1 and M_2 as the mentor-optimal and student-optimal stable matchings respectively. It is sufficient to show that $\xi(I) \leq 2n - 1$.

We may consider the following subinstances: $I_1 \equiv I_{(\emptyset, M_1)}$, $I_2 \equiv I_{(M_2, \emptyset)}$, and $I_3 \equiv I_{(M_1, M_2)}$. (Note that these subinstances are still satisfactory - M_1 is a perfect matching that is stable over I_1 and I_3, and M_2 is a perfect matching that is stable over I_2.) We note that $\psi^k_{I'}(\emptyset) = \psi^k_I(\emptyset) \cap G(I')$ for any $k \in \mathbb{N}$ and $I' \in \{I_1, I_2, I_3\}$ by Theorem 4.32, Corollary 4.33, and Corollary 4.34 respectively. In addition, every edge $e \in G(I)$ that doesn't appear in $G(I_1)$, $G(I_2)$, or $G(I_3)$ must fit into one of four categories:

1. m_e prefers $p_{M_1}(m_e)$ to d_e and d_e prefers $p_{M_1}(d_e)$ to m_e.

2. m_e prefers d_e to $p_{M_1}(m_e)$ and d_e prefers m_e to $p_{M_1}(d_e)$.

3. m_e prefers $p_{M_2}(m_e)$ to d_e and d_e prefers $p_{M_2}(d_e)$ to m_e.

4. m_e prefers d_e to $p_{M_2}(m_e)$ and d_e prefers m_e to $p_{M_2}(d_e)$.

Any edge in category 2 or 4 would destabilize M_1 or M_2 respectively, so no such edge can exist. There can exist edges that appear in category 1 or 3; however, we can make the following observation about them.

Lemma 4.62. *Let I be any instance, and S be the set of all edges (m, d) with the property that there exists a stable matching M over I such that m strictly prefers $p_M(m)$ to d and d strictly prefers $p_M(d)$ to m. Then, for every set of edges E such that $\psi^2(\emptyset) \subseteq E \subseteq G(I)$, $S \cap \psi(E) = \emptyset$.*

Proof. We first show that $\psi^3(\emptyset)$ contains no element of E by contradiction. Assume that there exists some $e \in E$ such that $e \in \psi^3(\emptyset)$; then, there must be a 2-stable matching M_e that contains E. Since $M \subseteq \psi^2(\emptyset)$, M_e is also M-stable. M is a stable matching, so it is M_e-stable, implying that M and M_e are costable; this means that m_e prefers $p_{M_e}(m_e) = d_e$ to $p_M(m_e)$ iff d_e prefers $p_M(d_e)$ to $p_{M_e}(d_e) = m_e$. This contradicts the fact that m_e and d_e prefer their respective partners in M to each other, so no such e can exist.

For any $E \supseteq \psi^2(\emptyset)$, $\psi(E) \subseteq \psi^3(\emptyset)$, by Theorem 4.9, so $S \cap \psi(E) \subseteq S \cap \psi^3(\emptyset) = \emptyset$. $\square$

As a result, no edge in category 1 or 3 appears in $\psi(E)$ for any $E \supseteq \psi^2(\emptyset)$; however, $\psi^i(\emptyset) \supseteq \psi^2(\emptyset)$ for all $i \geq 1$, implying that no such edge appears in $\psi^k(\emptyset)$ for all

$k \geq 2$. As such, either $\xi(I) \leq 1$, or $\xi(I) = max\{\xi(I_1), \xi(I_2), \xi(I_3)\}$. We will show that $\xi(I') \leq 2n - 1$ for all $I' \in \{I_1, I_2, I_3\}$.

To show that $\xi(I_1) \leq 2i - 1$, we note that G_{I_1} contains exactly the edges in I over which a proposal is made during the mentor-optimal Gale-Shapley algorithm; therefore, performing the mentor-optimal Gale-Shapley algorithm proceeds in exactly the same way in I_1 as in I, and the resulting mentor-optimal stable matching M_1 has every student partnered with their top partner. As a result, M_1 is also the student-optimal (and therefore only) stable matching, and so $\psi^2_{I_1}(\emptyset) = M_1$. Let d_0 be any student that is proposed to last in some procedure of the mentor-optimal Gale-Shapley algorithm.

Lemma 4.63. *$(p_{M_1}(d_0), d_0) \in \psi^3_{I_1}(\emptyset)$, and no other edge $\in \psi^3_{I_1}(\emptyset)$ is incident with d_0 or $p_{M_1}(d_0)$.*

Proof. In the aforementioned procedure of the Gale-Shapley algorithm, d_0 does not reject a previous suitor in response to the final proposal - otherwise, the rejected suitor would make a new proposal right after, since the Gale-Shapley algorithm only terminates on a satisfactory instance when every vertex has a partner. As a result, d_0 has only one possible partner in I_1, and since M_1 is a perfect matching, this partner is $p_{M_1}(d_0)$.

Since M_1 is a perfect matching, every 2-stable matching over I_1 is perfect by Theorem 4.18. As a result, every such matching contains $(p_{M_1}(d_0), d_0)$ as an edge, and so this is the only edge in $\psi^3_{I_1}(\emptyset)$ that contains either of $p_{M_1}(d_0)$ and d_0. □

As a result, $\psi^3_{I_1}(\emptyset)$ is the vertex-disjoint union of $\{(p_{M_1}(d_0), d_0)\}$ and $G' \equiv \psi^3_{I_1}(\emptyset) - \{(p_{M_1}(d_0), d_0)\}$. If $I' = I_1[\psi^3_{I_1}(\emptyset)]$, then, by Corollary 4.60:

$$\xi(I') = max\{\xi(I'[\{(p_{M_1}(d_0), d_0)\}]), \xi(I'[G'])\}.$$

However, both of these instances are satisfactory; I_{d_0} is a 1×1 instance and $I_{G'}$ is a $n \times n$ instance, so $\xi(I_{d_0}) = 1$ and $\xi(I_{G'}) \leq 2n - 3$ by our inductive assumption. This implies that $\xi(I') \leq 2n - 3$; by Lemma 4.61, either $\xi(I_1) \leq 2 \leq 2n - 1$ (as $n \geq 3$), or $\xi(I_1) = \xi(I') + 2 \leq 2n - 1$. In either case, $\xi(I_1) \leq 2n - 1$.

By a similar argument, we may show that $\xi(I_2) \leq 2n - 1$. Finally, I_3 is an instance where the mentor-optimal matching has every mentor partnered with their top preference, and the student-optimal matching has every student partnered with their top preference. By Corollary 4.21, $\xi(I_3) \leq 2 \leq 2n - 1$ (since $n \geq 3$). As such, $\xi(I) \leq max\{2n-1, 2n-1, 2n-1\} = 2n-1$; however, I is an arbitrary satisfactory $(n+1) \times (n+1)$ instance, so $\Xi^*(n+1) \leq 2n-1$.

Thus, we have shown that $\Xi^*(3) = 3 = 2 * 3 - 3$, and that $\Xi^*(n) \leq 2n - 3 \Rightarrow \Xi^*(n+1) \leq 2n - 1 = 2(n+1) - 3$ for all $n \geq 3$. By induction, $\Xi^*(n) \leq 2n - 3$ for all $n \geq 3$. □

4.6.3 A Proof of Lemma 4.61

As noted previously, our proof of Lemma 4.57 requires Lemma 4.61. In this subsection, we prove this lemma.

Proposition 4.64. *Let I_1 and I_2 be two instances on disjoint sets of vertices, and I be the instance with vertex set $V(I_1) \cup V(I_2)$, where each vertex from I_1 and I_2 has the same preference list as in I_1 and I_2 respectively. Then, for all $S_1 \subseteq G(I_1)$ and $S_2 \subseteq G(I_2)$, $\psi_I(S_1 \cup S_2) = \psi_{I_1}(S_1) \cup \psi_{I_2}(S_2)$.*

Proof. We prove this by showing that the set of $S_1 \cup S_2$-stable matchings over I is the set of every union of an S_1-stable matching over I_1 and an S_2-stable matching over I_2. If M_1 is an S_1-stable matching over I_1 and M_2 is an S_2-stable matching over I_2, then these matchings are S_1-stable and S_2-stable over I, respectively. Since M_1 and M_2 are vertex-disjoint, their union is a matching and partners each vertex with its preferred partner over M_1 and M_2; consequentially, an edge can only destabilize $M_1 \cup M_2$ if it destabilizes both M_1 and M_2. No edge in S_1 destabilizes M_1, and no edge in S_2 destabilizes M_2, so $M_1 \cup M_2$ is $S_1 \cup S_2$-stable. As such, any union of an S_1-stable matching over I_1 and an S_2-stable matching over I_2 is an S-stable matching over I.

Now, let M be any $S_1 \cup S_2$-stable matching over I. We define $M_1 \equiv M \cap G(I_1)$ and $M_2 \equiv M \cap G(I_2)$; since $G(I)$ is the disjoint union of $G(I_1)$ and $G(I_2)$, M is the disjoint union of M_1 and M_2. For every $e \in S_1$, $(m_e, p_M(m_e)), (p_M(d_e), d_e) \in G(I_1)$ (as I_1 and

I_2 are vertex-disjoint); furthermore, by the fact that M is S_1-stable, at least one of m_e and d_e prefers their partner in M to the other. These partners are preserved in M_1, so M_1 remains e-stable. Since E is any edge in S_1, M_1 is S_1-stable over I, and therefore S_1-stable over I_1. Similary, M_2 is S_2-stable over I_2, and so M must be a union of an S_1-stable matching over I_1 and an S_2-stable matching over I_2.

As a result, the set of $S_1 \cup S_2$-stable matchings over I is the set of every union of an S_1-stable matching over I_1 and an S_2-stable matching over I_2. This implies that $\psi_I(S_1 \cup S_2) = \psi_{I_1}(S_1) \cup \psi_{I_2}(S_2)$. □

We may now prove Lemma 4.61.

Proof. We consider the values $k \in \mathbb{N}$ such that $\psi_I^k(\emptyset)$ is a hub. We set $S_1 \equiv \psi_I^k(\emptyset) \cap G(I_1)$ and $S_2 \equiv \psi_I^k(\emptyset) \cap G(I_2)$; by Proposition 4.64, $\psi_I(\psi_I^k(\emptyset)) = \psi_{I_1}(S_1) \cup \psi_{I_2}(S_2)$. Since $\psi_{I_1}(S_1) \subseteq G(I_1)$ and $\psi_{I_2}(S_2) \subseteq G(I_2)$, this equals $\psi_I^k(\emptyset)$ iff $\psi_{I_1}(S_1) = \psi_I^k(\emptyset) \cap G(I_1) = S_1$ and $\psi_{I_2}(S_2) = \psi_I^k(\emptyset) \cap G(I_2) = S_2$. This happens iff k is greater than or equal to both $\xi(I_1)$ and $\xi(I_2)$, so the minimum such k - i.e. $\xi(I)$ - is $max\{\xi(I_1), \xi(I_2)\}$. □

4.7 An Improvement to Theorem 4.52 for Nonsatisfactory Instances

In the previous section, we showed that if I is an $n \times n$ instance with $n \geq 3$, then $\xi(I) \leq 2n - 3$; furthermore, this upper bound is tight. However, if I is very far from complete, then we may be able to show that $\xi(I)$ is significantly smaller than $2n - 3$. In this section, we will show that if I is not satisfactory, then we can improve our upper bound on $\xi(I)$. Similarly, in the next section, we will show that if $G(I)$ is sparce, then we can make alternate improvements to our upper bound on $\xi(I)$.

Theorem 4.65. *If a vertex v has degree 0 in $\psi^2(\emptyset)$, then it has degree 0 in $\psi^k(\emptyset)$ for all $k \geq 2$.*

Proof. For every $k \geq 2$, $\psi^k(\emptyset)$ is the union of all $k-1$-stable matchings. Since $k-1 \geq 1$, every $k-1$-stable matching covers the same vertices as the 1-stable matchings by Theorem 4.18. As a result, $\psi^k(\emptyset)$ includes no edge in v iff no stable matching covers v - which occurs iff no edge covers v in $\psi^2(\emptyset)$. □

As a result, we see that if every stable matching over I has k edges, then $I[\psi^3(\emptyset)]$ is a $k \times k$ instance with some number of isolated vertices (by Theorem 2.4, we know that every stable matching covers the same k mentors and k students). This intuition on $I[\psi^3(\emptyset)]$ can be leveraged to say something about I using Corollary 4.60.

Theorem 4.66. *Let I be any instance, and M be any stable matching over I. Then, if $|M| \geq 2$, $\xi(I) \leq 2|M| - 1$.*

Proof. If $|M| = 2$, we may assume WLOG that I has mentors $\{m_1, m_2, \dots, m_{n_1}\}$ and students $\{d_1, d_2, \dots, d_{n_2}\}$, and $M = \{(m_1, d_1), (m_2, d_2)\}$ is a stable matching. We note that $\{(m_1, d_2), (m_2, d_1)\}$ is the only other possible perfect matching on mentors $\{m_1, m_2\}$ and students $\{d_1, d_2\}$. In addition, all of $\psi_I^2(\emptyset)$, $\psi_I^3(\emptyset)$, and ψ_I^∞ are unions of such perfect matchings by Theorem 4.18, and must contain the stable matching $\{(m_1, d_2), (m_2, d_2)\}$; this means that the only possibilities for these sets are:

- $\{(m_1, d_2), (m_2, d_2)\}$
- $\{(m_1, d_2), (m_1, d_2), (m_2, d_1), (m_2, d_2)\}$

By the pigeonhole principle, some pair of $\psi_I^2(\emptyset)$, $\psi_I^3(\emptyset)$, and ψ_I^∞ are equal. However, if $\psi_I^3(\emptyset) \neq \psi_I^\infty$, then $\psi_I^2(\emptyset)$ must be distinct from both of them, creating a contradiction. Since $\psi_I^3(\emptyset) = \psi_I^\infty$ thereby, $\xi(I) \leq 3 = 2|M| - 1$.

Now, let us consider the case when $|M| \geq 3$. If $\xi(I) \leq 3$, then the statement obviously holds. Otherwise, we define the instance I^* to be the restriction of I such that $G(I^*) = \psi_I^3(\emptyset)$. As is shown in Theorem 4.65, I^* is the union of an $|M| \times |M|$ instance I' with the same vertex set as M, and some number of isolated vertices with empty preference lists; as a consequence of Proposition 4.64, $\xi(I^*) = \xi(I')$. By Theorem 4.52, $\xi(I') \leq 2|M| - 3$ (since $|M| \geq 3$). This implies by Corollary 4.60 that $\xi(I) = \xi(I^*) + 2 = \xi(I') + 2 \leq 2|M| - 1$. □

4.8 The Convergence Rate of ψ for Sparse Instances

By Proposition 4.49, for an instance I such that $E(G(I)) = k$ and every stable matching has size q, $\xi(I) \leq k - q + 1$. Here, we will improve on this upper bound for the case when $k < 4q - 5$; this will allow us to improve on Theorem 4.52 and Theorem 4.66 for any instance I where $G(I)$ is sufficiently sparce.

Theorem 4.67. *If the lattice of hub-stable matchings $\mathcal{L}_h$ has $r-1$ join-irreducible elements, then $|\psi_I^\infty| \geq q + 2(r-1)$.*

Proof. Since the lattice of hub-stable matchings is a distributive lattice with $r-1$ join-irreducible elements, we can find a chain of length r in the lattice. The least element of this chain - the mentor-optimal hub-stable matching - contains q edges, and each subsequent element contains at least 2 edges that were not in any previous term (since it differs from the next-most student-optimal matching by performing a rotation that matches at least 2 students with strictly more desired partners). Each edge in such a matching must appear in K, so $|K| \geq n + 2(r-1)$. $\square$

Now, we can consider the lattices $\{\mathcal{L}_{\psi(\emptyset)}, \mathcal{L}_{\psi^3(\emptyset)}, \ldots, \mathcal{L}_{\psi^{2i+1}(\emptyset)}, \ldots\}$. Since these lattices are the lattices of S-stable matchings, where S decreases as the sequence goes on, each element of the sequence is a sublattice of the previous; as such, each lattice in the sequence has at least as many join- irreducible elements as the previous lattice.

Lemma 4.68. *If $\mathcal{L}_{\psi^{2i-1}(\emptyset)}$ and $\mathcal{L}_{\psi^{2i+1}(\emptyset)}$ both have r join-irreducible elements, then $\psi^{2i}(\emptyset) = \psi^\infty$.*

Proof. Since $\mathcal{L}_{\psi^{2i-1}(\emptyset)}$ is a distributive lattice with r join-irreducible elements, we can find a length $r+1$ maximal chain in it; since $\mathcal{L}_{\psi^{2i-1}(\emptyset)} \subseteq \mathcal{L}_{\psi^{2i+1}(\emptyset)}$, this chain must also exist in $\mathcal{L}_{\psi^{2i+1}(\emptyset)}$. However, since it is a chain of length $r+1$ in a distributive lattice with r join-irreducible elements, it must also be maximal in $\mathcal{L}_{\psi^{2i+1}(\emptyset)}$. By Theorem 3.10 and Corollary 2.15, the elements of this chain contain every edge that appears in at least one element of $\mathcal{L}_{\psi^{2i+1}(\emptyset)}$. Each element in the chain also appears in $\mathcal{L}_{\psi^{2i-1}(\emptyset)}$, so, by the definition of ψ, $\psi^{2i}(\emptyset) \supseteq \psi^{2i+2}(\emptyset)$.

However, since $\{\psi^{2j}(\emptyset) : j \in \mathbb{N}\}$ is an increasing sequence, $\psi^{2i}(\emptyset) \subseteq \psi^{2i+2}(\emptyset)$; therefore, $\psi^{2i}(\emptyset) = \psi^{2i+2}(\emptyset)$. By Theorem 4.1, this implies that $\psi^{2i}(\emptyset) = \psi_I^\infty$. $\square$

Corollary 4.69. *If $\mathcal{L}_h$ has at most r join-irreducible elements, then $\xi(I) \leq 2r + 2$.*

Proof. For all i, $\mathcal{L}_{\psi^{2i+1}(\emptyset)} \subseteq \mathcal{L}_K$, so each such lattice has at most r join-irreducible elements. (Since they are all nonempty, they also contain at least 1.) If $\psi^{2r+2}(\emptyset)$ was not a hub, this would imply that $\mathcal{L}_{\psi(\emptyset)}, \mathcal{L}_{\psi^3(\emptyset)}, \ldots, \mathcal{L}_{\psi^{2r+3}(\emptyset)}$ all have a different number of join-irreducible elements; however, this gives $r+2$ different lattices, each with a number of join-irreducible elements in $[r] \cup \{0\}$. By the pigeonhole principle, we have a contradiction, so $\psi^{2r+2}(\emptyset)$ is a hub, and $\xi(I) \leq 2r+2$ by Theorem 4.1. $\square$

Theorem 4.70. *For an instance I such that every stable matching over I has k edges and $|E(G(I))| = b$, $\xi(I) \leq \frac{2}{3}(b-k+2)$.*

Proof. Setting $r = \lceil \frac{\xi(I)}{2} \rceil$ gives us that $\psi^{2(r-1)}(\emptyset) \neq \psi^{\infty}$. By the contrapositive of Corollary 4.69, $\mathcal{L}_h$ has at least $r-1$ join-irreducible elements, which means that ψ^{∞} has at least $k+2(r-1)$ edges by Theorem 4.67. However, for each $i \in [\lfloor \frac{\xi(I)}{2} \rfloor]$, $\psi^{2i-1}(\emptyset)$ has a different number of edges, each of which is greater than the number in ψ^{∞}; consequentially, the largest of them has at least $k+2r-2+\lfloor \frac{\xi(I)}{2} \rfloor$ edges, and so $b \geq k-2+\lceil \frac{3\xi(I)}{2} \rceil$. As a result, $\frac{3\xi(I)}{2} \leq \lceil \frac{3\xi(I)}{2} \rceil \leq b-k+2$, so $\xi(I) \leq \frac{2}{3}(b-k+2)$. $\square$

Combining this result with Theorem 4.51 and Theorem 4.66, we see that for an $n \times n$ instance I such that $G(I)$ has b edges and any stable matching M over I has k edges, $\xi(I) \leq \min(2n-3, 2k-1, \lfloor \frac{2}{3}(b-k+2) \rfloor)$. In our final result, we show an instance where this is tight on all three measurements.

Example 4.71. *For any integer $n \geq 3$, we define I'_n as follows:*

- $V_m(I'_n) = \{m_1, m_2, \ldots, m_n\}$ *and* $V_d(I'_n) = \{d_1, d_2, \ldots, d_n\}$.
- *The preference list of m_1 is empty.*
- *For all $i \in \{2, 3\}$, the preference list of m_i is $[d_i, d_{i-1}, d_{i+1}]$.*
- *For all $i \in \{4, 5, \ldots, n-1\}$, the preference list of m_i is $[d_i, d_{i-1}, d_2, d_{i+1}]$.*
- *The preference list of m_n is $[d_n, d_{n-1}, d_2]$.*
- *The preference list of d_1 is $[m_2]$.*
- *The preference list of d_2 is $[m_n, m_{n-1}, \ldots, m_2]$.*

- *For all* $i \in \{3, 4, \ldots, n-1\}$*, the preference list of* d_i *is* $[m_{i+1}, m_{i-1}, m_i]$*.*

- *The preference list of* d_n *is* $[m_{n-1}, m_n]$*.*

We note that I'_n is the same as I_n from Lemma 4.55, with the edge (m_1, d_1) removed; it is straightforward to see that $\xi(I'_n) = \xi(I_n)$, and so $\xi(I_n) = 2n - 3$. Furthermore, $b = |G(I)| = 4n-7$ and the stable matching $\{(m_2, d_2), \ldots, (m_n, d_n)\}$ has $k = n-1$ edges, so $2k-1 = 2(n-1)-1 = 2n-3 = \xi(I_n)$, and $\lfloor \frac{2}{3}(b-k+2) \rfloor = \lfloor \frac{2}{3}(3n-4) \rfloor = 2n-3 = \xi(I_n)$.

Chapter 5

Representations of Lattice Flags

Given a stable matching instance I, there are a number of ways that we can associate I with a distributive lattice $\mathcal{L}$. The standard way is to associate I with $\mathcal{L}_s(I)$, the lattice of stable matchings over I; another way is by associating I with $\mathcal{L}_h(I)$, the lattice of hub-stable matchings. Furthermore, every distributive lattice is isomorphic to $\mathcal{L}_s(I)$ for some (non-unique) instance I, and $\mathcal{L}_h(I')$ for some (non-unique) instance I'. However, for a single instance, $\mathcal{L}_s(I)$ and $\mathcal{L}_h(I)$ are not independent structures, as noted by Theorem 4.15.

We define a **lattice flag** to be a pair $(\mathcal{L}_0, \mathcal{L}_1)$ of distributive lattices such that $\mathcal{L}_0$ is a sublattice of $\mathcal{L}_1$; more generally, we define a **lattice z-flag** to be a sequence $(\mathcal{L}_0, \mathcal{L}_1, \ldots, \mathcal{L}_z)$ of distributive lattices such that $\mathcal{L}_{r-1} \subseteq \mathcal{L}_r$ for all $r \in [z]$. (In particular, a lattice flag is a lattice 1-flag.) We also define a lattice z-flag to be **covering** if $\mathcal{L}_{r-1}$ is a cover-preserving sublattice of $\mathcal{L}_r$ for all $r \in [z]$. Two lattice z-flags $(\mathcal{L}_0, \ldots, \mathcal{L}_z)$ and $(\mathcal{L}'_0, \ldots, \mathcal{L}'_z)$ are **isomorphic** if there exists an order-preserving bijection $\zeta : \mathcal{L}_z \to \mathcal{L}'_z$ such that $\zeta(\mathcal{L}_i) = \mathcal{L}'_i$ for all $i \in \{0, \ldots, i-1\}$.

It is natural to ask for what lattice flags $(\mathcal{L}_s, \mathcal{L}_h)$ we can find an instance I such that $(\mathcal{L}_s, \mathcal{L}_h)$ is isomorphic to $(\mathcal{L}_s(I), \mathcal{L}_h(I))$. By Theorem 4.15, $(\mathcal{L}_s(I), \mathcal{L}_h(I))$ is a covering lattice flag. In this chapter, we will show that this is the only constraint on the structure of this lattice flag.

Theorem 5.1. *Let $(\mathcal{L}_s, \mathcal{L}_h)$ be any covering lattice flag. Then, there exists an instance I such that $(\mathcal{L}_s(I), \mathcal{L}_h(I))$ is isomorphic to $(\mathcal{L}_s, \mathcal{L}_h)$.*

There are other ways to associate lattice flags to a stable matching instance, which give rise to similar representation questions which will also be considered in this chapter.

5.1 Representation Theorems for Lattice Flags

In preparation for proving Theorem 5.1, we review representation theorems for lattice flags that are analogous to the Birkhoff Representation Theorem ([Sig14], [RS]). We define a **pointed order** $(P, \leq)$ as a poset with a minimum element $\hat{0}_P$ and a maximum element $\hat{1}_P$ - in other words, P is a finite set of elements (including $\hat{0}_P$ and $\hat{1}_P$) and $\leq$ is a binary relation that obeys the reflexive, antisymmetric, and transitive properties such that for all $p \in P$, $\hat{0}_P \leq p \leq \hat{1}_P$. (In cases where P is implied, we shorten $\hat{0}_P$ to $\hat{0}$ and $\hat{1}_P$ to $\hat{1}$.) [1]

A **pointed quasi-order** $(P, \leq^*)$ is defined in the same way, except that we no longer require that the binary relation be antisymmetric (i.e. we can have distinct $p_1, p_2 \in P$ such that $p_1 \leq p_2$ and $p_2 \leq p_1$). The elements of a quasi-order split into equivalence classes, where each equivalence class consists of some $p \in P$ and all $p' \in P$ such that $p \leq^* p'$ and $p' \leq^* p$; we note that $\leq^*$ induces a pointed order on the equivalence classes. (In particular, a pointed order is a pointed quasi-order where every equivalence class has one element.) An **extension** $(P, \leq^{**})$ of $(P, \leq^*)$ is a pointed quasi-order where $\leq^{**}$ is at least as strong as $\leq^*$ - i.e. if $p_1, p_2 \in P$ and $p_1 \leq^* p_2$, then $p_1 \leq^{**} p_2$.

Proposition 5.2. *Given a sequence of pointed quasi-orders $(P, \leq^0), \ldots, (P, \leq^z)$ such that for all $i \in [z]$, $(P, \leq^{i-1})$ is an extension of $(P, \leq^i)$, we can label the elements of P as $p_0, \ldots, p_{|P|-1}$ such that for all $i \in [z]$ and $j, j' \in \{0, \ldots, |P|-1\}$ such that $j < j'$, either $p_j \not\geq^i p_{j'}$ or p_j and $p_{j'}$ are in the same equivalence class of $(P, \leq^i)$.*

Proof. For each $i \in \{0, \ldots, z\}$, we define $(P, \leq^{*i})$ to be the relation such that $p \leq^{*i} p'$ iff $p \leq^i p'$ and $p' \not\leq^i p$; it is straightforward to see that $\leq^{*i}$ upholds the transitive and asymmetric property necessary to be a partial order. We further define $(P, \leq)$ to be the relation such that $p \leq p'$ iff $p \leq^{*i} p'$ for some $i \in \{0, \ldots, z\}$. This also upholds the transitive property (since if $j \leq i$, then $p \leq^{*i} p' \leq^{*j} p'' \Rightarrow p \leq^{*j} p''$, and $p \leq^{*j} p' \leq^{*i} p'' \Rightarrow p \leq^{*j} p''$), so it is a partial order as well; hence we may extend $(P, \leq)$ to a total ordering $(P, \leq')$. Let $[p_1, \ldots, p_{|P|}]$ be the elements of P ordered in terms of

[1] Since a distributive lattice $\mathcal{L}$ is also a pointed order, we can use the same notation for the least and greatest element of $\mathcal{L}$.

$\leq'$. By the definition of $(P, \leq)$, we note that for all $i \in [z]$ and $j, j' \in \{0, \ldots, |P|-1\}$ such that $j < j'$, either $p_j \not\geq^i p_{j'}$ or p_j and $p_{j'}$ are in the same equivalence class of $(P, \leq^i)$. $\square$

We refer to any total ordering of P as given by Proposition 5.2 as a **reference ordering** of P. (Note that for any reference ordering of P, if $(P, \leq^0)$ is an order, then $\hat{0}_{(P,\leq^0)} = p_0$ and $\hat{1}_{(P,\leq^0)} = p_{|P|-1}$.)

Given any pointed quasi-order $(P, \leq^*)$, we define $\mathcal{D}(P, \leq^*)$ as the collection of downsets of P that contain $\hat{0}$ and not $\hat{1}$. We can restate the Birkhoff Representation Theorem (Theorem 2.10) as follows:

Theorem 5.3. *Given a distributive lattice $\mathcal{L}$, there exists a pointed order $(P, \leq)$ such that $\mathcal{D}(P, \leq)$ is isomorphic to $\mathcal{L}$.*

In this case, we identify an isomorphism of $\mathcal{L}$ with $\mathcal{D}(P, \leq)$, the collection of downsets in the pointed order $(P, \leq)$. In particular, we note that $(P - \{\hat{0}, \hat{1}\}, \leq)$ is isomorphic to the poset of join-irreducible elements of $\mathcal{L}$. Mark Siggers showed that there is a correspondence between the distributive sublattices of $\mathcal{L}$ and the extensions of $(P, \leq)$:

Theorem 5.4. *Given a distributive lattice $\mathcal{L}_1$, let $(P, \leq)$ be a pointed quasi-order such that $\mathcal{L}_h = \mathcal{D}(P, \leq)$. Then, there exists a bijection Γ from the set of all distributive sublattices $\mathcal{L}_0$ of $\mathcal{L}_1$ to the extensions $(P, \leq^*)$ of $(P, \leq)$ such that $\Gamma(\mathcal{L}_1)) = (P, \leq)$, and the lattice flag $(\mathcal{D}(\Gamma(\mathcal{L}_0)), \mathcal{D}(\Gamma(\mathcal{L}_1)))$ is isomorphic to $(\mathcal{L}_0, \mathcal{L}_1)$. ([Sig14] Corollary 4.2)*

Corollary 5.5. *Given a lattice z-flag $(\mathcal{L}_0, \ldots, \mathcal{L}_z)$, there exists a pointed order $(P, \leq^z)$ and a sequence of extensions $(P, \leq^{z-1}), \ldots, (P, \leq^0)$ with the property that $(P, \leq^{i-1})$ is a extension of $(P, \leq^i)$ for all $i \in [z]$, such that $(\mathcal{D}(P, \leq^0), \ldots, \mathcal{D}(P, \leq^z))$ is isomorphic to $(\mathcal{L}_0, \ldots, \mathcal{L}_z)$.*

The following theorem of Vladimir Retakh and Michael Saks ([RS]), which extends Theorem 5.4 and Corollary 5.5 from [Sig14], allows us to make a similar statement on cover-proeserving sublattices of $\mathcal{L}$ - thereby having important implications on covering lattice flags. We define a pointed quasi-order to be **separated** if every equivalence class other than the equivalence classes containing $\hat{0}$ and $\hat{1}$ contains exactly one element.

Theorem 5.6. *Given a distributive lattice $\mathcal{L}_1$, let $(P, \leq)$ be a separated quasi-order such that $\mathcal{L}_h = \mathcal{D}(P, \leq)$, and Γ be defined as in Theorem 5.4. Then, Γ maps the set of all cover-preserving sublattices of $\mathcal{L}_1$ to the set of all separated extensions $(P, \leq^*)$ of $(P, \leq)$. ([RS], Theorem 4.2)*

Corollary 5.7. *Given a covering lattice z-flag $(\mathcal{L}_0, \ldots, \mathcal{L}_z)$, there exists a pointed order $(P, \leq^z)$ and a sequence of extensions $(P, \leq^{z-1}), \ldots, (P, \leq^0)$ with the property that $(P, \leq^{i-1})$ is a separated extension of $(P, \leq^i)$ for all $i \in [z]$, such that $(\mathcal{D}(P, \leq^0), \ldots, \mathcal{D}(P, \leq^z)$ is isomorphic to $(\mathcal{L}_0, \ldots, \mathcal{L}_z)$.*

Proof. By Theorem 5.3, there exists a pointed order $(P, \leq^z)$ such that $\mathcal{D}(P, \leq^z)$ is isomorphic to $\mathcal{L}_z$; let $\gamma : \mathcal{D}(P, \leq^z) \to \mathcal{L}_z$ be the order-preserving bijection.

Now, we will show that for all $i \in [z]$, there exists a separated extension $(P, \leq^{z-i})$ of $(P, \leq^z)$ such that $(P, \leq^{z-i})$ is a separated extension of $(P, \leq^{z-i+1})$, and γ maps $\mathcal{D}(P, \leq^{z-i})$ to $\mathcal{L}_{z-i}$; we do this by induction on i. For our base case, when $i = 1$, such an extension exists by Theorem 5.6.

For our inductive step, for any given $i \in [z]$, assume that we have a separated extension $(P, \leq^{z-i+1})$ of $(P, \leq^z)$ such that γ maps $\mathcal{D}(P, \leq^{z-i+1})$ to $\mathcal{L}_{z-i+1}$. Then, by Theorem 5.6, there exists a separated extension $(P, \leq^{z-i})$ of $(P, \leq^{z-i+1})$ such that γ maps $\mathcal{D}(P, \leq^{z-i})$ to $\mathcal{L}_{z-i}$. Since every equivalence class of $(P, \leq^{i-1})$ other than those containing $\hat{0}$ and $\hat{1}$ has one element, $(P, \leq^{i-1})$ is also a separated extension of $(P, \leq^z)$. Thus, we have completed the inductive step, and by induction, we see that γ maps the lattice z-flag $(\mathcal{D}(P, \leq^0), \ldots, \mathcal{D}(P, \leq^z)$ to $(\mathcal{L}_0, \ldots, \mathcal{L}_z)$, and $(P, \leq^{i-1})$ is a separated extension of $(P, \leq^i)$ for all $i \in [z]$. $\square$

Proposition 5.8. *Let $(P, \leq)$ be a pointed order and $(P, \leq^*)$ be a separated extension of $(P, \leq)$. Then, $\hat{0}_{\mathcal{D}(P,\leq^*)}$ is the set of $p \in P$ in the equivalence class of $\hat{0}$ in $\leq^*$, and $\hat{1}_{\mathcal{D}(P,\leq^*)}$ is the set of $p \in P$ not in the equivalence class of $\hat{1}$ in $\leq^*$.*

5.1.1 The Rotations as a Pointed Order

We will apply the above representation theorems - especially Corollary 5.7 - in the context of the lattice of stable (or hub-stable) matchings. If $\mathcal{L}_z$ is isomorphic to the lattice of stable matchings for a given instance I, we recall that by Theorem 2.18, $\mathcal{L}_z$ is isomorphic to the lattice of all downsets of the rotation poset of I. Combining this with Theorem 5.3, we see the following:

Proposition 5.9. *Let I be an instance, and $(P, \leq)$ be a pointed order such that $\mathcal{D}(P, \leq)$ is isomorphic to $\mathcal{L}_s(I)$. Then, there exists a bijection μ from $P - \{0, 1\}$ to $\Pi(I)$ such that $p_1 \leq p_2$ iff $\mu(p_1) \leq \mu(p_2)$ in $\Pi(I)$.*

In particular, we can compose μ with the order-preserving bijection ν^{-1} (with ν described as in Theorem 2.18) from the downsets of $\Pi(I)$ to the stable matchings over I. If θ is the function with domain $\mathcal{D}(P, \leq)$ such that $\theta(D) = D - \{\hat{0}\}$, then we may consider the mapping $\gamma : \mathcal{D}(P, \leq) \to \mathcal{L}_s(I)$ such that $\gamma = \mu \circ \nu^{-1} \circ \theta$.

Proposition 5.10. *Let I be an instance, and $(P, \leq)$ be a pointed order such that $\mathcal{D}(P, \leq)$ is isomorphic to $\mathcal{L}_s(I)$. Then, the mapping $\gamma : \mathcal{D}(P, \leq) \to \mathcal{L}_s(I)$ is an order preserving bijection such that for all $D \in \mathcal{D}(P, \leq)$, $\gamma(D) = M_0 \cup (\cup_{p \in D}(\mu(p))_d) - (\cup_{p \in D}(\mu(p))_m)$. Furthermore, for all $D, D' \in \mathcal{D}(P, \leq)$, $\gamma(D)$ dominates $\gamma(D')$ iff $D \supseteq D'$.*

We note that Corollary 5.7 allows us to represent $\mathcal{L}_h(I)$ and $\mathcal{L}_s(I)$ for a given instance I as respectively representing downsets of the set P under a pointed order $(P, \leq^h)$ and a quasi-order $(P, \leq^s)$ which is a separated extension of $(P, \leq^h)$.

5.2 Background on the Construction of the Representative Instance

In [Bla84], Charles Blair gave an algorithm to construct an instance such that the lattice of stable matchings is isomorphic to a given distributive lattice $\mathcal{L}$ (see immediately after Theorem 2.9). An improvement on this result appears in [GILS87], which provides an algorithm that, for any distributive lattice $\mathcal{L}$ with O as its poset of join-irreducible elements, gives an instance I_0 of relatively small size such that $\mathcal{L}_s(I) = \mathcal{L}$. The algorithms that we use here will use the algorithm in [GILS87] as a foundation, and so we review the algorithm here.

One tool that the construction uses is the *Hasse diagram* of a poset P. The Hasse diagram of P is the digraph $H(P)$ with vertex set P such that $e = (p_1, p_2)$ is an edge in $H(P)$ iff p_1 covers p_2; in such a case, we say that e is incident with p_1 from below, and incident with p_2 from above. (In pictures of the Hasse diagram, we generally don't show directed edges as having an arrow, and instead position the vertices such that if $p_1 \geq p_2$, then p_1 appears higher in the picture than p_2.)

Algorithm 5.11. *Let $(P, \leq)$ be a pointed order, and Q be a list of elements of $P - \{\hat{1}\}$ (potentially with repeated elements);*[2] *we construct a set of mentors V_m and a set of students V_d with preference lists of the opposite type as follows:*

1. *Let $k = |P| - 2$, and $P = \{p_0, \ldots, p_{k+1}\}$ be any reference ordering of P, as given by Proposition 5.2. (Note that $\hat{0}_P = p_0$ and $\hat{1} = p_{k+1}$.)*

2. *Let $H(P)$ be the Hasse diagram of P. Let $E' = [(\hat{1}, p) : p \in Q]$ (potentially with repeated edges), and E be the disjoint union of $E(H(P))$ and E'. The instance I_0 will have $V_m = \{m_e : e \in E\}$ and $V_d = \{d_e : e \in E\}$.*

3. *In this step and the next one, we construct preference lists for each mentor m_e and each student d_e for $e \in E$. For each $e \in E$, initialize the list of m_e by placing d_e on their preference list, and initialize the list of d_e by placing m_e on their preference list.*

4. *For i from 1 to k, iterate the following: Let $A_i = \{a_i(1), \ldots, a_i(r_i)\}$ be an arbitrary ordering of the edges in E incident with $p_i \in P$. Let $B_i = \{b_i(1), \ldots, b_i(r_i)\}$ such that for all $j \in [r_i]$, $d_{b_i(j)}$ be the last choice on $m_{a_i(j)}$'s current preference list. Then, for all $j \in [r_i]$, place $d_{b_i(j+1)}$ at the bottom of $m_{a_i(j)}$'s preference list and $m_{a_i(j)}$ at the top of $d_{b_i(j+1)}$'s preference list, where $j+1$ is taken $\mod r_i$.*

Theorem 5.12. *Let $\mathcal{L}$ be a distributive lattice, and $(P, \leq)$ be a pointed order such that $\mathcal{D}(P, \leq)$ is isomorphic to $\mathcal{L}$. Then, the set of preference lists I_0 constructed from $(P, \leq)$ by Algorithm 5.11 is a stable matching instance, and $\mathcal{L}$ is isomorphic to $\mathcal{L}_s(I_0)$. [GILS87]*

---

[2] *This construction is a generalization of the one given by [GILS87]; in the original construction, $Q = []$.*

In later sections we will adapt the algorithm and theorem to other contexts involving lattice flags. It is therefore useful to review the details of the proof.

For all $i \in [k]$, we define $\rho(i) = (\rho_m(i), \rho_d(i))$, where $\rho_m(i) = \{(m_{a_i(j)}, d_{b_i(j)}) : j \in [r_i]\}$ and $\rho_d(i) = \{(m_{a_i(j)}, d_{b_i(j+1)}) : j \in [r_i - 1]\} \cup \{m_{a_i(r_i)}, d_{b_i(1)})\}$. (We will show in Theorem 5.23 that $\rho(i)$ is a rotation over I_0, as defined in Section 2.3.) Furthermore, for $i \in \{0, \ldots, k\}$, we define M_i be the set of all edges (m, d) such that d appears last on m's preference list after the ith iteration of step 4. (For M_0, this is the set of edges such that d appears last on m's preference list after step 3.)

Proposition 5.13. *For all $i \in \{0, \ldots, k\}$, M_i is a perfect matching, and for all $d \in V_d(I_0)$, $p_{M_i}(d)$ appears first on d's preference list after the ith iteration of step 4.*

Proof. We prove this result by induction on i. For the base case, when $i = 0$, the statement is trivial. For the inductive step, assume for $i \geq 0$ that M_i is a perfect matching such that, for all $d \in V_d(I_0)$, $p_{M_i}(d)$ appears first on d's preference list after the ith iteration of step 4. Then, since the $(i+1)$th iteration of step 4 adds exactly one student to the bottom of the preference lists of each $m \in \{m_a : a \in A_{i+1}\}$, we see that $M_{i+1} = M_i \cup \rho_d(i+1) - \rho_m(i+1)$. We note that for all $b \in B_{i+1}$, M_{i+1} matches d_b with a different element of $\{m_a : a \in A_{i+1}\}$, and that element was added to the top of d_b's preference list in the $(i+1)$th iteration of step 4. For all $d \in V_d(I_0) - \{d_b : b \in B_{i+1}\}$, M_{i+1} matches d to the same element of $V_d(I_0) - \{m_a : a \in A_{i+1}\}$ as M_i - all of which, by the inductive assumption, are distinct and appear at the top of the corresponding d's preference after the ith iteration of step 4. The $(i+1)$th iteration does not change this, so M_{i+1} is a perfect matching such that, for all $d \in V_d(I_0)$, $p_{M_{i+1}}(d)$ appears first on d's preference list after the $(i+1)$th iteration of step 4. $\square$

It is not immediately obvious that the preference ists constructed in Algorithm 5.11 produce a stable matching instance. In order for this to be the case, we need each vertex's preference list to consist of distinct elements.

Proposition 5.14. *Given any pointed order $(P, \leq)$, let V_m and V_d (and their corresponding preference lists) be defined as in Algorithm 5.11. Then, for all $m \in V_m, d \in V_d$, m and d appear in one another's preference lists at most once.*

Proof. By symmetry, it is sufficient to show that no mentor appears on the preference list of any student more than once. Let d_e be an arbitrary element of V_d, and $m_{e_1} = m_e, m_{e_2}, \ldots, m_{e_c}$ be the mentors in d_e's preference list, in the order that they are added to d_e's preference list; for $j \geq 2$, let i_j be the iteration of step 4 where m_{e_j} is added to d_e's preference list. Since the above algorithm only adds vertices to the top of d_e's preference list, d_e's preference list is $[m_{e_c}, m_{e_{c-1}}, \ldots, m_{e_1}]$.

By the description of step 4 above and the fact that every $p \in P - \{\hat{0}, \hat{1}\}$ has at least two edges incident with it in $H(P)$ (one above and one below), $e_i \neq e_{i+1}$ for all $i \in [c-1]$. In particular, if $c = 2$, $m_{e_{c-1}} \neq m_{e_c}$.

If $c = 1$, then d_e's preference list trivially cannot include any element more than once; if $c = 2$, then $m_1 \neq m_2$, so all of the elements on the preference list of d_e are distinct. Now, assume $c \geq 3$; for $2 \leq j \leq c$, we define $i_j \in [k]$ such that m_{e_j} is added to d_e's preference list in the i_jth iteration of step 4. We show the following lemma:

Lemma 5.15. *For all $2 \leq j \leq c$, e_j and e_{j-1} are incident with p_{i_j}.*

Proof. Since m_{e_j} and d_e add each other to their respective preference lists in the i_jth iteration of step 4, $e_j \in A_{i_j}$ and $e \in B_{i_j}$. The former fact immediately implies that e_j is incident with p_{i_j}. We also note that, since d_e's preference list is constructed from bottom to top, their top choice prior to the i_jth iteration of step 4 was $m_{e_{j-1}}$ - and at that time, $m_{e_{j-1}}$'s bottom choice was d_e by Proposition 5.13. Since $e \in B_{i_j}$, this tells us that e_{j-1} must be in A_{i_j},and so e_{j-1} is incident with p_{i_j}. □

Corollary 5.16. *For all $2 \leq j \leq c-1$, $e_j \in E(H(P))$ and is incident with p_{i_j} from above and $p_{i_{j+1}}$ from below.*

Proof. By Lemma 5.15, e_j is incident with both p_{i_j} and $p_{i_{j+1}}$; since these vertices are distinct, e_j must be incident with one from above and the other from below, with the former covered by the latter in P. However, since m_{e_j} is added to d_e's preference list before $m_{e_{j+1}}$, $i_j \leq i_{j+1}$, and so $p_{i_j} \ngeq p_{i_{j+1}}$. As a result, p_{i_j} cannot cover $p_{i_{j+1}}$, implying that e_j is incident with p_{i_j} from above and $p_{i_{j+1}}$ from below. Since $i_{j+1} \in [k]$ for all $j \in \{2, \ldots, c-1\}$, no such e_j is incident with $\hat{1}$ from below, and so every such

$e_j \in E(H(P))$. $\square$

By Corollary 5.16, $\{p_{i_j} : 2 \leq j \leq c\}$ forms a maximal chain in $(P, \leq)$, and so every element of $\{e_j : 2 \leq j \leq c-1\}$ is distinct. We only need to show that e_1 and e_c are also distinct from these elements and one another.

Lemma 5.17. *For all $j \geq 3$, e_1 is not incident with p_{i_j}.*

Proof. By Lemma 5.15, e_1 is incident with p_{i_2}. For all $j \geq 4$, $p_{i_2} \leq p_{i_3} \leq p_{i_j}$, so $(p_{i_j}, p_{i_2}) \notin H(P)$; in addition, since $i_j \in [k]$, $p_{i_j} \neq \hat{1}$, so $(p_{i_j}, p_{i_2}) \notin E'$. As a result, $(p_{i_j}, p_{i_2}) \notin E$, and so e_1 cannot be incident with p_{i_j}. We now need only to show that e_1 is not incident with p_{i_3}.

Assume for the sake of contradiction that e_1 is incident with p_{i_3}; then, e_1 and e_2 are both incident with p_{i_2} and p_{i_3}. However, neither vertex equals $\hat{1}$, so $e_1, e_2 \notin E'$, and thus $e_1, e_2 \in H(P)$. Since they are incident with the same two vertices, and $H(P)$ has no repeated edges, $e_1 = e_2$; however, this contradicts the fact that $e_{j+1} \neq e_j$ for all $j \in [c-1]$, so e_1 is not incident with p_{i_3} and we are done. $\square$

By Lemma 5.15, $e_{j'}$ is incident with some element of $\{p_{i_j} : 3 \leq j \leq c\}$ for all $j' \geq 2$, so e_1 is distinct from every element of $\{e_2, \ldots, e_c\}$.

Lemma 5.18. *For all $j \leq c-1$, e_c is not incident with p_{i_j}.*

Proof. By Lemma 5.15, e_c is incident with p_{i_c}. For all $j \leq c-2$, $p_{i_j} \leq p_{i_{c-1}} \leq p_{i_c}$, so $(p_{i_c}, p_{i_j}) \notin H(P)$; in addition, since $i_c \in [k]$, $p_{i_c} \neq \hat{1}$, so $(p_{i_c}, p_{i_j}) \notin E'$. As a result, $(p_{i_c}, p_{i_j}) \notin E$, and so e_c cannot be incident with p_{i_j}. We now need only to show that e_c is not incident with $p_{i_{c-1}}$.

Assume for the sake of contradiction that e_c is incident with $p_{i_{c-1}}$; then, e_{c-1} and e_c are both incident with $p_{i_{c-1}}$ and p_{i_c}. However, neither vertex equals $\hat{1}$, so $e_{c-1}, e_c \notin E'$, and thus $e_{c-1}, e_c \in H(P)$. Since they are incident with the same two vertices, and $H(P)$ has no repeated edges, $e_{c-1} = e_c$; however, this contradicts the fact that $e_{j+1} \neq e_j$ for all $j \in [c-1]$, so e_c is not incident with $p_{i_{c-1}}$ and we are done. $\square$

By Lemma 5.15, $e_{j'}$ is incident with some element of $\{p_{i_j} : 2 \leq j \leq c-1\}$ for all $j' \leq c-1$, so e_c is distinct from every element of $\{e_1, \ldots, e_{c-1}\}$.All together, these imply that the elements $\{e_1, \ldots, e_c\}$ are distinct. $\square$

Thus, Algorithm 5.11 produces a stable matching instance. Knowing this, we can make a few observations on the structure of I_0, and in particular on the overlap between elements of $\{\rho(i) : i \in [k]\}$.

Proposition 5.19. *For any $m_e \in V_m(I_0)$, $\{i \in [k] : m_e \in \rho(i)\}$ has at most two elements.*

Proof. Consider any $i \in [k]$. Since $\cup_{\epsilon \in \rho_m(i)}\{m_\epsilon, d_\epsilon\} = \cup_{\epsilon \in \rho_d(i)}\{m_\epsilon, d_\epsilon\}$, $m_e \in \rho(i)$ iff m_e appears in some element of $\rho_m(i)$ - which occurs iff $e = a_i(j)$ for some $j \in [r_i]$. By the definition of $\{a_i(1), \ldots, a_i(r_i)\}$, this occurs iff the edge $e \in H(P)$ is incident with the vertex $p_i \in P - \{0, 1\}$. However, e is incident with at most two vertices in $P - \{0, 1\}$, so $\{i \in [k] : m_e \in \rho(i)\}$ has at most two elements. $\square$

Lemma 5.20. *For all $i_1, i_2 \in [k]$, $\rho_d(i_1) \cap \rho_m(i_2) \neq \emptyset$ iff p_{i_2} covers p_{i_1} in $(P, \leq)$.*

Proof. Because step 4 only adds students to the bottom of mentors's preference lists, $\rho_d(i_1) \cap \rho_m(i_2) \neq \emptyset$ iff there exists some mentor m_e such that m_e adds to their preference list in the i_1th and i_2th iterations of step 4, but not any iteration between them. By Proposition 5.19, this occurs iff $i_1 < i_2$ and m_e adds to their preference list in the i_1th and i_2th iterations of step 4.

For any $i \in [k]$, m_e adds to their preference list in the ith iteration of step 4 iff e is incident with p_i in $H(P)$; as a result, $\rho_d(i_1) \cap \rho_m(i_2) \neq \emptyset$ iff $i_1 < i_2$ and there exists some $e \in E(H(P))$ that is incident with p_{i_1} and p_{i_2}. (Any such edge must have e incident with p_{i_1} from above and p_{i_2} from below, since $i_1 < i_2$.) By the definition of a Hasse diagram, this occurs iff p_{i_2} covers p_{i_1} in $(P, \leq)$. $\square$

To complete the proof of Theorem 5.12, we need to show that the poset of join-irreducibles of $\mathcal{L}_s(I_0)$ is isomorphic to $(P, \leq)$. The strategy is to use Theorem 2.18, which says that the poset of join irreducibles of $\mathcal{L}_s(I_0)$ is isomorphic to the rotation poset $\Pi(I_0) = (R(I_0), \leq^R)$. Therefore, Theorem 5.12 follows if we can show that $\Pi(I_0)$ is isomorphic to $(P, \leq)$, and this is how we proceed.

Proposition 5.21. *For all $i \in \{0, \ldots, k\}$, M_i is stable over I_0.*

Proof. Let $(m, d) \in E(G(I))$ be arbitrary; we only need to show that (m, d) does not destabilize M_i. If m and d add each other to their preference lists at or before the ith iteration of step 4, then, by Proposition 5.13, d prefers $p_{M_i}(d)$ to m. On the other hand, if m and d add each other to their preference lists after the ith iteration of step 4, then, because Algorithm 5.11 only adds students to the bottom of mentors's preference lists, we note that m prefers $p_{M_i}(m)$ to d. Either way, (m, d) does not destabilize M_i, and so M_i is stable over I_0. $\square$

Corollary 5.22. *For all $i \in [k]$, M_i covers M_{i-1} in $\mathcal{L}_s(I_0)$.*

Proof. By Proposition 5.21, M_{i-1} is a stable matching over I_0. In the faithful truncation $I_{(M_{i-1}, \emptyset)}$, each mentor in $\rho(i)$ has their partner in $\rho_m(i)$ as their top choice and their partner in $\rho_d(i)$ as their second choice. By Proposition 2.12, $\rho(i)$ is a rotation exposed by M_{i-1} and M_i covers M_{i-1} in $\mathcal{L}_s(I_0)$. $\square$

Theorem 5.23. *Let $(P, \leq)$ be a pointed order, and I_0 be the stable matching instance constructed in Algorithm 5.11 such that $\mathcal{D}(P, \leq)$ is isomorphic to $\mathcal{L}_s(I_0)$. Then, the set $R(I_0)$ of rotations over I_0 is $\{\rho(i) : i \in [k]\}$, and the bijection $\mu : P - \{\hat{0}, \hat{1}\} \to R(I_0)$ such that $\mu(p_i) = \rho(i)$ is an order isomorphism between $(P - \{\hat{0}, \hat{1}\}, \leq)$ and $\Pi(I_0)$.*

Proof. We note that $M_0 = \{(m_e, d_e) : e \in E(H(P))\}$ is the mentor-optimal stable matching over I_0, since it matches each mentor with their top choice; similarly, M_c is the student-optimal stable matching over I_0. In addition, by Corollary 5.22, for all $i \in [k]$, M_i covers M_{i-1} in $\mathcal{L}_s(I_0)$. By Lemma 2.14, $R(I_0) = \{\rho(i) : i \in [k]\}$.

By Theorem 2.17, the rotation poset $\Pi(I_0) = (\{\rho(i) : i \in [k]\}, \leq^r)$, where $\leq^r$ is the transitive closure of $\mathcal{R}(I_0)$ - the digraph containing all edges of the form (ρ, ρ') such that at least one of the following occurs:

- $\rho_d \cap \rho'_m \neq \emptyset$.

- There exists a mentor $m_0 \in \rho'$ and a student $d_0 \in \rho$ such that (m_0, d_0) does not appear in any stable matching over I_0 and, in I_0, m_0 prefers $p_{\rho'_m}(m_0)$ to d_0 to $p_{\rho'_d}(m_0)$ and d_0 prefers $p_{\rho_d}(d_0)$ to m_0 to $p_{\rho_m}(d_0)$.

Since every edge in $G(I_0)$ appears in some stable matching over I_0 by Proposition 5.21, $\mathcal{R}(I_0)$ contains no edges of the second type. For edges of the first type, by Lemma 5.20, $\rho_d \cap \rho'_m \neq \emptyset$ iff there exist $i_{1,2} \in [k]$ such that $\rho = \rho(i_1)$, $\rho' = \rho(i_2)$, and p_{i_2} covers p_{i_1} in $(P, \leq)$. Therefore, we see that $(P - \{\hat{0}, \hat{1}\}, \leq)$ is isomorphic to the transitive closure $\Pi(I_0)$ via the bijection μ. ☐

We take particular note of the fact that any vertex v is added to the preference list of any over vertex v' at most once. Furthermore, we note the following property of I_0 as defined above, which will be very useful in a later section.

Proposition 5.24. *Every hub-stable matching in I_0 is stable.*

Proof. As a result of Theorem 5.23, every edge in $G(I_0)$ appears in some stable matching; consequentially, $\psi_{I_0}(G(I_0)) = G(I_0)$, so $G(I_0)$ is the unique hub over I_0 and the hub-stable matchings are all $G(I_0)$-stable. However, every $G(I_0)$-stable matching is stable by definition, so every hub-stable matching in I_0 is also stable. ☐

5.3 The Structure of the Edge-Specific Sublattice

Before looking at the possible representations of $(\mathcal{L}_s(I), \mathcal{L}_h(I))$ as a lattice flag, we look at the style of reasoning that we will use to determine all such representations on a similar but simpler problem. If two stable matchings both contain a particular edge (m, d), then their join and meet do as well; consequentially, the set of all stable matchings that contain (m, d) is closed under join and meet, and the following results trivially. For this, we let $\mathcal{K}_e = \mathcal{K}_e(I)$ be the set of all stable matchings over I that contain the edge e.

Theorem 5.25. *For a given instance I and edge (m, d), the structure $(\mathcal{K}_e, \preceq)$ is a distributive sublattice of $(\mathcal{L}_s, \preceq)$.*

Proof. In order for $(\mathcal{K}_e, \preceq)$ to be a distributive lattice, it must be closed under $\vee$ and $\wedge$; by the definitions provided in Theorem 2.8, it is straightforward to see that this is the case. $\square$

In this section, we show that $(\mathcal{K}_e(I), \mathcal{L}_s(I))$ is a covering lattice flag, and identify the necessary and sufficient conditions on a lattice flag $(\mathcal{L}_0, \mathcal{L}_1)$ for it to be $(\mathcal{K}_e(I), \mathcal{L}_s(I))$ for some instance I and edge e. (If $\mathcal{L}_0 = \emptyset$, then $\mathcal{L}_1$ can be any distributive lattice - by Theorem 5.12, there exists an instace I such that $\mathcal{L}_s(I)$ is isomorphic to $\mathcal{L}_1$, and setting e to be any edge $\notin E(G(I))$ will make $(\mathcal{K}_e(I), \mathcal{L}_s(I))$ isomorphic to $(\mathcal{L}_0, \mathcal{L}_1)$. For the remainder of the section, we will assume that $\mathcal{L}_0 \neq \emptyset$).)

For this section, we label the mentor-optimal and student-optimal matchings in $\mathcal{K}_e(I)$ as M_0 and M_1 respectively. The interval $[M_0, M_1]$ is the sublattice of $\mathcal{L}_s(I)$ that contains every matching M such that $M_0 \preceq M \preceq M_1$.

Proposition 5.26. *$(\mathcal{K}_e(I), \mathcal{L}_s(I))$ is a lattice flag, and $\mathcal{K}_e(I) = [M_0, M_1]$.*

Proof. If M is any stable matching $\succeq M_0$ and $\preceq M_1$, then m ranks $p_M(m)$ between their partners in M_0 and M_1; however, they is matched with d in both of those matchings, so they must be matched with d in M as well. As a result, $\mathcal{K}_e(I)$ contains the set of all matchings that $\succeq M_0$ and $\preceq M_1$. Furthermore, since M_0 and M_1 are the mentor-optimal and student-optimal matchings respectively in $\mathcal{K}_e(I)$, any element that $\not\succeq M_0$ or $\not\preceq M_1$ cannot be in $\mathcal{K}_e(I)$; therefore, $\mathcal{K}_e(I) = [M_0, M_1]$. $\square$

We note that $M_0 \preceq M_1$ obviously. We recall that an element l of a distributive lattice $\mathcal{L}$ is join-irreducible iff it cannot be represented as the join of two elements $l_1, l_2 \prec l$ and $\neq \hat{0}_{\mathcal{L}}$, and is meet-irreducible iff it cannot be represented as the meet of two elements $l_1, l_2 \succ l$ and $\neq \hat{1}_{\mathcal{L}}$. We define $IJ(\mathcal{L})$ to be the union of $\{\hat{0}_{\mathcal{L}}\}$ and the set of all join-irreducible elements of $\mathcal{L}$, and $IM(\mathcal{L})$ to be the union of $\{\hat{1}_{\mathcal{L}}\}$ and the set of all meet-irreducible elements of $\mathcal{L}$.

Proposition 5.27. *As elements of $\mathcal{L}_s$, $M_0 \in IJ(\mathcal{L}_s)$ and $M_1 \in IM(\mathcal{L}_s)$.*

Proof. If M and M' are two stable matchings that do not contain e, then $M \vee M', M \wedge M' \subseteq M \cup M'$ cannot contain e either. As a result, if we express M_0 as the join of two elements that dominate it, at least one must be in $\mathcal{K}_e$; however, by the definition of M_0, the only such matching that dominates M_0 is itself. Consequentially, M_0 must be join-irreducible or $\hat{0}_{\mathcal{L}_s}$.

Similarly, if we express M_1 as the meet of two elements that it dominates, at least one must be in $\mathcal{K}_e$; however, by the definition of M_1, the only such matching that M_1 dominates is itself. Consequentially, M_1 must be meet-irreducible or $\hat{0}_{\mathcal{L}_s}$. □

Proposition 5.28. *Every element of $\mathcal{L}_s$ either $\succeq M_0$ or $\preceq M_1$.*

Proof. Assume for the sake of contradiction that there exists a matching M such that $M \not\succeq M_0$ and $M \not\preceq M_1$. Since every element of $\mathcal{K}_e$ dominates M_1, M is not in $\mathcal{K}_e$, and $p_M(m) \neq d$. Let M' be any element of $\mathcal{K}_e$. Since $M \wedge M' \prec M'$, m prefers $p_{M \wedge M'}(m)$ to $p_{M'}(m) = d$; however, since $M \wedge M' \prec M \not\succeq M_0$, $M \wedge M' \not\succeq M_0$ and so $\notin \mathcal{K}_e$, implying m strictly prefers $p_{M \wedge M'}(m)$ to d. Since $M \wedge M'$ can only match m with d or $p_M(m)$, m strictly prefers $p_M(m)$ to d.

Similarly, since $M \vee M' \succ M'$, m prefers $p_{M'}(m) = d$ to $p_{M \vee M'}(m)$. As a result, since $M \vee M'$ can only match m with d or $p_M(m)$, m prefers d to $p_M(m)$. This creates a contradiction, so no such M can exist, and so every element of $\mathcal{L}_s$ either $\succeq M_0$ or $\preceq M_1$. □

We will prove that the above three propositions give the only restrictions on the structure of $(\mathcal{K}_e(I), \mathcal{L}_s(I))$.

Theorem 5.29. *Let $(\mathcal{L}_0, \mathcal{L}_1)$ be a lattice flag. Then, there exists an instance I and edge $e \in E(G(I))$ such that $(\mathcal{L}_0, \mathcal{L}_1)$ is isomorphic to $(\mathcal{K}_e(I), \mathcal{L}_s(I))$ iff:*

1. $\hat{0}_{\mathcal{L}_0} \in IJ(\mathcal{L}_1)$ *and* $\hat{1}_{\mathcal{L}_0} \in IM(\mathcal{L}_1)$.
2. $\{l \in \mathcal{L}_1 : \hat{0}_{\mathcal{L}_0} \not\preceq l \not\preceq \hat{1}_{\mathcal{L}_1}\} = \emptyset$.
3. $\mathcal{L}_0 = [\hat{0}_{\mathcal{L}_0}, \hat{1}_{\mathcal{L}_0}]$ *in* $\mathcal{L}_1$.

5.3.1 Proof of Theorem 5.29

In this subsection, we show that, given any lattice flag $(\mathcal{L}_0, \mathcal{L}_1)$ that upholds conditions 1-3 in Theorem 5.29, we can find an instance and an edge e such that $(\mathcal{K}_e(I), \mathcal{L}_s(I))$ is isomorphic to $(\mathcal{L}_0, \mathcal{L}_1)$. We use Corollary 5.7 to represent $(\mathcal{L}_0, \mathcal{L}_1)$ as $(\mathcal{D}(P, \leq^*), \mathcal{D}(P, \leq))$ - in doing so, we need to consider how conditions 1-3 translate into this new representation.

Proposition 5.30. *Let $(P, \leq)$ and $(P, \leq^*)$ be a pointed order and separated extension (see definition prior to Theorem 5.6) respectively. Then, $\mathcal{D}(P, \leq^*)$ is an interval of $\mathcal{D}(P, \leq)$ iff for all $p_1, p_2 \in P$ not in the equivalence class of $\hat{0}$ or $\hat{1}$ in $\leq^*$, $p_1 \leq^* p_2 \Rightarrow p_1 \leq p_2$.*

Proof. We note that $\mathcal{D}(P, \leq^*)$ is an interval of $\mathcal{D}(P, \leq)$ iff for all $d \in \mathcal{D}(P, \leq^*)$ such that $\hat{0}_{\mathcal{D}(P,\leq^*)} \preceq d \preceq \hat{1}_{\mathcal{D}(P,\leq^*)}$, $d \in \mathcal{D}(P, \leq^*)$. If $p_1 \leq^* p_2 \Rightarrow p_1 \leq p_2$ for all $p_1, p_2 \in P$ not in the equivalence class of $\hat{0}$ or $\hat{1}$ in $\leq^*$, we see that every such d remains in $\mathcal{D}(P, \leq^*)$, and $\mathcal{D}(P, \leq^*)$ is an interval of $\mathcal{D}(P, \leq)$.

Otherwise, take any such p_1, p_2 such that $p_1 \not\leq p_2$ and $p_1 \leq^* p_2$; the set $D \subseteq P$, consisting of every element of P that is either $\leq p_2$ or in the equivalence class of $\hat{0}$ in $\leq^*$, is in $\mathcal{D}(P, \leq)$, but not $\mathcal{D}(P, \leq^*)$ (since D contains p_2 but not p_1). However, $\hat{0}_{\mathcal{D}(P,\leq^*)} \preceq D \preceq \hat{1}_{\mathcal{D}(P,\leq^*)}$, so by our note at the beginning of the proof, $\mathcal{D}(P, \leq^*)$ is not an interval of $\mathcal{D}(P, \leq)$. □

Proposition 5.31. *Let $(P, \leq)$ and $(P, \leq^*)$ be a pointed order and separated extension respectively. Then, $\hat{0}_{\mathcal{D}(P,\leq^*)} \in IJ(\mathcal{D}(P, \leq))$ iff the equivalence class of $\hat{0}$ in $\leq^*$ is $\{p \in P : p \leq p_\alpha\}$ for some $p_\alpha \in P - \{\hat{1}\}$, and $\hat{1}_{\mathcal{D}(P,\leq^*)} \in IM(\mathcal{D}(P, \leq))$ iff the equivalence class of $\hat{1}$ in $\leq^*$ is $\{p \in P : p \geq p_\beta\}$ for some $p_\beta \in P - \{\hat{0}\}$.*

Proof. The equivalence class of $\hat{0}$ in $(P, \leq^*)$ is $\hat{0}_{\mathcal{D}(P,\leq^*)}$, so the first statement holds iff $IJ(\mathcal{D}(P, \leq)) = \{\{p \in P : p \leq p_\alpha\} : p_\alpha \in P - \{\hat{1}\}\}$. However, as noted by Birkhoff, the set of join-irreducible elements of $\mathcal{D}(P, \leq)$ is $\{\{p \in P : p \leq p_\alpha\} : p_\alpha \in P - \{\hat{0}, \hat{1}\}\}$; in addition, the only other element of $IJ(\mathcal{D}(P, \leq))$ is $\hat{0}_{\mathcal{D}(P,\leq)} = \{\hat{0}\} = \{p \in P : p \leq \hat{0}\}$.

Similarly, the equivalence class of $\hat{1}$ in $(P, \leq^*)$ is $\hat{1}_{\mathcal{D}(P,\leq^*)}$, so the second statement holds iff $IM(\mathcal{D}(P, \leq)) = \{\{p \in P : p \geq p_\alpha\} : p_\alpha \in P - \{\hat{0}\}\}$. However, as noted by Birkhoff, the set of meet-irreducible elements of $\mathcal{D}(P, \leq)$ is $\{\{p \in P : p \geq p_\alpha\} : p_\alpha \in P - \{\hat{0}, \hat{1}\}\}$; in addition, the only other element of $IM(\mathcal{D}(P, \leq))$ is $\hat{1}_{\mathcal{D}(P,\leq)} = \{\hat{1}\} = \{p \in P : p \geq \hat{1}\}$. $\square$

Proposition 5.32. *Let $(P, \leq)$ and $(P, \leq^*)$ respectively be a pointed order and separated extension such that $\hat{0}_{\mathcal{D}(P,\leq^*)} \in IJ(\mathcal{D}(P, \leq))$ and $\hat{1}_{\mathcal{D}(P,\leq^*)} \in IM(\mathcal{D}(P, \leq))$, and p_α, p_β be defined as in Proposition 5.31. Then, $p_\alpha < p_\beta$ iff every element of $\mathcal{D}(P, \leq)$ is $\succeq \hat{0}_{\mathcal{D}(P,\leq^*)}$ or $\preceq \hat{1}_{\mathcal{D}(P,\leq^*)}$.*

Proof. Every element of $\mathcal{D}(P, \leq)$ is $\succeq \hat{0}_{\mathcal{D}(P,\leq^*)}$ or $\preceq \hat{1}_{\mathcal{D}(P,\leq^*)}$ iff every element of $\mathcal{D}(P, \leq)$ that contains p_β (or any $p_j \geq p_\beta$) also contains p_α (and every element $p_i \leq p_\alpha$). This occurs iff $p_\alpha \leq p_\beta$. In addition, $\alpha \neq \beta$ - otherwise $\hat{1} \leq^* p_\beta = p_\alpha \leq^* \hat{0}$, which contradicts $(P, \leq^*)$ being a pointed quasi-order. $\square$

We therefore see that conditions 1-3 of Theorem 5.29 are equivalent to the existence of two elements $p_\alpha, p_\beta \in P$ satisfying:

1'. The equivalence class of $\hat{0}$ in $\leq^*$ is $\{p :\in P : p \leq p_\alpha\}$, and the equivalence class of $\hat{1}$ in $\leq^*$ is $\{p :\in P : p \geq p_\beta\}$.

2'. $p_\alpha < p_\beta$.

3'. For all $p_1, p_2 \in P$ not in the equivalence class of $\hat{0}$ or $\hat{1}$ in $\leq^*$, $p_1 \leq^* p_2 \Rightarrow p_1 \leq p_2$.

To complete the proof of Theorem 5.29, we need to show that there is an instance I and an edge e so that $(\mathcal{K}_e(I), \mathcal{L}_s(I))$ is isomorphic to $(\mathcal{D}(P, \leq^*), \mathcal{D}(P, \leq))$.

In the case that p_β cover p_α, we claim that the instance I_0 obtained by applying Algorithm 5.11 to $(P, \leq)$ achieves this for the edge $(p_\alpha, p_\beta) \in E(H(P))$. By Theorem 5.12, the algorithm in Algorithm 5.11 generates an instance I_0 such that $\mathcal{L}_1$ is isomorphic to $\mathcal{L}_s(I_0)$. Furthermore, taking e_0 to be the edge of the Hasse diagram $H(P)$ that is incident to both p_α and p_β, let $m' = m_{e_0}$, and d' be the student that m' is partnered with in $\rho_d(\alpha)$ (if $p_\alpha \neq \hat{0}$) and $\rho_m(\beta)$ (if $p_\beta \neq \hat{1}$); we observe that

$(\mathcal{K}_{(m',d')}(I_0), \mathcal{L}_s(I_0))$ is isomorphic to $(\mathcal{L}_0, \mathcal{L}_1)$, by noting that the set of elements in $\mathcal{K}_{(m',d')}(I_0)$ is $\{M \in \mathcal{L}_s(I_0) : \rho_\alpha \in \nu^{-1}(M), \rho_\beta \notin \nu^{-1}(M)\} = \{\nu(\mu(D)) : D \in \mathcal{D}(P, \leq^*)\}$. (For this, μ and ν are defined as in Proposition 5.9 and Theorem 2.18 respectively.)

Now, suppose p_β does not cover p_α in $(P, \leq)$. Then, we need to modify the algorithm as follows. (For this, $H(P, \leq)$ **augmented by** s is the Hasse diagram with the edge s added; the meaning of edges being incident with vertices from above or below remains the same.)

1. Perform step 1 of Algorithm 5.11.

2. Let H be the Hasse diagram $H(P, \leq)$ augmented by $s = (p_\beta, p_\alpha)$, and $E = E(H)$. The instance I will have $V_m = \{m_e : e \in E\}$ and $V_d = \{d_e : e \in E\}$. (This is the same as step 2 of Algorithm 5.11, with an extra m_s and d_s.)

3. Perform step 3 of Algorithm 5.11.

4. For i from 1 to k, iterate the following:

   (a) If $i \neq \alpha$ or β, perform step 4 of Algorithm 5.11.

   (b) If $i = \alpha$ or β, let $a_i(1) = s$, and $\{a_i(2), \ldots, a_i(r)\}$ be an arbitrary ordering of the edges incident with p_i such that $a_i(2)$ is incident with node i from above and $a_i(r)$ is incident with p_i from below.[3] For $j \in [r]$, let $d_{b_i(j)}$ be the last choice on $m_{a_i(j)}$'s current preference list. Then, for $j \in [r]$, place $d_{b_i(j+1)}$ at the bottom of $m_{a_i(j)}$'s preference list and $m_{a_i(j)}$ at the top of $d_{b_i(j+1)}$'s preference list, where $j+1$ is taken $\mod r$.

It is not immediately obvious that these preference lists produce a stable matching instance; in order for this to be the case, we need to show that no vertex appears on the preference list of another vertex more than once.

Lemma 5.33. *No vertex appears on the preference list of another vertex more than once.*

[3] We know that at least one edge is incident with p_i from above, because $p_i \leq \hat{0}$ - implying that $\{p \in P : p_i < p\}$ is nonempty, and so the element of this set with the smallest index covers p_i. Similarly, we know that at least one edge is incident with p_i from below.

Proof. By symmetry, it is sufficient to show that no mentor appears on the preference list of any student more than once. Let d_e be an arbitrary element of V_d, and $m_{e_1} = m_e, m_{e_2}, \dots, m_{e_c}$ be the vertices in d_e's preference list, in the order that they are added to d_e's preference list. (Since the above algorithm only adds vertices to the top of d_e's preference list, d_e's preference list is $[m_{e_c}, m_{e_{c-1}}, \dots, m_{e_1}]$.) If $e \neq s$, then by the same proof as the one presented in Proposition 5.14, no mentor appears on d_e's preference more than once.

Now, suppose that $e = s$. We note that $c > 1$ iff $(m_s, d_s) \in \rho_m(i)$ for some $i \in [k]$; this occurs iff $p_\alpha \neq \hat{0}$ or $p_\beta \neq \hat{1}$. Every element of d_s's preference list is obviously distinct if $c = 1$, so assume that $c \geq 2$; we set $\gamma = \beta$ if $p_\alpha = \hat{0}$ and $= \alpha$ otherwise. Then, because $e_1 = s = a_\gamma(1)$, $e_2 = a_\gamma(r)$ is incident to p_γ from below; the only other node that e_2 is incident to has a lower index than γ, so no subsequent operation of step 4 will add another element to d_s's preference list, and c must equal 2. Now, $e_2 \neq s$, so every element of d_s's preference list is distinct - whether $c = 1$ or ≥ 2. $\square$

Now that we know that the above algorithm produces an instance, we consider the structure of $\mathcal{L}_s(I)$. We do this by constructing the rotation poset and showing that it is isomorphic to $(P - \{\hat{0}, \hat{1}\}, \leq)$ via the bijection μ, analogously to Theorem 5.23. For $i \in [k]$, let

$$\rho(i) = (\{(m_{a_i(1)}, d_{b_i(1)}), \dots, (m_{a_i(r)}, d_{b_i(r)})\}, \{(m_{a_i(1)}, d_{b_i(2)}),$$

$$\dots, (m_{a_i(r-1)}, d_{b_i(r)}), (m_{a_i(r)}, d_{b_i(1)})\}).$$

Lemma 5.34. *$R(I) = \{\rho(i) : i \in [k]\}$, and the bijection $\mu : P - \{\hat{0}, \hat{1}\} \to R(I_0)$ such that $\mu(p_i) = \rho(i)$ is an order isomorphism between $(P - \{0, 1\}, \leq)$ and $\Pi(I_0)$.*

Proof. We note that $M_0 = \{(m_e, d_e) : e \in E(H)\}$ is the mentor-optimal stable matching over I, since it matches each mentor with their top choice. Given this, we may show that the set of all rotations over I is $\{\rho(i) : i \in [k]\}$ by the same argument used in Theorem 5.23.

By Theorem 2.17 and the argument presented in Theorem 5.23, $\Pi(I)) = (\{\rho(i) : i \in [k]\}, \leq^r)$, where $\leq^r$ is the transitive closure of the digraph containing all edges of the form (ρ, ρ') such that $\rho_d \cap \rho'_m \neq \emptyset$. Since every mentor appears in at most two rotations, $\rho_d \cap \rho'_m \neq \emptyset$ iff there exists a mentor in ρ and ρ'; this occurs iff $\rho = \rho(i_1)$ and $\rho' = \rho(i_2)$, where p_{i_2} covers p_{i_1} in $(P, \leq)$ or $(i_1, i_2) = (\alpha, \beta)$. Since $p_\alpha \leq p_\beta$, the effect of $(\rho(\alpha), \rho(\beta))$ on the transitive closure is redundant, and we see that $(P - \{\hat{0}, \hat{1}\}, \leq)$ is isomorphic to the transitive closure $\Pi(I)$ via the bijection μ. □

It remains to select an edge e and show that γ maps $\mathcal{D}(P, \leq^*)$ to $\mathcal{K}_e(I)$. We note that μ maps $\mathcal{D}(P, \leq^*) = \{D \in \mathcal{D}(P, \leq) : p_\alpha \in D, p_\beta \notin D\}$ to $\kappa \equiv \{D \in \mathcal{D}(\Pi(I)) : \rho(\alpha) \in D, \rho(\beta) \notin D\}$.

Lemma 5.35. *Let ν be as defined in Theorem 2.18. Then, ν maps κ to $\mathcal{K}_{(m_s, d')}(I)$ for some $d' \in V_d(I)$.*

Proof. If $\alpha = 0$, we set $d' = d_s$; otherwise, we set $d' = p_{\rho_d(\alpha)}(m_s)$. Now, consider any $D \in \mathcal{D}(\Pi(I))$. The edge (m_s, d') appears in the following rotations over I:

- If $\alpha = 0$, then (m_s, d') is in the mentor-optimal stable matching, and appears in no $\rho_d \in R(I)$; otherwise, $(m_s, d') \in \rho_d(\alpha)$.
- If $\beta = k + 1$, then (m_s, d') is in the student-optimal stable matching (since Algorithm 5.11 does not change the preference lists of m_s or d' after the αth iteration of step 4). Otherwise, $(m_s, d') \in \rho_m(\beta)$ (since after the αth iteration of step 4, the βth iteration is the first thime that m_s - and correspondingily d' - sees its preference list altered).

By Theorem 2.18, $(m_s, d') \in \nu(D)$ iff $\rho(\alpha) \in D$ and $\rho(\beta) \notin D$ - i.e. iff $D \in \kappa$. □

As a result, γ maps $\mathcal{D}(P, \leq^*)$ to $\mathcal{K}_e(I)$ when $e = (m_s, d')$ (with d' defined as in Lemma 5.35). Therefore, $(\mathcal{K}_{(m_s,d')}(I), \mathcal{L}_s(I))$ is isomorphic to $(\mathcal{D}(P, \leq^*), \mathcal{D}(P, \leq)) = (\mathcal{L}_0, \mathcal{L}_1)$.

5.4 Proof of Theorem 5.1

In this section, we prove Theorem 5.1, which states that every covering lattice flag $(\mathcal{L}_s, \mathcal{L}_h)$ can be realized as $(\mathcal{L}_s(I), \mathcal{L}_h(I))$ for some stable matching instance I, and construct such an instance I. Our full construction is outlined in the following algorithm - we note that this algorithm follows steps analogous to those in Algorithm 5.11, with step 4 in particular significantly expanded upon. The crux of the construction is to create an instance I such that $I[\psi_I^\infty]$ is equal to the instance that would be created by Algorithm 5.11 on input of $(P, \leq^h)$ (see Proposition 5.38), and add additional edges that change which matchings are stable without affecting the hub of I. For this section, we define P_0 and P_1 to be the equivalence classes of $\hat{0}$ and $\hat{1}$ respectively in $(P, \leq^s)$.

Algorithm 5.36. *Let $(P, \leq^h)$ be a pointed order, and $(P, \leq^s)$ be a separated extension (defined prior to Theorem 5.6). Then, we construct a set of mentors V_m and a set of students V_d such that each vertex has a preference list consisting of vertices of the other type as follows:*

1. *Let $k = |P| - 2$, and $P = \{p_0, \ldots, p_{k+1}\}$ be any reference ordering of P as defined by Proposition 5.2.*

2. *Let $H(P)$ be the Hasse diagram of $(P, \leq^h)$. The instance I will have $V_m = \{m_e : e \in E\} \cup \{m_\tau\}$ and $V_d = \{d_e : e \in E\} \cup \{d_\tau\}$.*

3. *Perform step 3 of Algorithm 5.11. In addition, initialize the list of m_τ by placing d_τ on their preference list, and initialize the list of d_τ by placing m_τ on their preference list. Set $V'_d = V_d(I) - \{d_\tau\}$.*

4. *For i from 0 to $k+1$, iterate the following steps:*

   (a) *If $0 < i < k+1$, let $a_i(1), \ldots, a_i(r_i)$ be an arbitrary ordering of the edges incident with node i in $H(P)$. For $j \in [r_i]$, let $d_{b_i(j)}$ be the last element of V'_d that appears on $m_{a_i(j)}$'s current preference list. Then, for $j \in [r_i]$, place $d_{b_i(j+1)}$ at the bottom of $m_{a_i(j)}$'s preference list and $m_{a_i(j)}$ at the top of $d_{b_i(j+1)}$'s preference list, where $j+1$ is taken $\mod r_i$. (This is functionally the same as step 4 of Algorithm 5.11 applied to $(P, \leq^h)$, ignoring d_τ - see Lemma 5.47.)*

   (b) *If $p_i \in P - P_0 - P_1$, then we define $y(i) \in E$ to be any edge incident to p_i from above, $x'(i) \in E$ to be any edge incident to p_i from below, and $x(i) \in E$ to* *be the index of the last student on $m_{x'(i)}$'s preference list.*[4] *Then, for every $p_j \in P - P_0 - P_1$ such that $j < i$, $p_j \not\leq^h p_i$, and p_i covers p_j in $(P, \leq^s)$, place $d_{x(j)}$ second from the bottom on $m_{y(i)}$'s preference list, and $m_{y(i)}$ second from the top on $d_{x(j)}$'s preference list. (This ensures that rotations corresponding to elements that are totally ordered in $(P, \leq^s)$ but not $(P, \leq^h)$ are totally ordered in $\Pi(I)$ but not $\Pi(I[\psi_I^\infty])$ - see Lemma 5.48 and Lemma 5.53.)*

(c) *If p_i is the last element of P_0, then, for every $e \in E$, place m_τ second from the top of d_e's preference list and d_e at the top of m_τ's preference list (in any order). (This ensures that rotations corresponding to elements of P that are $\leq^s \hat{0}$ don't appear in $\Pi(I)$ - see Lemma 5.48 and Lemma 5.54.)*

(d) *If p_i is the last element of $P - P_1$, then, for every $e \in E$, place d_τ at the bottom of m_e's preference list and m_e at the top of d_τ's preference list (in any order). (This ensures that rotations corresponding to elements of P that are $\geq^s \hat{1}$ don't appear in $\Pi(I)$ - see Lemma 5.48 and Lemma 5.54.)*

Proposition 5.37. *For all $m \in V_m(I)$ and $d \in V_d(I)$, m and d appear on each other's preference list in I an equal number of times.*

Proof. Each step of Algorithm 5.36 adds m to d's preference list iff it adds d to m's preference list. □

For the instance I output by Algorithm 5.36, let G_h be the set of edges (m, d) such that m and d add each other to their preference lists in step 3 or 4a. In order to show that this construction creates an instance I where $(\mathcal{L}_s(I), \mathcal{L}_h(I))$ is isomorphic to $(\mathcal{D}(P, \leq^s), \mathcal{D}(P, \leq^h))$, we show the following lemmas centered around the restriction $I[G_h]$, in this order. (Recall that ψ_I^∞ is the unique hub of I.)

1. The preference lists restricted to G_h are the same as the preference list created by Algorithm 5.11 on input of $(P, \leq^h)$ (see Proposition 5.38).

[4] *We note that, since $x'(i) = a_i(j)$ for some $j \in [r_i]$, $d_{x(i)}$ was added to $m_{x'(i)}$'s preference list during the ith iteration of step 4a.*

2. The preference lists created by Algorithm 5.36 do not repeat any elements. (This implies that I is an instance - see Theorem 5.44).

3. The lattice of stable matchings over $I[G_h]$ is isomorphic to $\mathcal{D}(P, \leq^h)$ via an order isomorphism γ (see Lemma 5.47).

4. The lattice of the stable matchings over I which are $\subseteq I[G_h]$ is isomorphic to $\mathcal{D}(P, \leq^s)$ via γ (see Lemma 5.48).

5. $\psi_I^\infty = G_h$, and the set of hub-stable matchings over I is the set of stable matchings over $I[G_h]$ (see Theorem 5.55 and Corollary 5.56).

5.4.1 The Structure of $I[G_h]$

In this subsection, we will carry out step 1 in our outline of the proof of Theorem 5.1.

Proposition 5.38. *Given any pointed order $(P, \leq^h)$ and separated extension $(P, \leq^s)$, the preference lists created by Algorithm 5.36, restricted to the elements added during steps 3 and 4a, is the set of preference lists created by Algorithm 5.11 on input of $(P, \leq^h)$, with $E = \{\tau\} = \{(\hat{0}, \hat{1})\}$.*

Proof. The instance created by running Algorithm 5.36 without running steps 4b, 4c, and 4d is the same as the instance created by Algorithm 5.11 on input of $(P, \leq^h)$, with $Q = [\hat{0}]$ - implying $E' = [\tau] = [(\hat{1}, \hat{0})]$ (since m_τ and d_τ are never added to another preference list by step 4a). Thus, to prove the proposition, we need only show that steps 4b, 4c, and 4d never change the element of V'_d that appears last in any mentor's preference list (since then, steps 4b, 4c, and 4d will not have any impact on step 4a).

In the ith iteration of step 4, for each relevant $j < i$, step 4b places a student second from the top of $m_{y(i)}$'s preference list; however, $m_{y(i)}$ already has a preference list with at least 2 terms $\neq d_\tau$ (one from step 3, and one from the ith iteration of step 4a - the latter is the case because $y(i) \in E$ is incident with p_i), so this does not change the element of V'_d that appears last in any mentor's preference list. Step 4c only adds elements to the top of m_τ's preference list, and step 4d can only add d_τ to any mentor's preference list, so we are done. □

The preservation of the structure of Algorithm 5.11 in Algorithm 5.36 immediately implies that Proposition 5.13 is preserved in the latter.

Proposition 5.39. *At any given point in creation of I, the mentor-association that matches m_τ with d_τ and each mentor in V'_m with the last element of V'_d on their preference list is a matching, and this matching matches each student in V'_d with the top element on their preference list.*

In particular, Proposition 5.39 implies that every element of V'_d is at the bottom of the preference list of a unique element of V'_m throughout the process of Algorithm 5.36.

5.4.2 The Proof That Algorithm 5.36 Creates an Instance

In this subsection, we carry out step 2 in our outline of the proof of Theorem 5.1. As with Algorithm 5.11, it is not immediately obvious that the preference lists produced by Algorithm 5.36 describe a stable matching instance - in order for this to be the case, we need every such preference list to consist of distinct elements. By Proposition 5.38, steps 3 and 4a don't produce any repeated elements on any preference list, so we only need to make certain that steps 4b, 4c, and 4d don't produce any repeated elements, or add an element already added by step 3 or 4a.

We recall that for all $i \in [k]$ such that $p_i \notin P_0 \cup P_1$, $y(i)$ is an edge incident with p_i from above, $x'(i)$ is an edge incident with p_i from below, and $d_{x(i)}$ is the last element of V'_d to appear on $m_{x'(i)}$'s preference list (as of the ith iteration of step 4b).

Proposition 5.40. *For all $i \in [k]$ such that $p_i \notin P_0 \cup P_1$, in the ith iteration of step 4a, $m_{x'(i)}$ is added to the top of $d_{x(i)}$'s preference list, and $d_{x(i)}$ is added to the bottom of $m_{x'(i)}$'s preference list.*

Proof. As stated in step 4b, $x'(i) \in E$ is incident with p_i, so in the ith iteration of step 4a, $m_{x'(i)}$ has a student $d_e \in V'_d$ added to the bottom of their preference list, and is added to the top of d_e's preference list. By the definition of $x(i)$ and the fact that this occurs immediately before the assignment of $x'(i)$ and $x(i)$ in the ith iteration of step 4b, $d_e = d_{x(i)}$. $\square$

Proposition 5.41. *For all $i \in [k]$ such that $p_i \in P - P_0 - P_1$, from the ith iteration of step 4a onward, $d_{x(i)}$ has $m_{x'(i)}$ as the first element of their preference list, and $m_{x'(i)}$ has $d_{x(i)}$ as the last element of their preference list in $\{d_e : e \in E\}$.*

Proof. Since $d_{x(i)}$ and $m_{x'(i)}$ were added to one another's preference lists in the ith iteration of step 4a by Proposition 5.40, $d_{x(i)}$ has $m_{x'(i)}$ as the first element of their preference list and $m_{x'(i)}$ has $d_{x(i)}$ as the last element of their preference list immediately after the ith iteration of step 4a. For all $e \in E$, steps 4b, 4c, and 4d cannot add an element to the top of d_e's preference list, or put d_e at the bottom of any mentor's preference list; thus, we see that only the jth iteration of step 4a (for some $j > i$) could introduce an element that breaks the property in the proposition.

Assume for the sake of contradiction that there exists a minimum $j > i$ such that at least one of $m_{x'(i)}$ and $d_{x(i)}$ changes their preference list in the jth iteration of step 4a. Since $d_{x(i)}$ is still the last element of V'_d on $m_{x'(i)}$'s preference list before this (since, as the proof of Proposition 5.38 shows, no other step can change the last element of $m_{x'(i)}$'s preference list), they cannot change their preference list unless $m_{x'(i)}$ changes their as well.

In order for $m_{x'(i)}$ to change their preference list in the jth iteration of step 4a, p_j must be incident with $x'(i)$; however, since $x'(i)$ is incident with p_i from below, the only other vertex $x'(i)$ is incident with must have index $< i$. This creates a contradiction, so neither vertex expands its preference list in step 4a after the ith iteration, and so we are done. □

Proposition 5.42. *For all $i \in [k]$ such that $p_i \in P - P_0 - P_1$, $m_{y(i)}$ and $d_{y(i)}$ do not add any element of $\{d_e : e \in E - \{y(i)\}\}$ or $\{m_e : e \in E - \{y(i)\}\}$ to their preference lists before the ith iteration of step 4a.*

Proof. Since such an e has $e \notin \{y(i), \tau\}$, we note that any such addition can only occur in step 4a or 4b. Let j be the smallest natural number such that $m_{y(i)}$ or $d_{y(i)}$ adds to their preference list in the jth iteration of step 4a. Since step 4a is the only time in step 4 that $m_{y(i)}$ can change the last element of V'_d on their preference list, this implies that the last element of V'_d on $m_{y(i)}$'s preference list is the same prior to the jth iteration of step 4a as it is at the end of step 3 - namely, $d_{y(i)}$. Thus, we see that $y(i)$ must be incident with p_j and p_k for some $k > j$. However, since $y(i)$ is incident with p_i from above, it is incident with p_i and p_k for some $k > i$ - thereby implying that $i = j$.

Furthermore, since the only vertices that $y(i)$ is incident to have index $\geq i$, $y(i) \neq x(j)$ or $y(j)$ for any $j < i$. As a result, we see that the preference lists of $m_{y(i)}$ and $d_{y(i)}$ are not changed by steps 4a and 4b before the ith iteration, and so we are done. □

Proposition 5.43. *As functions from $[k]$ to E, $x(i)$ and $y(i)$ are both injections.*

Proof. Consider any $i, j \in [k]$ such that $i < j$. Then, $y(i)$ and $y(j)$ are incident to p_i and p_j from above respectively; since $p_i \neq p_j$, $y(i) \neq y(j)$.

Now, suppose that $x(i) = x(j)$. This implies that $m_{x'(j)}$ was added to the top of $d_{x(i)}$'s preference list during the jth iteration of step 4a. However, this contradicts Proposition 5.41, so this cannot happen. □

Theorem 5.44. *Given any pointed order $(P, \leq)$, let V_m and V_d (and their corresponding preference lists) be defined as in Algorithm 5.36. Then, for all $m \in V_m, d \in V_d$, m and d appear in one another's preference lists at most once.*

Proof. We begin by showing that the following two lemmas hold.

Lemma 5.45. *For all $m \in V_m(I), d \in V_d(I)$, d appears on m_τ's preference list at most once, and m appears on d_τ's preference list at most once.*

Proof. Initially, m_τ only has d_τ on their preference list. The only time that m_τ adds to their preference list thereafter is during the ith iteration of step 4c, where p_i is the last element of P_0; during that step, every other student is added to m_τ's preference list exactly once. Therefore m_τ's preference list has no repeated elements. By a similar argument (with step 4d replacing step 4c and p_i being the last element of $P - P_1$), d_τ's preference list has no repeated elements. □

Lemma 5.46. *For any $e, e' \in E$, $m_{e'}$ appears on d_e's preference list at most once.*

Proof. By Proposition 5.14, steps 3 and 4a together don't add any $m_{e'}$ to d_e's preference list more than once, so we only need to show that step 4b does not cause any duplicates.

(Steps 4c and 4d cannot add any $m_{e'}$ to d_e's preference list when $e, e' \in E$, so we only need to show that we don't create duplicates with steps 3, 4a, and 4b.)

Consider any edge $(m_{e'}, d_e)$ such that d_e adds $m_{e'}$ to their preference list in step 4b; then, $e = x(j)$ and $e' = y(i)$ for some $i, j \in [k]$ such that $j < i$. We note that d_e's preference list is added to in the jth iteration of step 4a and - by Proposition 5.41 - not in any subsequent one. By Proposition 5.42, $m_{e'}$'s preference list is not changed by any iteration of step 4a before the ith one; therefore, since $i > j$, no iteration of step 4a changes both preference lists, which is necessary in order to add $m_{e'}$ to d_e's preference list. Similary, by Proposition 5.42, $d_{e'}$'s preference list is not changed in the jth iteration of step 4a, so $e \neq e'$ and so $m_{e'}$ is not added to d_e's preference in step 3.

As a result, we see that the theorem holds iff for any given $e, e' \in E$, step 4b adds $m_{e'}$ to d_e's preference list at most once. In total, for each $p_i, j \in P - P_0 - P_1$ such that $i > j$, step 4b adds $m_{y(i)}$ to $d_{x(j)}$'s preference list at most once. Furthermore, by Proposition 5.43, the function that maps (i, j) to $(y(i), x(j))$ is an injection, so we are done. $\square$

By Proposition 5.37, the above two lemmas imply that for all $e, e' \in E$, d_e and d_τ each appear on $m_{e'}$'s preference list at most once, and m_τ appears on d_e's preference list at most once. As a result, we see that for all $m \in V_m, d \in V_d$, m and d appear in one another's preference lists at most once. $\square$

5.4.3 The Structure of $\mathcal{L}_h(I)$ and $\mathcal{L}_s(I)$

In this subsection, we carry out steps 3 through 5 of the proof of Theorem 5.1, and tie all of the results together. For all $i \in [k]$, we define $\rho_m(i) = \{(m_{a_i(t)}, d_{b_i(t)}) : t \in [r_i]\}$ and $\rho_d(i) = \{(m_{a_i(t)}, d_{b_i(t+1)}) : t \in [r_i - 1]\} \cup \{m_{a_i(r)}, d_{b_i(1)})\}$. We will show that for all $i \in [k]$, $\rho(i) = (\rho_m(i), \rho_d(i))$ is a rotation over $I[G_h]$ (as defined prior to Theorem 2.11).

Lemma 5.47. *The lattice $\mathcal{L}_s(I[G_h])$ is isomorphic to $\mathcal{L}_h$. Furthermore:*

1. *$R(I[G_h]) = \{\rho(i) : i \in [k]\}$.*

2. *the bijection $\mu : P - \{\hat{0}, \hat{1}\} \to R(I[G_h])$ such that $\mu(p_i) = \rho(i)$ is an order isomorphism between $(P - \{0, 1\}, \leq^h)$ and $\pi(I[G_h])$.*

3. every edge in G_h appears in some stable matching over $I[G_h]$.

Proof. We note by Proposition 5.38 that $I[G_h]$ is the instance created by Algorithm 5.11 on input of $(P, \leq^h)$, with $E' = \{\tau\} = \{(\hat{1}, \hat{0}\}$; as such, by Theorem 5.12, $\mathcal{L}_s(I[G_h])$ is isomorphic to $\mathcal{D}(P, \leq^h)$, which is isomorphic to $\mathcal{L}_h$. In addition, by Theorem 5.23, the following properties hold:

1. $R(I[G_h]) = \{\rho(i) : i \in [k]\}$.

2. the bijection $\mu : P - \{\hat{0}, \hat{1}\} \to R(I[G_h])$ such that $\mu(p_i) = \rho(i)$ is an order isomorphism between $(P - \{0, 1\}, \leq^h)$ and $\pi(I[G_h])$.

3. every edge in G_h appears in some stable matching over $I[G_h]$.

□

Since $\mathcal{L}_s(I[G_h])$ is isomorphic to $\mathcal{D}(P, \leq^h)$, by Proposition 5.10, we can identify a mapping $\gamma : \mathcal{D}(P, \leq) \to \mathcal{L}_s(I)$ is a bijection such that for all $D \in \mathcal{D}(P, \leq)$, $\gamma(D) = M_0 \cup (\cup_{p \in D}(\mu(p))_d) - (\cup_{p \in D}(\mu(p))_m)$. Furthermore, for all $D, D' \in \mathcal{D}(P, \leq^h)$, $\gamma(D)$ dominates $\gamma(D')$ iff $D \supseteq D'$. (Note that, for all $D \in \mathcal{D}(P, \leq^h)$, $\gamma(D)$ is perfect with respect to both $I[G_h]$ and I.)

Lemma 5.48. *Let $S \in \mathcal{D}(P, \leq^h)$. Then, $\gamma(S)$ is stable in I iff $S \in \mathcal{D}(P, \leq^s)$.*

In order to prove this lemma, we look at each $e \in E(G(I)) - G_h$, and determine the necessary and sufficient conditions for $\gamma(S)$ to be $\{e\}$-stable. Let S_b (resp. S_c and S_δ) be the set of all edges $e \in E(G(I))$ such that m_e and d_e add one another to their preference lists in step 4b (resp. 4c and 4d); since G_h consists of every $e \in E(G(I))$ such that m_e and d_e add one another to their preference lists in step 3 (when the preference lists are first created) or 4a, $S_b \cup S_c \cup S_\delta = E(G(I)) - G_h$. We will determine the necessary and sufficient conditions for $\gamma(S)$ to be S_b-stable in Lemma 5.49, and the necessary and sufficient conditions for $\gamma(S)$ to be S_c-stable and/or S_δ-stable in Lemma 5.51.

For the following lemmas, we recall the notation:

$$\rho(i) = \{\rho_m(i), \rho_d(i)\} = \{\{(m_{a_i(1)}, d_{b_i(1)}), \ldots, (m_{a_i(r_i)}, d_{b_i(r_i)})\},$$

$$\{(m_{a_i(1)}, d_{b_i(2)}), \ldots, (m_{a_i(r_i-1)}, d_{b_i(r)}), (m_{a_i(r_i)}, d_{b_i(1)})\}\}.$$

In addition, for all $i, j \in [k]$ such that $\hat{1} \not\leq^s p_i, p_j \not\leq^s \hat{0}$, $m_{y(i)} \in \rho(i)$, $d_{x(j)} \in \rho(j)$, and $m_{y(i)}$ and $d_{x(j)}$ add one another to their preference lists in the ith iteration of step 4b iff $j < i$, $p_j \not\leq^h p_i$, and p_i covers p_j in $(P, \leq^s)$.

Lemma 5.49. *Let $i, j \in [k]$ such that $p_i, p_j \in P - P_0 - P_1$, $i > j$, $p_i \not\geq^h p_j$, and p_i covers p_j in $(P, \leq^s)$. Then, $m_{y(i)} \in \rho(i)$ and $d_{x(j)} \in \rho(j)$. Furthermore, $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$ to $p_{\rho_d(i)}(m_{y(i)})$, and $d_{x(j)}$ prefers $p_{\rho_d(j)}(d_{x(j)})$ to $m_{y(i)}$ to $p_{\rho_m(j)}(d_{x(j)})$.*

Proof. Since $y(i)$ and $x(j)$ are incident with p_i and p_j respectively, $y(i) = a_i(s)$ and $x(j) = b_j(t)$ for some $s \in r_i$, $t \in r_j$. We note the following:

- $m_{a_i(s)}$ has $d_{b_i(s)} = p_{\rho_m(i)}(m_{y(i)})$ as the last element of V'_d on their preference list before the ith iteration of step 4a (by the definition of $b_i(s)$ given in Algorithm 5.36).

- $m_{a_i(s)}$ adds $d_{b_i(s+1)} = p_{\rho_d(i)}(m_{y(i)})$ (with $s+1$ taken $\mod r_i$) to the bottom of their preference list during the ith iteration of step 4a.

- $m_{a_i(s)}$ adds $d_{x(j)}$ second from the bottom of their preference list in the ith iteration of step 4b.

Consequentially, $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$ to $p_{\rho_d(i)}(m_{y(i)})$.

Similarly, we note that:

- $d_{b_j(t)}$ has $m_{a_j(t)} = p_{\rho_m(j)}(d_{x(j)})$ as the last element on their preference list before the jth iteration of step 4a (by Proposition 5.39.

- $d_{b_j(t)}$ adds $m_{a_j(t-1)} = p_{\rho_d(j)}(d_{x(j)})$ (with $s-1$ taken $\mod r_j$) to the top of their preference list during the jth iteration of step 4a.

- $d_{b_j(t)}$ adds $m_{y(i)}$ second from the top of their preference list in the ith iteration of step 4b.

By Proposition 5.41, the top element of $d_{x(j)}$'s preference list does not change after the ith iteration of step 4a, so $d_{x(j)}$ prefers $p_{\rho_d(j)}(d_{x(j)})$ to $m_{y(i)}$ to $p_{\rho_m(j)}(d_{x(j)})$. $\square$

Corollary 5.50. *Let $S \in \mathcal{D}(P, \leq^h)$. Then, $\gamma(S)$ is S_b-stable iff for all i, j such that $j < i$, $p_j \nleq^h p_i$, and $p_j \leq^s p_i$, $p_i \in S \Rightarrow p_j \in S$.*

Proof. We note that S_b is the set of all edges of the form $e = (m_{y(i)}, d_{x(j)})$, where $p_i, p_j \in P - P_0 - P_1$, $i > j$, $p_i \ngeq^h p_j$, and p_i covers p_j in $(P, \leq^s)$. By Lemma 5.49, $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$ to $p_{\rho_d(i)}(m_{y(i)})$, and $d_{x(j)}$ prefers $p_{\rho_d(j)}(d_{x(j)})$ to $m_{y(i)}$ to $p_{\rho_m(j)}(d_{x(j)})$; therefore, by the definition of γ, d_a prefers $p_{\gamma(S)}(d_a)$ to m_b iff S contains p_i, and m_b prefers $p_{\gamma(S)}(m_b)$ to d_a iff S does not contain p_j. Therefore, $\gamma(S)$ is $\{e\}$-stable iff $p_i \in S \Rightarrow p_j \in S$, and $\gamma(S)$ is S_b-stable iff $p_i \in S \Rightarrow p_j \in S$ for all $i, j \in [k]$ such that $j < i$, $p_j \nleq^h p_i$, and p_i covers p_j in $(P, \leq^s)$.

If $p_i \in S \Rightarrow p_j \in S$ for all i, j such that $j < i$, $p_j \nleq^h p_i$, and $p_j \leq^s p_i$, then $\gamma(S)$ upholds the necessary conditions to be S_b-stable. Conversely, suppose that $\gamma(S)$ is S_b-stable; for any $i, j \in [k]$ such that $j < i$, $p_j \nleq^h p_i$, and $p_j \leq^s p_i$, there exists a sequence $p_j, p', p'', \ldots, p^(l), p_i$ such that each element other than p_j covers the previous one in $(P, \leq^s)$. As a result, $p_i \in S \Rightarrow p^(l) \in S \Rightarrow \ldots \Rightarrow p_j \in S$, so $p_i \in S \Rightarrow p_j \in S$ for all i, j such that $j < i$, $p_j \nleq^h p_i$, and $p_j \leq^s p_i$. $\square$

Lemma 5.51. *For all $i \in \{0, \ldots, k\}$, $\gamma(\{p_0, \ldots, p_i\})$ matches each mentor in V_m' (resp. student in V_d') with their bottom choice in V_d' (resp. their top choice) after the ith iteration of step 4.*

Proof. We prove this result by induction. For the base case, when $i = 0$, $\gamma(\{p_0\}) = \{(m_e, d_e) : e \in E \cup \{\tau\}\}$ matches each mentor in V_m' (resp. student in V_d') with their bottom choice in V_d' (resp. their top choice) after step 3. The only steps that can occur nonvacuously during the 0th iteration of step 4 are 4c and 4d; however, as seen in the proof of Proposition 5.38, neither step can change the bottom (resp. top) element of V_d' (resp. V_m) on the preference list of any mentor in V_m' (resp. student in V_d'), so $\gamma(\{p_0,\})$ matches each mentor in V_m' (resp. student in V_d') with their bottom choice in V_d' (resp. their top choice) after the 0th iteration of step 4.

To show our inductive step, assume that for any $i \in [k]$, $\gamma(\{p_0, \dots, p_{i-1}\})$ matches each mentor in V'_m (resp. student in V'_d) with their bottom choice in V'_d (resp. their top choice) after the $(i-1)$th iteration of step 4. The ith iteration of step 4a makes the following changes to preference lists:

- For all $m \in V'_m$ such that $m \in \rho(i)$, $p_{\rho_d(i)}(m) \in V'_d$ is added to the bottom of m's preference list. (Previously, by the inductive assumption, the last element of V'_d on m's preference list was $p_{\gamma(\{p_0,\dots,p_{i-1}\}}(m) = p_{\rho_m(i)}(m)$.)

- For all $d \in V'_m$ such that $d \in \rho(i)$, $p_{\rho_d(i)}(d)$ is added to the top of d's preference list. (Previously, by the inductive assumption, the first element on m's preference list was $p_{\gamma(\{p_0,\dots,p_{i-1}\}}(d) = p_{\rho_m(i)}(d)$.)

As a result, we see that there exists a perfect matching over I that matches each mentor in V'_m (resp. student in V'_d) with their bottom choice in V'_d (resp. their top choice) after the ith iteration of step 4a, and this matching is $\gamma(\{p_0, \dots, p_{i-1}\}) \cup \rho_d(i) - \rho_m(i) = \{(m_e, d_e) : e \in E \cup \{\tau\}\} \cup (\cup_{j=1}^{i} \rho_d(j)) - (\cup_{j=1}^{i} \rho_m(j)) = \gamma(\{p_0, \dots, p_i\})$. As seen in the proof of Proposition 5.38, steps 4b through 4d do not change the bottom element of V'_d (resp. top element) in the preference list of any $m \in V'_m$ (resp. $d \in V'_d$), so we see that $\gamma(\{p_0, \dots, p_i\})$ matches each mentor in V'_m (resp. student in V'_d) with their bottom choice in V'_d (resp. their top choice) after the ith iteration of step 4. By indiction, we are done. $\square$

Corollary 5.52. *Let $S \in \mathcal{D}(P, \leq^h)$. Then, $\gamma(S)$ is S_c-stable (resp. S_δ-stable) iff $p_i \in S$ for all $p_i \leq^s \hat{0}$ (resp. $p_i \notin S$ for all $p_i \geq^s \hat{1}$).*

Proof. Recall that $S_c = \{(m_\tau, d) : d \neq d_\tau\}$. Since $(m_\tau, d_\tau) \in \gamma(S)$ for all $S \in \mathcal{D}(P, \leq^h)$, and m_τ prefers any other possible partner to d_τ, $\gamma(S)$ is S_c-stable iff every student other than d_τ prefers their partner in $\gamma(S)$ to m_τ. Since m_τ is added to each student's preference list second from the top at the ith iteration of step 4c (where i is the greatest index such that $p_i \leq^s \hat{0}$), this occurs iff each such student weakly prefers their partner in $\gamma(S)$ to their top choice at that point. By Lemma 5.51, that top choice is their partner in $\gamma(T_m)$, where $T_m = \{p \in P : p \leq^s \hat{0}\}$ - and every student prefers their partner in $\gamma(S)$ iff $S \supseteq T_m$.

Similary, recall that $S_\delta = \{(m, d_\tau) : m \neq m_\tau\}$. Since $(m_\tau, d_\tau) \in \gamma(S)$ for all $S \in \mathcal{D}(P, \leq^h)$, and d_τ prefers any other possible partner to m_τ, $\gamma(S)$ is S_δ-stable iff every mentor other than m_τ prefers their partner in $\gamma(S)$ to d_τ. Since d_τ is added to each mentor's preference list at the bottom during the ith iteration of step 4d (where i is the greatest index such that $p_i \not\geq^s \hat{1}$), this occurs iff each such mentor weakly prefers their partner in $\gamma(S)$ to their bottom choice in V_d' at that point. By Lemma 5.51, that bottom choice is their partner in $\gamma(T_d)$, where $T_d = \{p \in P : p \not\geq^s \hat{1}\}$ - and every mentor prefers their partner in $\gamma(S)$ iff $S \subseteq T_d$. □

We can now prove Lemma 5.48.

Proof. We see that S is $\{e\}$-stable for every $e \in G(I)$ iff the following hold:

- $p_i \in S \Rightarrow p_j \in S$ for all $i, j \in [k]$ such that $j < i$, $p_j \not\leq^h p_i$, and $p_j \leq^s p_i$. (This is the necessary and sufficient condition for $\gamma(S)$ to be S_b-stable, by Corollary 5.50.)

- $p_i \in S$ for all $p_i \leq^s \hat{0}$. (This is the necessary and sufficient condition for $\gamma(S)$ to be S_c-stable, by Corollary 5.52.)

- $p_i \notin S$ for all $p_i \geq^s \hat{1}$. (This is the necessary and sufficient condition for $\gamma(S)$ to be S_δ-stable, by Corollary 5.52.)

However, this is the list of necessary and sufficient conditions for any $S \subseteq P$ to be in $\mathcal{D}(P, \leq^s)$, so we are done. □

Since every stable matching in I that consists entirely of edges in G_h is also stable in $I[G_h]$, we see that γ is a bijection from $\mathcal{D}(P, \leq^s)$ to the stable matchings of I that consist entirely of edges in G_h. This set of stable matchings, which we define as $\mathcal{S}$, is obviously closed under join and meet (as the join and meet of two stable matchings consist of edges from those matchings).

We now aim to show that $\psi^\infty(I) = G_h$. Since every edge in G_h appears in some element of $\mathcal{S}$, $\psi^\infty(I) \supseteq \cup_{S \in \mathcal{S}} S$. Therefore, we need only to show that for every edge

$e_0 = (m_e, d_e) \notin \cup_{S \in \mathcal{S}} S$, $e \notin \psi^\infty(I)$.

Lemma 5.53. *If $e \in S_b$, then $e \notin \psi_I^\infty$.*

Proof. Since $e \in S_b$, $e = (m_{y(i)}, d_{x(j)})$, where $i, j \in [k]$ such that $i > j$, $p_i, p_j \in P - P_0^* - P_1^*$, p_i and p_j are independent in $\mathcal{D}(P, \leq^h)$, and p_i covers p_j in $\mathcal{D}(P, \leq^s)$. Since p_i covers p_j in $\mathcal{D}(P, \leq^s)$, there exists some $D \in \mathcal{D}(P, \leq^s)$ such that $p_j \in D$ and $p_i \notin D$. As a result, if $M' = \gamma(D)$, then $d_{x(j)}$ prefers $p_{M'}(d_{x(j)})$ to $p_{\rho_d(i)}(d_{x(j)})$, and $m_{y(i)}$ prefers $p_{M'}(m_{y(i)})$ to $p_{\rho_m(i)}(m_{y(i)})$. However, by Lemma 5.49, $d_{x(j)}$ also prefers $p_{\rho_d(i)}(d_{x(j)})$ to $m_{y(i)}$, and $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$.

As a result, if M is any matching that includes e, then $m_e = m_{y(i)}$ and $d_e = d_{x(j)}$ each prefer their partner in M' to their partner in M, so M and M' are not costable. However, by Lemma 5.48, M' is stable - and therefore hub-stable - so M cannot be hub-stable. Since M is any arbitrary matching that includes e, $e \notin \psi_I^\infty$. □

Lemma 5.54. *If $e \in S_c \cup S_\delta$, then, $e \notin \psi_I^\infty$.*

Proof. Assume, for the sake of contradiction, that the lemma is false; then, there exists a hub-stable matching M^* such that m_τ and d_τ are not matched with each other. Let M_0 be any stable matching over I that includes (m_τ, d_τ) as an edge - we know such a matching exists by Lemma 5.48. Since M^* and M_0 are hub-stable, they must be costable as well. However, m_τ and d_τ are partnered in M_0, and both prefer their respective partners in M^* to each other; this creates a contradiction with M^* and M_0 being costable, and so no such M^* can exist. □

Theorem 5.55. $\psi_I^\infty = G_h$.

Proof. As a consequence of Lemma 5.53 and Lemma 5.54, $\psi_I^\infty \subseteq G_h$, and so $G_h \supseteq \psi_I(G_h)$. However, by Lemma 5.47, every edge $e \in G_h$ appears in some G_h-stable matching M_e over $I[G_h]$; this matching remains G_h-stable over I, so $e \in \psi_I(G_h)$. This means that $G_h \subseteq \psi_I(G_h)$, so $G_h = \psi_I(G_h)$ - implying that $G_h = \psi_I^\infty$. □

Corollary 5.56. *The set of hub-stable matchings over I is the set of stable matchings over $I[G_h]$.*

Since every stable matching is hub-stable, every stable matchings over I appears in $\mathcal{S}$, as defined in Lemma 5.48. Consequentially, $\mathcal{L}_h(I)$ and $\mathcal{L}_s(I)$ have the desired structure.

We are now ready to finish proving Theorem 5.1.

Proof. By Theorem 5.6, we may find a pointed order $(P, \leq^h)$ and a separated extension $(P, \leq^s)$ such that $(\mathcal{D}(P, \leq^s), \mathcal{D}(P, \leq^h))$ is isomorphic to $(\mathcal{L}_s, \mathcal{L}_h)$. Let I be the instance created by Algorithm 5.36 given $(P, \leq^h)$ and $(P, \leq^s)$. By Corollary 5.56 and Lemma 5.47, the lattice of hub-stable matchings over I is isomorphic to $\mathcal{L}_h$, and the bijection γ maps $\mathcal{D}(P, \leq^h)$ to the set of all hub-stable matchings over I. Furthermore, by Lemma 5.48, γ also maps $\mathcal{D}(P, \leq^s)$ to the set of all stable matchings over I that are $\subseteq G_h$; however, every stable matching is hub-stable, and every hub-stable matching is $\subseteq G_h$ by Theorem 5.55, so γ maps $\mathcal{D}(P, \leq^s)$ to the set of all stable matchings over I. Therefore, $(\mathcal{L}_s(I), \mathcal{L}_h(I))$ is isomorphic to $(\mathcal{D}(P, \leq^s), \mathcal{D}(P, \leq^h))$ - which is isomorphic to $(\mathcal{L}_s, \mathcal{L}_h)$. $\square$

5.5 Lattices of the Odd-Stable Matchings

We recall that the operator $\psi : 2^{E(G(I))} \to 2^{E(G(I))}$ maps any $S \subseteq E(G(I))$ to the union of all S-stable matchings over I, and that a matching is k-stable over I if it is $\psi_I^k(\emptyset)$-stable. By Theorem 4.9, for all $r \in \mathbb{N}$, $\psi_I^{2r+2}(\emptyset) \subseteq \psi_I^{2r+1}(\emptyset)$, so $\psi_I^{2r+1}(\emptyset)$ is stable-closed. As a result, by Theorem 3.10, the $(2r+1)$-stable matchings over I form a distributive lattice $\mathcal{L}_r(I)$ under the domination ordering. As an extension of Theorem 5.1, we may consider what the sequence of lattices $(\mathcal{L}_0(I) = \mathcal{L}_s(I), \mathcal{L}_1(I), \ldots, \mathcal{L}_z(I) = \mathcal{L}_h(I))$ can look like, where $z \in \mathbb{N}$.

Proposition 5.57. *For all $r \leq z$, $\mathcal{L}_r(I)$ is a distributive lattice under the domination ordering.*

Proof. Every $(2r+1)$-stable matching M is $\subseteq \psi_I^{2r+2}(\emptyset) \subseteq \psi_I^{\infty} \subseteq \psi_I^{2r+1}(\emptyset)$ (by Theorem 4.9); consequentially, by Theorem 3.10, $\mathcal{L}_r(I)$ is a distributive lattice under the domination ordering. $\square$

Proposition 5.58. *For all $r \in [z]$, $\mathcal{L}_{r-1}(I)$ is a cover-preserving sublattice of $\mathcal{L}_r(I)$.*

Proof. Consider the instance $I_{r-1} = I[\psi_I^{2r-1}(\emptyset)]$; over this instance, $\mathcal{L}_{r-1}(I)$ is the lattice of stable matchings and $\mathcal{L}_r(I)$ is the lattice of 3-stable matchings. By Theorem 4.9 and Theorem 3.12, $\mathcal{L}_{r-1}(I)$ is a sublattice of $\mathcal{L}_r$ that preserves the property of covering. □

Corollary 5.59. *For all $r < r'$, $\mathcal{L}_r(I)$ is a cover-preserving sublattice of $\mathcal{L}_{r'}(I)$.*

As a result, we see that $(\mathcal{L}_0(I), \mathcal{L}_1(I), \ldots, \mathcal{L}_z(I))$ is a covering lattice z-flag. This naturally leads to the question of, for any given $z \in \mathbb{N}$, what lattice z-flags are isomorphic to $(\mathcal{L}_0(I), \mathcal{L}_1(I), \ldots, \mathcal{L}_z(I))$ for some instance I. The proofs of Theorem 5.12 and Theorem 5.1 are sufficient to determine that any lattice z-flag can be represented in such a fashion when $z = 1$ and 2 respectively; we therefore focus our efforts on the cases where $z \geq 3$.

By the Birkhoff Representation Theorem, we construct a pointed order $(P, \leq^z)$ to create an isomorphism between $\mathcal{L}_z = \mathcal{L}_h$ and $\mathcal{D}(P, \leq^z)$. Since $\mathcal{L}_r$ is a cover-preserving sublattice of $\mathcal{L}_h$ for all $r \in [z]$, we may invoke Corollary 5.7 and identify, for each $\mathcal{L}_r$, a corresponding extension $(P, \leq^r)$ of $(P, \leq^z)$. By the fact that every $\mathcal{L}_{r-1}$ is a cover-preserving sublattice of $\mathcal{L}_r$ for $r \in [z]$, $(P, \leq^{r-1})$ is an extension of $(P, \leq^r)$ with the property that all equivalence classes other than those that include $\hat{0}$ and $\hat{1}$ have size exactly 1. This representation of the lattice z-flag $(\mathcal{L}_0(I), \mathcal{L}_1(I), \ldots, \mathcal{L}_z(I))$ allows us to note the following necessary properties.

Proposition 5.60. *For all $r \in [z-1]$, if the equivalence classes of $\hat{0}$ in $(P, \leq^r)$ and $(P, \leq^{r-1})$ are the same, then both are $\{\hat{0}\}$.*

Proof. In order to prove ths statement, we prove the following equivalent statement: for any $r \in [z-1]$, if a matching M_0 is the mentor-optimal matching in $\mathcal{L}_{r-1}$ and $\mathcal{L}_r$, then it is the mentor-optimal matching in $\mathcal{L}_{r'}$ for all $r \leq r' \leq z$.

Consider the faithful truncation $I' \equiv I_{(M_0,\emptyset)}$, the restriction of I to all edges $e \in G(I)$ such that m_e weakly prefers d_e to $p_{M_0}(m_e)$. We note that $\psi_I^{2r}(\emptyset) \cap G(I') = \psi_I^{2r+2}(\emptyset) \cap$

$G(I') = M_0$, so by Theorem 4.29 and the fact that M_0 is $(2r-1)$-stable, $\psi_{I'}^2(M_0) = M_0$. This also informs us that $\psi_{I'}^{2s}(M_0) = M_0$ for all $s \in \mathbb{N}$.

M_0 is trivially in $\mathcal{L}_{r'}$ for all $r' > r$ as well, so by Theorem 4.29, $\psi_I^{2r'+2}(\emptyset) \cap G(I') = \psi_I^{2(r'-r+1)}(\psi_I^{2r}(\emptyset)) \cap G(I') = \psi_{I'}^{2(r'-r+1)}(M_0) = M_0$. However, since $M_0 \in \mathcal{L}_{r'}$, the mentor-optimal matching in $\mathcal{L}_{r'}$ must weakly dominate it - i.e. consist only of edges in $G(I')$. The only such edges that can appear in a $\psi_I^{2r'+1}$-stable matching are those in M_0, so M_0 is left as the only candidate for the mentor-optimal matching in $\mathcal{L}_{r'}$. $\square$

The following proposition follows analogously.

Proposition 5.61. *For all $r \in [z-1]$, if $\{p \in P : \hat{1} \leq^{r-1} p\} \supset \{\hat{1}\}$, then $\{p \in P : \hat{1} \leq^r p\} \subset \{p \in P : \hat{1} \leq^{r-1} p\}$.*

The final necessary constraint on such a lattice z-flag restricts the conditions under which two elements can be covering in $(P, \leq^r)$ and independent in $(P, \leq^z)$.

Theorem 5.62. *Let $p_1, p_2 \in P$, and $r \in [z]$ be any element such that both p_1 and p_2 are in their own equivalence classes in $(P, \leq^r)$. If p_2 covers p_1 in $(P, \leq^r)$ and $p_1 \not\leq^z p_2$, then $p_1, p_2 \leq^{r-1} \hat{0}$ or $p_1, p_2 \geq^{r-1} \hat{1}$.*

We ultimately prove this by the following proposition. In this proposition, $\mathcal{L}_s(I)$ is the lattice of stable matchings, $\mathcal{L}_c(I)$ is the lattice of 3-stable matchings, and $\mathcal{L}_h(I)$ is the lattice of hub-stable matchings. Using Proposition 5.58, we note that $(\mathcal{L}_s(I), \mathcal{L}_c(I), \mathcal{L}_h(I))$ is a lattice 3-flag. By Corollary 5.5, we can identify a pointed order $(P, \leq^h)$, a (pointed quasi-order) extension $(P, \leq^c)$, and a (pointed quasi-order) $(P, \leq^s)$ of $(P, \leq^c)$ such that $(\mathcal{L}_s(I), \mathcal{L}_c(I), \mathcal{L}_h(I))$ is isomorphic to $(\mathcal{D}(P, \leq^s), \mathcal{D}(P, \leq^c), \mathcal{D}(P, \leq^h))$. Recall that γ is an order-preserving bijection from $\mathcal{D}(P, \leq^h)$ to $\mathcal{L}_h(I)$ that also maps $\mathcal{D}(P, \leq^s)$ (resp. $\mathcal{D}(P, \leq^c)$) to $\mathcal{L}_s(I)$ (resp. $\mathcal{L}_c(I)$).

Proposition 5.63. *Let $p_1, p_2 \in P$. If neither p_1 nor p_2 are in the same equivalence class as $\hat{0}$ or $\hat{1}$ in $(P, \leq^c)$, p_2 covers p_1 in $(P, \leq^c)$, and $p_1 \not\leq^h p_2$, then either $p_1, p_2 \leq^s \hat{0}$ or $p_1, p_2 \geq^s \hat{1}$.*

Proof. Let $\rho_1 = \mu(p_1), \rho_2 = \mu(p_2)$ be rotations over $I[\psi_I^\infty]$ as defined by Proposition 5.9.

Since p_1 and p_2 aren't ordered in $(P, \leq^h)$, we know that ρ_1 and ρ_2 don't share any vertices; WLOG, we say $\rho_1 = (\{(m_1, d_1), \ldots, (m_a, d_a)\}, \{(m_1, d_2), \ldots, (m_{a-1}, d_a), (m_a, d_1)\})$ and $\rho_2 = (\{(m_{a+1}, d_{a+1}), \ldots, (m_b, d_b)\}, \{(m_{a+1}, d_{a+2}), \ldots, (m_{b-1}, d_b), (m_b, d_{a+1})\})$.

Let $D = \{p \in P : p \leq^c p_2, p \notin p_1, p_2\}$; then, D and $D \cup \{p_1, p_2\}$ are both $\in \mathcal{D}(P, \leq^c)$ (since p_2 covers p_1 in $(P, \leq^c)$, whereas $D \cup \{p_2\} \in \mathcal{D}(P, \leq^h)$ but not $\mathcal{D}(P, \leq^c)$. As a result, we note that $M \equiv \gamma(D)$ and $M'' \equiv \gamma(D \cup \{p_1, p_2\}) = M \cup (\rho_1)_d \cup (\rho_2)_d - (\rho_1)_m - (\rho_2)_m$ are 3-stable, whereas $M' \equiv \gamma(D \cup \{p_2\}) = M \cup (\rho_2)_d - (\rho_2)_m$ is hub-stable but not 3-stable. This implies the existence of an edge $e \in \psi_I^3(\emptyset)$ that destabilizes M', but not M or M''. We make the following observations about e.

- If $m_e \notin \rho_2$, then $p_{M'}(m_e) = p_M(m_e)$, and d_e prefers $p_{M'}(d_e)$ to $p_M(d_e)$; therefore, if e destabilizes M', it also destabilizes M. This contradicts M being 3-stable, so $m_e \in \rho_2$.

- If $d_e \notin \rho_1$, then $p_{M'}(d_e) = p_{M''}(d_e)$, and m_e prefers $p_{M'}(m_e)$ to $p_{M''}(m_e)$; therefore, if e destabilizes M', it also destabilizes M'. This contradicts M'' being 3-stable, so $d_e \in \rho_1$.

Therefore, $m_e \in \rho_2$ and $d_e \in \rho_1$. (WLOG, we may assume that $m_e = m_b$ and $d_e = d_1$.) In order to ensure that e destabilizes only M', we must have m_b prefer d_b to d_1 to d_{a+1}, and d_1 prefer m_a to m_b to m_1. However, since M' is hub-stable, $e \notin \psi_I^\infty$, and so $e \notin \psi_I^2(\emptyset)$ by Theorem 4.9. Since $e \in \psi_I^3(\emptyset)$, by Corollary 4.47, e must uphold one of the following:

- $m_e = m_b$ prefers $d_e = d_1$ to their partner in the mentor-optimal stable matching M_0 over I. However, d_{a+1} is m_b's top choice among the students that they prefers d_1 to, and can be partnered with in a 3-stable matching. M_0 is also a 3-stable matching, so m_b also prefers d_{a+1} to $p_{M_0}(m_b)$; this means that m_b strictly prefers d_b to $p_{M_0}(m_b)$. However, this means that $\mu^{-1}(M_0)$ contains p_2; since $\mu^{-1}(M_0)$ is the smallest downset in $\mathcal{D}(P, \leq^s)$, then every downset in it contains p_2, so $p_2 \leq^s \hat{0}$. In addition, $p_1 \leq^s p_2$ (since $\leq^s$ is an extension of $\leq^c$), so $p_1 \leq^s \hat{0}$ as well.

- $d_e = d_1$ prefers $m_e = m_b$ to their partner in the student-optimal stable matching M_1 over I. However, m_1 is d_1's top choice among the mentors that they prefers

m_b to, and can be partnered with in a 3-stable matching. M_1 is also a 3-stable matching, so d_1 also prefers m_1 to $p_{M_1}(d_1)$; as a result, d_1 does not strictly prefer $p_{M_1}(d_1)$ to m_1. However, this means that $\mu^{-1}(M_1)$ does not contain p_1; since $\mu^{-1}(M_1)$ is the largest downset in $\mathcal{D}(P, \leq^s)$, then every downset in it does not contain p_1, so $p_1 \geq^s \hat{1}$. In addition, $p_2 \geq^s p_1$ (since $\leq^s$ is an extension of $\leq^c$), so $p_2 \geq^s \hat{1}$ as well.

In either case, we see that $p_1, p_2 \leq^s \hat{0}$ or $p_1, p_2 \geq^s \hat{1}$, so we are done. $\square$

We note that this generalizes to Theorem 5.62.

Proof. Consider the instance $I' = I[\psi_I^{2r-1}(\emptyset)]$. By Corollary 4.28 and Theorem 4.9, $\mathcal{L}_s(I') = \mathcal{L}_{r-1}(I)$, $\mathcal{L}_c(I') = \mathcal{L}_r(I)$, and $\mathcal{L}_h(I') = \mathcal{L}_z(I')$; they correspond to $\mathcal{D}(P, \leq^{r-1})$, $\mathcal{D}(P, \leq^r)$, and $\mathcal{D}(P, \leq^z)$ respectively. By Proposition 5.63, we see that $p_1, p_2 \leq^{r-1} \hat{0}$ or $p_1, p_2 \geq^{r-1} \hat{1}$. $\square$

5.6 The Representation Theorem For Lattice 3-Flags

Our explorations into the relationships between the lattices of $(2r+1)$-stable matchings over a given instance I allows us to conjecture on the possible structures created by the lattice flag of all such lattices for a given instance.

Conjecture 5.64. *Let $(\mathcal{D}(P, \leq^0), \mathcal{D}(P, \leq^1), \ldots, \mathcal{D}(P, \leq^z))$ be a lattice z-flag. Then, there exists an instance I such that $(\mathcal{L}_0(I), \mathcal{L}_1(I), \ldots, \mathcal{L}_z(I) = \mathcal{L}_h(I))$ is isomorphic to $(\mathcal{D}(P, \leq^0), \mathcal{D}(P, \leq^1), \ldots, \mathcal{D}(P, \leq^z))$ iff the following properties hold:*

1. *For all $r \in [z]$, $(P, \leq^{r-1})$ is a separated extension of $(P, \leq^r)$. (By Theorem 5.6, this condition is necessary and sufficient for $(\mathcal{D}(P, \leq^0), \mathcal{D}(P, \leq^1), \ldots, \mathcal{D}(P, \leq^z))$ to be a covering lattice z-flag.)*

2. *For any $r \in [z]$, if the equivalence classes of $\hat{0}$ in $(P, \leq^r)$ and $(P, \leq^{r-1})$ are the same, then both are $\{\hat{0}\}$.*

3. *For any $r \in [z]$, if the equivalence classes of $\hat{1}$ in $(P, \leq^r)$ and $(P, \leq^{r-1})$ are the same, then both are $\{\hat{1}\}$.*

4. *Let* $r \in [z]$. *Then, for every* $p, p' \in P$ *not in the same equivalence class of* $(P, \leq^r)$, p *covers* p' *in* $(P, \leq^z)$ *iff* p *covers* p' *in* $(P, \leq^{r+1})$.

The necessity of properties 1, 2, 3, and 4 can be seen from Proposition 5.58, Proposition 5.60, Proposition 5.61, and Theorem 5.62 respectively. Therefore, in order to prove the conjecture, we only need to show that the conditions are sufficient. While we have not yet been able to show that this is the case, we have determined that these are the only necessary conditions for such a lattice 3-flag. For the following, $\mathcal{L}_c(I) \equiv \mathcal{L}_1(I)$ is the lattice of 3-stable matchings over I.

Theorem 5.65. *Let* $(P, \leq^h)$ *be a pointed order,* $(P, \leq^c)$ *be a separated extension of* $(P, \leq^h)$*, and* $(P, \leq^s)$ *be a separated extension of* $(P, \leq^c)$ *such that the following conditions are upheld:*

- *For any* $p_1, p_2 \in P$ *such that* $\hat{1} \not\leq^c p_1, p_2 \not\leq^c \hat{0}$, p_2 *covers* p_1 *in* $(P, \leq^c)$, *and* $p_1 \not\leq^h p_2$, *either* $p_1, p_2 \leq^s \hat{0}$ *or* $p_1, p_2 \geq^s \hat{1}$.
- *If the equivalence classes of* $\hat{0}$ *are the same for* $\leq^s$ *and* $\leq^c$, *then both are* $\{\hat{0}\}$.
- *If the equivalence classes of* $\hat{1}$ *are the same for* $\leq^s$ *and* $\leq^c$, *then both are* $\{\hat{1}\}$.

Then, there exists an instance I *such that* $(\mathcal{L}_s(I), \mathcal{L}_c(I), \mathcal{L}_h(I))$ *is isomorphic to* $(\mathcal{D}(P, \leq^s), \mathcal{D}(P, \leq^c), \mathcal{D}(P, \leq^h))$.

We show this via the following construction. For this construction, P_0 and P_1 are the equivalence classes of $\hat{0}$ and $\hat{1}$ respectively in $(P, \leq^c)$. In addition, P_0^* and P_1^* are the equivalence classes of $\hat{0}$ and $\hat{1}$ respectively in $(P, \leq^s)$.

Algorithm 5.66. *Let* $(P, \leq^h)$ *be a pointed order,* $(P, \leq^c)$ *be a separated extension of* $(P, \leq^h)$*, and* $(P, \leq^s)$ *be a separated extension of* $(P, \leq^c)$ *such that the conditions in Theorem 5.65 are upheld. We construct a set of mentors* V_m *and a set of students* V_d *such that each vertex has a preference list consisting of vertices of the other type as follows:*

1. *Let* $k = |P| - 2$, *and* $P = \{p_0, \ldots, p_{k+1}\}$ *be any reference ordering of* P *as defined by Proposition 5.2.*

(a) *If* $P_0^* - P_0 = \emptyset$*, then set* $i_0 = -2$*; otherwise, set* i_0 *to be the least index among the elements of* $P_0^* - P_0$*. (We note that if* $i_0 = -2$*, then steps 4a(i) and 4c will never occur nonvacuously.)*

(b) *If* $P_1^* - P_1 = \emptyset$*, then set* $i_1 = -2$*; otherwise, set* i_1 *to be the greatest index among the elements of* $P_1^* - P_1$*. (We note that if* $i_1 = -2$*, then steps 4a(ii) and 4f will never occur nonvacuously.)*

2. *Let* $H(P)$ *be the Hasse diagram of* $(P, \leq^h)$*, and* $E = E(H(P))$*. The instance* I *will have* $V_m(I) = \{m_e : e \in E\} \cup \{m_\tau, m_\sigma, m_{\sigma'}, m_{\sigma''}, m_\upsilon, m_{\upsilon'}, m_{\upsilon''}\}$*, and* $V_d(I) = \{d_e : e \in E\} \cup \{d_\tau, d_\sigma, d_{\sigma'}, d_{\sigma''}, d_\upsilon, d_{\upsilon'}, d_{\upsilon''}\}$*.*

3. *Perform step 3 of Algorithm 5.36. In addition, for* $c \in \{\sigma, \sigma', \sigma'', \upsilon, \upsilon', \upsilon''\}$*, initialize the preference list of* m_c *by placing* d_c *on it, and the preference list of* d_c *by placing* m_c *on it. Set* $V_d' = V_d(I) - \{d_\tau\}$*.*

4. *For* i *from* 0 *to* $k+1$*, iterate the following steps:*

   (a) *Perform the following:*

      i. *If* $i = i_0$*, then let* $A_i = \{a_i(1), \ldots, a_i(r_i)\}$ *such that* $a_i(1) = \sigma$*,* $a_i(2) = \sigma'$*,* $a_i(3) = \sigma''$*, and* $\{a_i(4), \ldots, a_i(r_i)\}$ *is an arbitrary ordering of the edges incident with node* i *in* H *such that* $a_i(4)$ *is incident with* p_i *from above. Let* $B_i = \{b_i(1), \ldots, b_i(r_i)\}$ *such that for all* $j \in [r_i]$*,* $d_{b_i(j)}$ *is the last element in* V_d' *on* $m_{a_i(j)}$*'s current preference list. Then, for all* $j \in [r_i]$*, place* $d_{b_i(j+1)}$ *at the bottom of* $m_{a_i(j)}$*'s preference list and* $m_{a_i(j)}$ *at the top of* $d_{b_i(j+1)}$*'s preference list, where* $j+1$ *is taken* $\bmod r_i$*.*

      ii. *If* $i = i_1$*, then let* $A_i = \{a_i(1), \ldots, a_i(r_i)\}$ *such that* $a_i(1) = \upsilon$*,* $a_i(2) = \upsilon'$*,* $a_i(3) = \upsilon''$*, and* $\{a_i(4), \ldots, a_i(r_i)\}$ *is an arbitrary ordering of the edges incident with node* i *in* H *such that* $a_i(r_i)$ *is incident with* p_i *from below. Let* $B_i = \{b_i(1), \ldots, b_i(r_i)\}$ *such that for all* $j \in [r_i]$*,* $d_{b_i(j)}$ *is the last element in* V_d' *on* $m_{a_i(j)}$*'s current preference list. Then, for all* $j \in [r_i]$*, place* $d_{b_i(j+1)}$ *at the bottom of* $m_{a_i(j)}$*'s preference list and* $m_{a_i(j)}$ *at the top of* $d_{b_i(j+1)}$*'s preference list, where* $j+1$ *is taken* $\bmod r_i$*. Lastly,* *remove $d_{v''}$ from V'_d.*

iii. *Otherwise, perform step 4a of Algorithm 5.36.*

(This is functionall the same as step 4 of Algorithm 5.11 applied to $(P, \leq^h)$, ignoring any elements of $V_d(I) - V'_d$ - see Lemma 5.78.)

(b) *If $p_i \in P - P_0 - P_1$, then let $x(i) \in E$ be any edge incident to p_i from below, $x'(i)$ be the last element of V'_d that appears on $m_{x(i)}$'s preference list, and $y(i) \in E$ be any edge incident to p_i from above. Then, do the following:*

i. *If $p_i \in P_0^*$, then, for every $p_j \in P_0^* - P_0$ such that $j < i$, $p_j \not\leq^h p_i$, and p_i covers p_j in $(P, \leq^c)$, place $d_{x(j)}$ second from the bottom on $m_{y(i)}$'s preference list and $m_{y(i)}$ second from the top on $d_{x(j)}$'s preference list.*

ii. *If $p_i \in P_1^*$, then, for every $p_j \in P_1^* - P_1$ such that $j < i$, $p_j \not\leq^h p_i$, and p_i covers p_j in $(P, \leq^c)$, place $d_{x(j)}$ second from the bottom on $m_{y(i)}$'s preference list and $m_{y(i)}$ second from the top on $d_{x(j)}$'s preference list. (This step and the previous one ensure that rotations corresponding to elements that are totally ordered in $(P, \leq^c)$ but not $(P, \leq^h)$ are totally ordered in $\Pi(I[\psi_I^3(\emptyset)])$ but not $\Pi(I[\psi_I^\infty])$ - see Lemma 5.85 and Lemma 5.86.)*

iii. *Otherwise, for each $p_j \in P - P_0^* - P_1^*$ such that $j < i$, $p_j \not\leq^h p_i$, and p_i covers p_j in $(P, \leq^s)$, place $d_{x(j)}$ second from the bottom on $m_{y(i)}$'s preference list, and $m_{y(i)}$ second from the top on $d_{x(j)}$'s preference list. (This ensures that rotations corresponding to elements that are totally ordered in $(P, \leq^s)$ but not $(P, \leq^h)$ are totally ordered in $\Pi(I)$ but not $\Pi(I[\psi_I^\infty])$ - see Lemma 5.79 and Lemma 5.83.)*

(c) *If $i = i_0 - 1$, let E_0 be the set of all edges $e \in E(H(P))$ that are incident with p_j for some $j \leq i$. Then, for every edge $e \in E_0$ (in any order), do the following: place $m_{\sigma''}$ second from the top of d_e's preference list and d_e at the top of $m_{\sigma''}$'s preference list . In addition, place $d_{\sigma'}$ at the top of m_e's preference list, and m_e at the bottom of $d_{\sigma'}$'s preference list. (This ensures that rotations corresponding to elements of P that are $\leq^c \hat{0}$ don't appear in $\Pi(I[\psi_I^3(\emptyset)])$ - see Lemma 5.85 and Lemma 5.87.)*

(d) *If* p_i *is the last element of* P_0^**, then, for every* $d_e \in V_d(I) - d_\tau$*, place* m_τ *second from the top of* d_e*'s preference list and* d_e *at the top of* m_τ*'s preference list (in any order). (This ensures that rotations corresponding to elements of* P *that are* $\leq^s \hat{0}$ *don't appear in* $\Pi(I)$ *- see Lemma 5.79 and Lemma 5.84.)*

(e) *If* p_i *is the last element of* $P - P_1^*$*, then, for every* $m_e \in V_m(I) - m_\tau$*, place* d_τ *at the bottom of* m_e*'s preference list and* m_e *at the top of* d_τ*'s preference list (in any order). (This ensures that rotations corresponding to elements of* P *that are* $\geq^s \hat{1}$ *don't appear in* $\Pi(I)$ *- see Lemma 5.79 and Lemma 5.84.)*

(f) *If* $i = i_1$*, let* E_1 *be the set of all edges* $e \in E(H(P))$ *that are incident with* p_j *for some* $j > i$*, and* D_1 *be the set of all students that appear as the last element of* V_d' *on* m_e*'s preference list for some* $e \in E_1$*. Then, for every edge* $e \in E_1$ *(in any order), place* $d_{\upsilon''}$ *at the bottom of* m_e*'s preference list and* m_e *at the top of* $d_{\upsilon''}$*'s preference list. In addition, for every* $d_{e'} \in D_1$*, place* $d_{e'}$ *at the bottom of* m_υ*'s preference list. (This step, along with step 5, ensures that rotations corresponding to elements of* P *that are* $\geq^c \hat{1}$ *don't appear in* $\Pi(I[\psi_I^3(\emptyset)])$ *- see Lemma 5.85 and Lemma 5.88.)*

5. *For every* $d_{e'} \in D_1$*, place* m_υ *at the top of* $d_{e'}$*'s preference list.*

Proposition 5.67. *For all* $m \in V_m(I)$ *and* $d \in V_d(I)$*,* m *and* d *appear on each other's preference list in* I *an equal number of times.*

Proof. Each step of Algorithm 5.66 adds m to d's preference list iff it adds d to m's preference list - with two families of exceptions:

- In the i_1th iteration of step 4f, for every $d_{e'} \in D_1$, $d_{e'}$ is added to m_υ's preference list.
- In step 5, for every $d_{e'} \in D_1$, m_υ is added to $d_{e'}$'s preference list.

However, for each (m, d) such that m adds d to their preference list as part of the first family, d adds m to their preference list as part of the second family (and vice versa). As a result, we are done. $\square$

Proposition 5.68. *Steps 4c, 4d, 4e, and 4f each occur nonvacuously for at most one value of i, and in that order.*

Proof. If it occurs nonvacuously, step 4c (resp. 4d, 4e, 4f) occurs nonvacuously only in the ith iteration of step 4, where i is the greatest index among the elements of P_0 (resp. P_0^*, $P - P_1^*$, $P - P_1$). Since $P_0 \subseteq P_0^*$, $P_1 \subseteq P_1^*$, and P_0^* is disjoint from P_1^*, we see that $P_0 \subseteq P_0^* \subseteq P - P_1^* \subseteq P - P_1$, implying the proposition. $\square$

For the instance I output by Algorithm 5.66, let G_h be the set of edges (m, d) such that m and d add each other to their preference lists in step 3 or 4a. We note that the restriction $I[G_h]$ is the same as the instance I_0 constructed by Algorithm 5.11 on input of $(P, \leq^h)$ and $Q = [\hat{0}, p_{i_0}, p_{i_0}, p_{i_0}, p_{i_1}, p_{i_1}, p_{i_1}]$; in this case, we use $E' = \{\tau, \sigma, \sigma', \sigma'', \upsilon, \upsilon', \upsilon''\}$, where $\tau = (\hat{1}, \hat{0})$, $\sigma = \sigma' = \sigma'' = (\hat{1}, p_{i_0})$, and $\upsilon = \upsilon' = \upsilon'' = (\hat{1}, p_{i_1})$. (We refer to this multi-digraph as H'.)

In order to show that this construction creates an instance I where the lattice 3-flag $(\mathcal{L}_s(I), \mathcal{L}_c(I), \mathcal{L}_h(I))$ is isomorphic to $(\mathcal{D}(P, \leq^s), \mathcal{D}(P, \leq^c), \mathcal{D}(P, \leq^h))$, we show the following lemmas centered around the restriction $I[G_h]$, in this order. (Recall that ψ_I^∞ is the unique hub of I.)

1. The preference lists restricted to G_h are the same as the preference lists created by Algorithm 5.11 on input of $(P, \leq^h)$, with some generality removed (see Theorem 5.69).

2. The preference lists created by Algorithm 5.66 do not repeat any elements. (This implies that I is an instance - see Theorem 5.75).

3. The lattice of stable matchings over $I[G_h]$ is isomorphic to $\mathcal{D}(P, \leq^h)$ via an order isomorphism γ (see Lemma 5.78).

4. The lattice of the stable matchings over I which are $\subseteq I[G_h]$ is isomorphic to $\mathcal{D}(P, \leq^s)$ via γ (see Lemma 5.79).

5. There exists a set of edges $S_2 \subseteq E(G(I))$ such that the lattice of the S_2-stable matchings over I which are $\subseteq I[G_h]$ is isomorphic to $\mathcal{D}(P, \leq^c)$ via γ (see Lemma 5.85).

6. $\psi_I^\infty = G_h$, and the set of hub-stable matchings over I is the set of stable matchings over $I[G_h]$ (see Lemma 5.89 and Corollary 5.90).

7. $S_2 = \psi_I^3(\emptyset)$ (see Lemma 5.91).

5.6.1 The Structure of $I[G_h]$ in Algorithm 5.66

In this subsection, we will carry out step 1 in our outline of the proof of Theorem 5.65.

Theorem 5.69. *Given any pointed order $(P, \leq^h)$ and separated extension $(P, \leq^s)$, the preference lists created by Algorithm 5.66, restricted to the elements added during steps 3 and 4a, is the set of preference lists created by Algorithm 5.11 on input of $(P, \leq^h)$ and $Q = [\hat{0}, p_{i_0}, p_{i_0}, p_{i_0}, p_{i_1}, p_{i_1}, p_{i_1}]$.*

Proof. We begin by proving the following lemma.

Lemma 5.70. *The instance created by running Algorithm 5.66 without running steps 4b through 4f is identical to that created by Algorithm 5.11 on input of $(P, \leq^h)$ and $Q = [\hat{0}, p_{i_0}, p_{i_0}, p_{i_0}, p_{i_1}, p_{i_1}, p_{i_1}]$.*

Proof. The only place where the modified Algorithm 5.66 differs from Algorithm 5.11 is in step 4 - for each $i \in [k]$ and $d \in V_d(I)$, $d \in B_i$ iff d is the last element of V_d' to appear in the preference list of some $m_{a_i(j)}$ (as opposed to being the last element of $V_d(I)$ in some such preference list). As a result, to prove the lemma, we need only to show that for all $a_i(j)$, any student $\notin V_d'$ is not the last element on $m_{a_i(j)}$'s preference list.

Since τ is not incident with any node in $\{p_1, \ldots, p_k\}$, m_τ has d_τ at the bottom of their preference list throughout the entirety of step 4 - which implies that no other mentor does by Proposition 5.13. Similarly, since $m_{\upsilon'}$ adds $d_{\upsilon''}$ to the bottom of their preference list during the i_1th iteration of step 4, and υ' is not incident with any node in $\{p_{i_1+1}, \ldots, p_k\}$, $m_{\upsilon'}$ has $d_{\upsilon''}$ at the bottom of their preference list for every iteration of step 4 past the i_1th one - which implies that no other mentor does by Proposition 5.13. $\square$

Given this lemma, we need only to show that steps 4b through 4f never change the element of V'_d that appears last in any mentor's preference list.

- In the ith iteration of step 4, for each relevant $j < i$, step 4b places a student second from the bottom of $m_{y(i)}$'s preference list; this does not change the element of $\{d_e : e \in E\}$ that appears last in any mentor's preference list.

- In step 4c, no mentor has any student added to the bottom of their preference list.

- The only mentor that sees their preference list change in step 4d is m_τ; however, since τ is not incident in H' with p_i for any $i \in [k]$, their preference list has no effect on the operation of step 4a.

- The only student that is added to any mentor's preference list in step 4e is d_τ; since $d_\tau \notin V'_d$, this has no impact on the operation of step 4a.

- In step 4f, some number of mentors have $d_{\upsilon''}$ added to the bottom of their preference lists, and m_υ has the bottom of their preference list changed. However, since $d_{\upsilon''}$ has already been removed from V'_d, and $\upsilon \in H'$ is not incident with any p_i for $i > i_0$, neither of these have any effect on the operation of step 4a.

As a result, we see that no change in the preference list from steps 4b through 4f changes the preference additions made by step 4a, and so we are done. □

For all $i \in [k]$, we define $\rho_m(i) = \{(m_{a_i(1)}, d_{b_i(1)}), \ldots, (m_{a_i(r_i)}, d_{b_i(r_i)})\}$ and $\rho_d(i) = \{(m_{a_i(1)}, d_{b_i(2)}), \ldots, (m_{a_i(r_i-1)}, d_{b_i(r_i)}), (m_{a_i(r_i)}, d_{b_i(1)})\}$. Furthermore, we set $\rho(i) = (\rho_m(i), \rho_d(i))$. As a consequence of Theorem 5.23 and Proposition 5.24, the following holds:

Corollary 5.71. *Every edge in G_h appears in some stable matching over $I[G_h]$. In addition, the set of rotations over $I[G_h]$ is $\{\rho(i) : i \in [k]\}$, and $\Pi(I[G_h])$ has the partial ordering that for $i, j \in [k]$, $\rho(j) \leq \rho(i)$ iff $p_j \leq^h p_i$.*

5.6.2 The Proof That Algorithm 5.66 Creates an Instance

In this subsection, we carry out step 2 in our outline of the proof of Theorem 5.65. As with Algorithm 5.36, it is not immediately obvious that the preference lists produced by Algorithm 5.66 describe a stable matching instance - in order for this to be the case, we need every such preference list to consist of distinct elements. By Theorem 5.69, the pruned Algorithm 5.66 is an instance that can be created by Algorithm 5.11; by Proposition 5.24, the steps that add to the vertices' preference lists in the pruned Algorithm 5.66 (i.e. steps 3 and 4a) don't produce any repeated elements on any preference list. As a result, we only need to make certain that steps 4b through 4f don't produce any repeated elements, or add an element already added by step 3 or 4a.

We recall that for all $i \in [k]$ such that $p_i \notin P_0 \cup P_1$, $y(i)$ is an edge incident with p_i from above, $x'(i)$ is an edge incident with p_i from below, and $d_{x(i)}$ is the last element of V_d' to appear on $m_{x'(i)}$'s preference list (as of the ith iteration of step 4b).

Proposition 5.72. *For all $i \in [k]$ such that $p_i \in P - P_0 - P_1$, from the ith iteration of step 4a onwards, $d_{x(i)}$ has $m_{x'(i)}$ as the first element of their preference list, and $m_{x'(i)}$ has $d_{x(i)}$ as the last element of their preference list in $\{d_e : e \in E\}$.*

Proof. Since $d_{x(i)}$ and $m_{x'(i)}$ were added to one another's preference lists in the ith iteration of step 4a, $d_{x(i)}$ has $m_{x'(i)}$ as the first element of their preference list and $m_{x'(i)}$ has $d_{x(i)}$ as the last element of their preference list. For all $e \in E$, steps 4b through 4f cannot add an element to the top of d_e's preference list, or put d_e at the bottom of any mentor's preference list; thus, we see that only step 4a could introduce an element that breaks the property in the proposition.

Assume for the sake of contradiction that there exists a minimum $j > i$ such that at least one of $m_{x'(i)}$ and $d_{x(i)}$ changes their preference list in the jth iteration of step 4a. Since $d_{x(i)}$ is still the last element of V_d' on $m_{x'(i)}$'s preference list before this (since, as the proof of Proposition 5.38 shows, no other step can change the last element of $m_{x'(i)}$'s preference list), they cannot change their preference list unless $m_{x'(i)}$ changes their as well.

In order for $m_{x'(i)}$ to change their preference list in the jth iteration of step 4a, p_j

must be incident with $x'(i)$; however, since $x'(i)$ is incident with p_i from below, the only other vertex $x'(i)$ is incident with must have index $< i$. This creates a contradiction, so neither vertex expands its preference list in step 4a after the ith iteration, and so we are done. $\square$

Proposition 5.73. *For all $i \in [k]$ such that $p_i \in P - P_0 - P_1$, $m_{y(i)}$ and $d_{y(i)}$ do not add any element of $\{d_e : e \in E - \{y(i)\}\}$ or $\{m_e : e \in E - \{y(i)\}\}$ to their preference lists before the ith iteration of step 4a.*

Proof. Since such an e has $e \notin \{y(i), \tau, \sigma, \sigma', \sigma'', \upsilon, \upsilon', \upsilon''\}$, we note that any such addition can only occur in step 4a or 4b. Let j be the smallest natural number such that $m_{y(i)}$ or $d_{y(i)}$ adds to their preference list in the jth iteration of step 4a. Since step 4a is the only time that $m_{y(i)}$ can change the last element of $\{d_e : e \in E\}$ on their preference list, this implies that $d_{y(i)}$ is the last such element prior to that step. Thus, we see that $y(i)$ must be incident with p_j and p_k for some $k > j$. However, since $y(i)$ is incident with p_i from above, it is incident with p_i and p_k for some $k > i$ - thereby implying that $i = j$.

Furthermore, since the only vertices that $y(i)$ is incident to have index $\geq i$, $y(i) \neq x(j)$ or $y(j)$ for any $j < i$. As a result, we see that the preference lists of $m_{y(i)}$ and $d_{y(i)}$ are unchanged by steps 4a and 4b before the ith iteration, and so we are done. $\square$

Proposition 5.74. *As functions from $[k]$ to E, $x(i)$, $x'(i)$, and $y(i)$ are all injections.*

Proof. Consider any $i, j \in [k]$ such that $i < j$. Then, $y(i)$ and $y(j)$ are incident to p_i and p_j from above respectively; since $p_i \neq p_j$, $y(i) \neq y(j)$.

Now, suppose that $x(i) = x(j)$. This implies that $m_{x'(j)}$ was added to the top of $d_{x(i)}$'s preference list during the jth iteration of step 4a. However, this contradicts Proposition 5.72, so this cannot happen. $\square$

The proofs of Proposition 5.72, Proposition 5.73, and Proposition 5.74 are analogous to the proofs of Proposition 5.41, Proposition 5.42, and Proposition 5.43 respectively.

Theorem 5.75. *Given any pointed order $(P, \leq)$, let V_m and V_d (and their corresponding preference lists) be defined as in Algorithm 5.66. Then, for all $m \in V_m, d \in V_d$, m and d appear in one another's preference lists at most once.*

Proof. We begin by showing that the following two lemmas hold. For the following, let $V_\tau = \{m_\tau, d_\tau, m_{\sigma''}, d_{\sigma'}, m_\upsilon, d_{\upsilon''}\}$.

Lemma 5.76. *If $v \in V_\tau$, then v's preference list has no repeated elements.*

Proof. We consider the preference list of each element of V_τ.

- m_τ (resp. d_τ) has its preference list added to twice - once in step 3, when it adds d_τ (resp. m_τ) to its preference list, and once in step 4d (resp. 4e), when it adds every student other than d_τ (resp. every mentor other than m_τ) to its preference list. As a result, no vertex appears in the preference list of m_τ (resp. d_τ) more than once.

- $m_{\sigma''}$ adds to their preference list at most three times - adding $d_{\sigma''}$ in step 3, d_e for every $e \in E_0$, and $d_{a_i(4)}$ in the i_0th iteration of step 4a. Since $a_i(4)$ is incident with p_{i_0} from above, the other vertex that it is incident with must have greater index, and so $a_i(4) \notin E_0$; consequentially, no vertex appears in the preference list of $m_{\sigma''}$ more than once.

- $d_{\sigma'}$ adds to their preference list at most three times - adding $m_{\sigma'}$ in step 3, m_e for every $e \in E_0$, and m_σ in the i_0th iteration of step 4a. As a result, no vertex appears in the preference list of $d_{\sigma'}$ more than once.

- m_υ adds to their preference list at most three times - adding d_υ in step 3, $d_{\upsilon'}$ in the i_1th iteration of step 4a, and d for every $d \in D_1$. We note that $d_\upsilon, d_{\upsilon'} \notin D_1$[5]. As a result, no vertex appears in the preference list of $m_{\sigma''}$ more than once.

[5] As of the i_1th iteration of step 4f, they appear last on the preference lists of $m_{a_{i_1}(r)}$ - which is incident with p_{i_1} from below - and m_υ respectively. We see that $a_{i_1}(r), \upsilon \notin E_1$; in addition, by applying Proposition 5.13 to the application of Algorithm 5.11 described in Proposition 5.38, we see that d_υ and $d_{\upsilon'}$ cannot appear as the last element of V_d' on m_e's preference list for any other $m_e \in V_m(I)$, and so cannot be the last such element for any m_e such that $e \in E_1$.

- $d_{v''}$ adds to their preference list at most three times - adding $m_{v''}$ in step 3, $m_{v'}$ in the i_1th iteration of step 4a, and m_e for every $e \in E_1$. As a result, no vertex appears in the preference list of $d_{v''}$ more than once.

□

Lemma 5.77. *For any $m_{e'} \in V_m(I) - V_\tau$ and $d_e \in V_d(I) - V_\tau$, $m_{e'}$ appears on d_e's preference list at most once.*

Proof. By Proposition 5.14, steps 3 and 4a together don't add any $m_{e'}$ to d_e's preference list more than once, so we only need to show that step 4b does not cause any duplicates. (Steps 4c through 4f cannot add any $m_{e'}$ to d_e's preference list when neither is in V_τ, so we only need to show that we don't create duplicates with steps 3, 4a, and 4b.)

Consider any edge $(m_{e'}, d_e)$ such that d_e adds $m_{e'}$ to their preference list in step 4b; then, $e = x(j)$ and $e' = y(i)$ for some $i, j \in [k]$ such that $j < i$. We note that d_e's preference list is added to in the jth iteration of step 4a and - by Proposition 5.72 - not in any subsequent one. By Proposition 5.73, $m_{e'}$'s preference list is not changed by any iteration of step 4a before the ith one; therefore, since $i > j$, no iteration of step 4a changes both preference lists, which is necessary in order to add $m_{e'}$ to d_e's preference list. Similary, by Proposition 5.73, $d_{e'}$'s preference list is not changed in the jth iteration of step 4a, so $e \neq e'$ and so $m_{e'}$ is not added to d_e's preference in step 3.

As a result, we see that the theorem holds iff for any given $e, e' \in E$, step 4b adds $m_{e'}$ to d_e's preference list at most once. In total, for each $p_i, p_j \in P - P_0 - P_1$ such that $i > j$, step 4b adds $m_{y(i)}$ to $d_{x(j)}$'s preference list at most once. Furthermore, by Proposition 5.74, the function that maps (i, j) to $(y(i), x(j))$ is an injection, so we are done. □

By Proposition 5.67, the above two lemmas imply that for all $m_{e'} \in V_m(I) - V_\tau$, $d_e \in V_d(I) - V_\tau$, $m \in V_m(I) \cap V_\tau$, and $d \in V_d(I) \cap V_\tau$, d_e and d each appear on $m_{e'}$'s preference list at most once, and m appears on d_e's preference list at most once. As a result, we see that for all $m \in V_m(I)$, $d \in V_d(I)$, m and d appear in one another's preference lists at most once. □

5.6.3 The Structure of $\mathcal{L}_h(I)$, $\mathcal{L}_s(I)$, and $\mathcal{L}_c(I)$

In this subsection, we carry out steps 3 and 4 of the proof of Theorem 5.65. For all $i \in [k]$, we define $\rho_m(i) = \{(m_{a_i(t)}, d_{b_i(t)}) : t \in [r_i]\}$ and $\rho_d(i) = \{(m_{a_i(t)}, d_{b_i(t+1)}) : t \in [r_i - 1]\} \cup \{m_{a_i(r)}, d_{b_i(1)})\}$. (Recall that $(m_a, d_b) \in \rho_m(i)$ iff a is incident with $p_i \in E(H(P))$ and d_b is the last element of V_d' on m_a's preference list just prior to the ith iteration of step 4a, whereas $(m_a, d_b) \in \rho_d(i)$ iff m_a adds d_b to the bottom of their preference list during the ith iteration of step 4a.) We begin by showing that for all $i \in [k]$, $\rho(i) = (\rho_m(i), \rho_d(i))$ is a rotation over $I[G_h]$ (as defined prior to Theorem 2.11). Specifically, the following holds as the application of Theorem 5.23 to $I[G_h]$:

Lemma 5.78. *The lattice $\mathcal{L}_s(I[G_h])$ is isomorphic to $\mathcal{L}_h$. Furthermore:*

1. *$R(I[G_h]) = \{\rho(i) : i \in [k]\}$.*
2. *the bijection $\mu : P - \{\hat{0}, \hat{1}\} \to R(I[G_h])$ such that $\mu(p_i) = \rho(i)$ is an order isomorphism between $(P - \{0, 1\}, \leq^h)$ and $\pi(I[G_h])$.*
3. *every edge in G_h appears in some stable matching over $I[G_h]$.*

In particular, by Proposition 5.10, we can identify a bijection $\gamma : \mathcal{D}(P, \leq) \to \mathcal{L}_s(I)$ such that for all $D \in \mathcal{D}(P, \leq)$, $\gamma(D) = M_0 \cup (\cup_{p \in D}(\mu(p))_d) - (\cup_{p \in D}(\mu(p))_m)$. Furthermore, for all $D, D' \in \mathcal{D}(P, \leq)$, $\gamma(D)$ dominates $\gamma(D')$ iff $D \supseteq D'$.

Lemma 5.79. *Let $S \in \mathcal{D}(P, \leq^h)$. Then, $\gamma(S)$ is stable in I iff $S \in \mathcal{D}(P, \leq^s)$.*

In order to prove this lemma, we look at each $e \in E(G(I)) - G_h$, and determine the necessary and sufficient conditions for $\gamma(S)$ to be $\{e\}$-stable. Let S_b (resp. S_b', S_b'', S_c, S_δ, S_ϵ, and S_f) be the set of all edges $e \in E(G(I))$ such that m_e and d_e add one another to their preference lists in steps 4b(iii) (resp. 4b(i), 4b(ii), 4c, 4d, 4e, and 4f/5); since G_h consists of every $e \in E(G(I))$ such that m_e and d_e add one another to their preference lists in step 3 (when the preference lists are first created) or 4a, $S_b \cup S_b' \cup S_b'' \cup S_c \cup S_\delta \cup S_\epsilon \cup S_f = E(G(I)) - G_h$. We will determine the necessary and sufficient conditions for $\gamma(S)$ to be S_b-stable in Lemma 5.80, and the necessary and sufficient conditions for $\gamma(S)$ to be S_δ and/or S_ϵ-stable in Lemma 5.81; we then show that these conditions ars sufficients to show that $\gamma(S)$ is also $(S'_b \cup S''_b \cup S_c \cup S_f)$-stable.

For the following lemmas, we recall the notation:

$$\rho(i) = \{\rho_m(i), \rho_d(i)\} = \{\{(m_{a_i(1)}, d_{b_i(1)}), \ldots, (m_{a_i(r_i)}, d_{b_i(r_i)})\},$$

$$\{(m_{a_i(1)}, d_{b_i(2)}), \ldots, (m_{a_i(r_i-1)}, d_{b_i(r)}), (m_{a_i(r_i)}, d_{b_i(1)})\}\}.$$

In addition, for all $i, j \in [k]$ such that $\hat{1} \not\leq^s p_i, p_j \not\leq^s \hat{0}$, $m_{y(i)} \in \rho(i)$, $d_{x(j)} \in \rho(j)$, and $m_{y(i)}$ and $d_{x(j)}$ add one another to their preference lists in the ith iteration of step 4b iff one of the following holds:

- $j < i$, $p_j \not\leq^c p_i$, and p_i covers p_j in $(P, \leq^s)$.
- $j < i$, $p_j \not\leq^h p_i$, and p_i covers p_j in $(P, \leq^c)$.

Lemma 5.80. *Let $i, j \in [k]$ such that $m_{y(i)}$ and $d_{x(j)}$ add one another to their preference lists in step 4b of Algorithm 5.66. Then, $m_{y(i)} \in \rho(i)$ and $d_{x(j)} \in \rho(j)$. Furthermore, $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$ to $p_{\rho_d(i)}(m_{y(i)})$, and $d_{x(j)}$ prefers $p_{\rho_d(j)}(d_{x(j)})$ to $m_{y(i)}$ to $p_{\rho_m(j)}(d_{x(j)})$.*

Proof. Since $y(i)$ and $x(j)$ are incident with p_i and p_j respectively, $y(i) = a_i(s)$ and $x(j) = b_j(t)$ for some $s \in r_i$, $t \in r_j$. We note the following:

- $m_{a_i(s)}$ has $d_{b_i(s)} = p_{\rho_m(i)}(m_{y(i)})$ as the last element of V'_d on their preference list before the ith iteration of step 4a (by the definition of $b_i(s)$ given in Algorithm 5.66).

- $m_{a_i(s)}$ adds $d_{b_i(s+1)} = p_{\rho_d(i)}(m_{y(i)})$ (with $s+1$ taken $\mod r_i$) to the bottom of their preference list during the ith iteration of step 4a.

- $m_{a_i(s)}$ adds $d_{x(j)}$ second from the bottom of their preference list in the ith iteration of step 4b.

Consequentially, $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$ to $p_{\rho_d(i)}(m_{y(i)})$.

Similarly, we note that:

- $d_{b_j(t)}$ has $m_{a_j(t)} = p_{\rho_m(j)}(d_{x(j)})$ as the last element on their preference list before the jth iteration of step 4a (by Proposition 5.39.

- $d_{b_j(t)}$ adds $m_{a_j(t-1)} = p_{\rho_d(j)}(d_{x(j)})$ (with $s-1$ taken $\mod r_j$) to the top of their preference list during the jth iteration of step 4a.

- $d_{b_j(t)}$ adds $m_{y(i)}$ second from the top of their preference list in the ith iteration of step 4b.

By Proposition 5.72, the top element of $d_{x(j)}$'s preference list does not change after the ith iteration of step 4a, so $d_{x(j)}$ prefers $p_{\rho_d(j)}(d_{x(j)})$ to $m_{y(i)}$ to $p_{\rho_m(j)}(d_{x(j)})$. $\square$

Lemma 5.81. *For all $i \in \{0, \dots, k\}$, $\gamma(\{p_0, \dots, p_i\})$ matches each mentor in V'_m (resp. student in V'_d) with their bottom choice in V'_d (resp. their top choice) after the ith iteration of step 4.*

Proof. We prove this result by induction. For the base case, when $i = 0$, $\gamma(\{p_0\}) = \{(m_e, d_e) : e \in E \cup \{\tau\}\}$ matches each mentor in V'_m (resp. student in V'_d) with their bottom choice in V'_d (resp. their top choice) after step 3. The only steps that can occur nonvacuously during the 0th iteration of step 4 are 4c, 4d, 4e, and 4f; however, as seen in the proof of Theorem 5.69, neither step can change the bottom (resp. top) element of V'_d (resp. V_m) on the preference list of any mentor in V'_m (resp. student in V'_d), so $\gamma(\{p_0,\})$ matches each mentor in V'_m (resp. student in V'_d) with their bottom choice in V'_d (resp. their top choice) after the 0th iteration of step 4.

To show our inductive step, assume that for any $i \in [k]$, $\gamma(\{p_0, \dots, p_{i-1}\})$ matches each mentor in V'_m (resp. student in V'_d) with their bottom choice in V'_d (resp. their top choice) after the $(i-1)$th iteration of step 4. The ith iteration of step 4a makes the following changes to preference lists:

- For all $m \in V'_m$ such that $m \in \rho(i)$, $p_{\rho_d(i)}(m) \in V'_d$ is added to the bottom of m's preference list. (Previously, by the inductive assumption, the last element of V'_d on m's preference list was $p_{\gamma(\{p_0,\dots,p_{i-1}\}}(m) = p_{\rho_m(i)}(m)$.)

- For all $d \in V'_m$ such that $d \in \rho(i)$, $p_{\rho_d(i)}(d)$ is added to the top of d's preference list. (Previously, by the inductive assumption, the first element on m's preference list was $p_{\gamma(\{p_0,\ldots,p_{i-1}\}}(d) = p_{\rho_m(i)}(d)$.)

As a result, we see that there exists a perfect matching over I that matches each mentor in V'_m (resp. student in V'_d) with their bottom choice in V'_d (resp. their top choice) after the ith iteration of step 4a, and this matching is $\gamma(\{p_0,\ldots,p_{i-1}\}) \cup \rho_d(i) - \rho_m(i) = \{(m_e, d_e) : e \in E \cup \{\tau, \sigma, \sigma', \sigma'', \upsilon, \upsilon', \upsilon''\}\} \cup (\cup_{j=1}^{i} \rho_d(j)) - (\cup_{j=1}^{i} \rho_m(j)) = \gamma(\{p_0,\ldots,p_i\})$. As seen in the proof of Theorem 5.69, steps 4b through 4f do not change the bottom element of V'_d (resp. top element) in the preference list of any $m \in V'_m$ (resp. $d \in V'_d$), so we see that $\gamma(\{p_0,\ldots,p_i\})$ matches each mentor in V'_m (resp. student in V'_d) with their bottom choice in V'_d (resp. their top choice) after the ith iteration of step 4. By indiction, we are done. □

The proofs of Lemma 5.80 and Lemma 5.81 are analogous to the proofs of Lemma 5.49 and Lemma 5.51 respectively. The stable matchings over $I[G_h]$ are perfect matchings, and retain this property over the larger instance I.

We can now prove Lemma 5.79.

Proof. Suppose $S \in \mathcal{D}(P, \leq^h)$. As noted by Lemma 5.78, the matching $\gamma(S)$ is G_h-stable. We consider whether S is S_b-stable, S_δ-stable, and S_ϵ-stable.

- We note that S_b is the set of all edges of the form $e = (m_{y(i)}, d_{x(j)})$, where $p_i, p_j \in P - P_0^* - P_1^*$, $i > j$, $p_i \not\geq^h p_j$, and p_i covers p_j in $(P, \leq^s)$. By Lemma 5.80, $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$ to $p_{\rho_d(i)}(m_{y(i)})$, and $d_{x(j)}$ prefers $p_{\rho_d(j)}(d_{x(j)})$ to $m_{y(i)}$ to $p_{\rho_m(j)}(d_{x(j)})$; therefore, by the definition of γ, d_a prefers $p_{\gamma(S)}(d_a)$ to m_b iff S contains p_i, and m_b prefers $p_{\gamma(S)}(m_b)$ to d_a iff S does not contain p_j. Therefore, $\gamma(S)$ is $\{e\}$-stable iff $p_i \in S \Rightarrow p_j \in S$, and $\gamma(S)$ is S_b-stable iff $p_i \in S \Rightarrow p_j \in S$ for all $i, j \in [k]$ such that $i < j$, $p_i \not\leq^h p_j$, and $p_i \leq^s p_j$.

- Recall that $S_\delta = \{(m_\tau, d) : d \neq d_\tau\}$. Since $(m_\tau, d_\tau) \in \gamma(S)$ for all $S \in \mathcal{D}(P, \leq^h)$, and m_τ prefers any other possible partner to d_τ, $\gamma(S)$ is S_δ-stable iff every student other than d_τ prefers their partner in $\gamma(S)$ to m_τ. Since m_τ is added to each student's preference list second from the top at the ith iteration of step 4d (where i is the greatest index such that $p_i \leq^s \hat{0}$), this occurs iff each such student weakly prefers their partner in $\gamma(S)$ to their top choice at that point. By Lemma 5.81, that top choice is their partner in $\gamma(T_m)$, where $T_m = \{p \in P : p \leq^s \hat{0}\}$ - and every student prefers their partner in $\gamma(S)$ iff $S \supseteq T_m$.

- Recall that $S_\epsilon = \{(m, d_\tau) : m \neq m_\tau\}$. Since $(m_\tau, d_\tau) \in \gamma(S)$ for all $S \in \mathcal{D}(P, \leq^h)$, and d_τ prefers any other possible partner to m_τ, $\gamma(S)$ is S_ϵ-stable iff every mentor other than m_τ prefers their partner in $\gamma(S)$ to d_τ. Since d_τ is added to each mentor's preference list at the bottom during the ith iteration of step 4e (where i is the greatest index such that $p_i \not\geq^s \hat{1}$), this occurs iff each such mentor weakly prefers their partner in $\gamma(S)$ to their bottom choice in V'_d at that point. By Lemma 5.81, that bottom choice is their partner in $\gamma(T_d)$, where $T_d = \{p \in P : p \not\geq^s \hat{1}\}$ - and every mentor prefers their partner in $\gamma(S)$ iff $S \subseteq T_d$.

Thus, we see the following conditions are necessary for S to be $\{e\}$-stable for every $e \in G(I)$:

- $p_i \in S \Rightarrow p_j \in S$ for all $i, j \in [k]$ such that $i < j$, $p_i \not\leq^h p_j$, and $p_i \leq^s p_j$.
- $p_i \in S$ for all $p_i \leq^s \hat{0}$.
- $p_i \notin S$ for all $p_i \geq^s \hat{1}$.

We note that this is the list of necessary and sufficient conditions for any $S \subseteq P$ to be in $\mathcal{D}(P, \leq^s)$. To show that the lemma is true, we need to show that the above list of conditions is also sufficient for S to be $\{e\}$-stable for every $e \in G(I)$ - i.e. if $S \in \mathcal{D}(P, \leq^s)$, then $\gamma(S)$ is $\{e\}$-stable for every $e \in S'_b \cup S''_b \cup S_c \cup S_f$.

- Suppose $e \in S'_b$, so there exist $p_j, p_i \in P^*_0 - P_0$ such that $e = (m_{y(i)}, d_{x(j)})$. Then, by Lemma 5.80, $d_e = d_{x(j)}$ prefers their partner in $\rho_d(j)$ to $m_{y(i)} = m_e$. Since $p_j \in S$ for any $S \in \mathcal{D}(P, \leq^s)$, d_e prefers their partner in $\gamma(S)$ to their partner in $\rho_d(j)$, and so prefers $p_{\gamma(S)}(d_e)$ to m_e - thereby showing that $\gamma(S)$ is $\{e\}$-stable.

- Suppose $e \in S_b''$, so there exist $p_j, p_i \in P_1^* - P_1$ such that $e = (m_{y(i)}, d_{x(j)})$. Then, by Lemma 5.80, $m_e = m_{y(i)}$ prefers their partner in $\rho_m(i)$ to $d_{x(j)} = d_e$. Since $p_i \notin S$ for any $S \in \mathcal{D}(P, \leq^s)$, m_e prefers their partner in $\gamma(S)$ to their partner in $\rho_m(i)$, and so prefers $p_{\gamma(S)}(m_e)$ to d_e - thereby showing that $\gamma(S)$ is $\{e\}$-stable.

- Suppose $e \in S_c$; then, either $m_e = m_{\sigma''}$ or $d_e = d_{\sigma'}$. If $m_e = m_{\sigma''}$, then d_e prefers m_τ to m_e by Proposition 5.68 (since m_e and m_τ are added to d_e's preference list second from the top); however, d_e also prefers $p_{\gamma(S)}(d_e)$ to m_τ - thereby proving that $\gamma(S)$ is $\{e\}$-stable. On the other hand, $d_e = d_{\sigma'}$, then d_e prefers $p_{\gamma(S)}(d_e) = m_\sigma$ to m_e, because m_e was added to the bottom of d_e's preference list and m_σ was added to the top in the i_0th iteration of step 4a. As a result, we see that $\gamma(S)$ is also $\{e\}$-stable in this case.

- Suppose $e \in S_f$; then, either $d_e = d_{\upsilon''}$ or $m_e = m_\upsilon$. If $d_e = d_{\upsilon''}$, then m_e prefers d_τ to d_e by Proposition 5.68 (since m_e and m_τ are added to d_e's preference list at the bottom); however, m_e also prefers $p_{\gamma(S)}(m_e)$ to d_τ - thereby proving that $\gamma(S)$ is $\{e\}$-stable. On the other hand, if $m_e = m_\upsilon$, then m_e prefers $p_{\gamma(S)}(m_e) = d_\upsilon$ to d_e, because d_e was added to the bottom of m_e's preference list. As a result, we see that $\gamma(S)$ is also $\{e\}$-stable in this case.

$\square$

Since every stable matching in I that consists entirely of edges in G_h is also stable in $I[G_h]$, we see that γ is a bijection from $\mathcal{D}(P, \leq^s)$ to the stable matchings of I that consisit entirely of edges in G_h. This set of stable matchings, which we define as $\mathcal{S}$, is obviously closed under join and meet (as the join and meet of two stable matchings consist of edges from those matchings). However, we still need to show that there are no other stable matchings over I.

5.6.4 The S_2-Stable Matchings

In this section, we carry out step 5 in our outline of the proof of Theorem 5.65. In particular, we need to identify a set of edges $S_2 \subseteq E(G(I))$ such that the lattice of the

S_2-stable matchings over I which are $\subseteq I[G_h]$ is isomorphic to $\mathcal{D}(P, \leq^c)$ via γ. We will show that $S_2 \equiv G(I) - S_b - S_\delta - S_\epsilon$ fulfills this condition.

Proposition 5.82. *$\psi_I^\infty \subseteq S_2$.*

Proof. We prove this statement by the following two lemmas.

Lemma 5.83. *If $e \in S_b$, then $e \notin \psi_I^\infty$.*

Proof. Since $e \in S_b$, $e = (m_{y(i)}, d_{x(j)})$, where $i, j \in [k]$ such that $i > j$, $p_i, p_j \in P - P_0^* - P_1^*$, p_i and p_j are independent in $\mathcal{D}(P, \leq^h)$, and p_i covers p_j in $\mathcal{D}(P, \leq^s)$. Since p_i covers p_j in $\mathcal{D}(P, \leq^s)$, there exists some $D \in \mathcal{D}(P, \leq^s)$ such that $p_j \in D$ and $p_i \notin D$. As a result, if $M' = \gamma(D)$, then $d_{x(j)}$ prefers $p_{M'}(d_{x(j)})$ to $p_{\rho_d(i)}(d_{x(j)})$, and $m_{y(i)}$ prefers $p_{M'}(m_{y(i)})$ to $p_{\rho_m(i)}(m_{y(i)})$. However, by Lemma 5.80, $d_{x(j)}$ also prefers $p_{\rho_d(i)}(d_{x(j)})$ to $m_{y(i)}$, and $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$.

As a result, if M is any matching that includes e, then $m_e = m_{y(i)}$ and $d_e = d_{x(j)}$ each prefer their partner in M' to their partner in M, so M and M' are not costable. However, by Lemma 5.79, M' is stable - and therefore hub-stable - so M cannot be hub-stable. Since M is any arbitrary matching that includes e, $e \notin \psi_I^\infty$. $\square$

Lemma 5.84. *If $e \in S_\delta \cup S_\epsilon$, then, $e \notin \psi_I^\infty$.*

Proof. Assume, for the sake of contradiction, that the lemma is false; then, there exists a hub-stable matching M^* such that m_τ and d_τ are not matched with each other. Let M_0 be any stable matching over I that includes (m_τ, d_τ) as an edge - we know such a matching exists by Lemma 5.79. Since M^* and M_0 are hub-stable, they must be costable as well. However, m_τ and d_τ are partnered in M_0, and both prefer their respective partners in M^* to each other; this creates a contradiction, and so no such M^* can exist. $\square$

By Lemma 5.83 and Lemma 5.84, every $e \in S_b \cup S_\delta \cup S_\epsilon$ is $\notin \psi_I^\infty$; as a result, every $e \in G(I)$ that is $\notin S_2$ is also $\notin \psi_I^\infty$. $\square$

Recall that, for all $i, j \in [k]$ such that $p_i, p_j \leq^s \hat{0}$ and $\not\leq^c \hat{0}$, $m_{y(i)} \in \rho(i)$, $d_{x(j)} \in \rho(j)$, and $m_{y(i)}$ and $d_{x(j)}$ add one another to their preference lists in the ith iteration of step

4c iff $j < i$, $p_j \not\leq^h p_i$, and p_i covers p_j in $(P, \leq^c)$. Similarly, for all $i, j \in [k]$ such that $p_i, p_j \geq^s \hat{1}$ and $\not\geq^c \hat{1}$, $m_{y(i)} \in \rho(i)$, $d_{x(j)} \in \rho(j)$, and $m_{y(i)}$ and $d_{x(j)}$ add one another to their preference lists in the ith iteration of step 4d iff $j < i$, $p_j \not\leq^h p_i$, and p_i covers p_j in $(P, \leq^c)$.

Lemma 5.85. *Let $S \in \mathcal{D}(P, \leq^h)$. Then, $\gamma(S)$ is S_2-stable iff $S \in \mathcal{D}(P, \leq^c)$.*

Proof. Suppose $S \in \mathcal{D}(P, \leq^h)$. As noted by Lemma 5.78, the matching $\gamma(S)$ is G_h-stable. Since $S_2 = G_h \cup S_b' \cup S_b'' \cup S_c \cup S_f$, we consider whether $\gamma(S)$ is S_b'-stable, S_b''-stable, S_c-stable, and S_f-stable.

- We note that $S_b' \cup S_b''$ is the set of all edges of the form $e = (m_{y(i)}, d_{x(j)})$, where $i, j \in [k]$ uphold one of the following:

  - $p_i, p_j \in P_0^* - P_0$, $i > j$, $p_i \not\geq^h p_j$, and p_i covers p_j in $(P, \leq^c)$.
  - $p_i, p_j \in P_1^* - P_1$, $i > j$, $p_i \not\geq^h p_j$, and p_i covers p_j in $(P, \leq^c)$.

  By Lemma 5.80, $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$ to $p_{\rho_d(i)}(m_{y(i)})$, and $d_{x(j)}$ prefers $p_{\rho_d(j)}(d_{x(j)})$ to $m_{y(i)}$ to $p_{\rho_m(j)}(d_{x(j)})$; therefore, by the definition of γ, d_a prefers $p_{\gamma(S)}(d_a)$ to m_b iff S contains p_i, and m_b prefers $p_{\gamma(S)}(m_b)$ to d_a iff S does not contain p_j. Therefore, $\gamma(S)$ is $\{e\}$-stable iff $p_i \in S \Rightarrow p_j \in S$, and $\gamma(S)$ is S_b-stable iff $p_i \in S \Rightarrow p_j \in S$ for all $i, j \in [k]$ such that $i < j$, $p_i \not\leq^h p_j$, and $p_i \leq^c p_j$.

- If $e \in S_c$, then there are two possibilities on the structure of e:

  - $m_e = m_{\sigma''}$. We note that $\gamma(S)$ is stable for every such e iff, for all $e' \in E_0$, $d_{e'}$ prefers their partner in $\gamma(S)$ to $m_{\sigma''}$. Since $m_{\sigma''}$ is added to each such student's preference list second from the top at the $(i_0 - 1)$th iteration of step 4c, this occurs iff each such student weakly prefers their partner in $\gamma(S)$ to their top choice at that point. By Lemma 5.81, that top choice is their partner in $\gamma(T_m')$, where $T_m' = \{p \in P : p \leq^c \hat{0}\}$ - and every such student prefers their partner in $\gamma(S)$ iff $S \supseteq T_m'$.

- $d_e = d_{\sigma'}$. If so, we note that $p_{\gamma(S)}(d_e) = m_{\sigma'}$ or m_σ (since $\rho(i_0)$ is the only rotation over $I[G_h]$ that d_e appears in), and d_e prefers these mentors to any mentors added to their preference list in step 4c (since $m_{\sigma'}$ was added to their preference list first, m_σ to the top, and every element from step 4c to the bottom). Thus, $\gamma(S)$ is e-stable for any possible $S \in \mathcal{D}(P, \leq^h)$.

Thus, we see that $\gamma(S)$ is S_c-stable iff $S \supseteq P_0$.

- If $e \in S_f$, then there are two possibilities on the structure of e:

  - $d_e = d_{\upsilon''}$. We note that $\gamma(S)$ is stable for every such e iff, for all $e' \in E_1$, $m_{e'}$ prefers their partner in $\gamma(S)$ to $d_{\upsilon''}$. Since $d_{\upsilon''}$ is added to each such mentor's preference list at the bottom during the i_1th iteration of step 4f, this occurs iff each such mentor weakly prefers their partner in $\gamma(S)$ to their bottom choice in V'_d prior to that point. By Lemma 5.81, that bottom choice is their partner in $\gamma(T'_d)$, where $T'_d = \{p \in P : p \not\geq^c \hat{1}\}$ - and every such mentor prefers their partner in $\gamma(S)$ iff $S \subseteq T'_d$.
  - $m_e = m_\upsilon$. If so, we note that $p_{\gamma(S)}(m_e) = d_\upsilon$ or $d_{\upsilon'}$ (since $\rho(i_1)$ is the only rotation over $I[G_h]$ that m_e appears in), and m_e prefers these students to any students added to their preference list in step 4f (since each such student is added to the bottom of their preference list after step 3 and the i_1th iteration of step 4a). Thus, $\gamma(S)$ is e-stable for any possible $S \in \mathcal{D}(P, \leq^h)$.

  Thus, we see that $\gamma(S)$ is S_f-stable iff $S \subseteq P - P_1$.

As a result, we see that $\gamma(S)$ is $\{e\}$-stable for every $e \in S_2$ iff:

- For all $p_i, p_j \in P$ such that $p_i \not\geq^h p_j$ and p_i covers p_j in $(P, \leq^c)$, $p_i \in S \Rightarrow p_j \in S$.
- $S \supseteq P_0 = \{p \in P : p \leq^c \hat{0}\}$.
- $S \subseteq P - P_1 = \{p \in P : p \geq^c \hat{1}\}$.

However, this list of conditions is precisely the list of conditions for elements of $\mathcal{D}(P, \leq^h)$ to be elements of $\mathcal{D}(P, \leq^c)$. $\square$

5.6.5 The Structure of $\mathcal{L}_h(I)$, $\mathcal{L}_s(I)$, and $\mathcal{L}_c(I)$

In this subsection, we carry out steps 6 and 7 of the outline of the proof of Algorithm 5.66, and tie together the results of the corresponding lemmas. The pivotal step of this stage of the proof is showing that G_h is the unique hub ψ_I^∞; we do this by showing that $G_h \subseteq \psi_I(G_h)$ and $\supseteq \psi_I(G_h)$. In order to show the latter statement holds, we show that $G_h \supseteq \psi_I^\infty$. We have already shown in Proposition 5.82 that $\psi_I^\infty \subseteq S_2$; the next three lemmas similarly observe that the elements of $S_2 - G_h$ are also $\notin \psi_I^\infty$.

Lemma 5.86. *If $e \in S_b' \cup S_b''$, then $e \notin \psi_I^\infty$.*

Proof. Since $e \in S_b' \cup S_b''$, $e = (m_{y(i)}, d_{x(j)})$, where $i, j \in [k]$ such that $i > j$, $p_i, p_j \in (P_0^* - P_0) \cup (P_1^* - P_1)$, p_i and p_j are independent in $\mathcal{D}(P, \leq^h)$, and p_i covers p_j in $\mathcal{D}(P, \leq^c)$. Since p_i covers p_j in $\mathcal{D}(P, \leq^c)$, there exists some $D \in \mathcal{D}(P, \leq^c)$ such that $p_j \in D$ and $p_i \notin D$. As a result, if $M' = \gamma(D)$, then $d_{x(j)}$ prefers $p_{M'}(d_{x(j)})$ to $p_{\rho_d(i)}(d_{x(j)})$, and $m_{y(i)}$ prefers $p_{M'}(m_{y(i)})$ to $p_{\rho_m(i)}(m_{y(i)})$. However, by Lemma 5.80, $d_{x(j)}$ also prefers $p_{\rho_d(i)}(d_{x(j)})$ to $m_{y(i)}$, and $m_{y(i)}$ prefers $p_{\rho_m(i)}(m_{y(i)})$ to $d_{x(j)}$.

As a result, if M is any matching that includes e, then $m_e = m_{y(i)}$ and $d_e = d_{x(j)}$ each prefer their partner in M' to their partner in M, so M and M' are not costable. However, by Lemma 5.85, $M' = \gamma(D)$ is S_2-stable, and therefore hub-stable byproposition 5.82. As a result, M cannot be hub-stable; since M is any arbitrary matching that includes e, $e \notin \psi_I^\infty$. □

Lemma 5.87. *If $e_0 \in S_c$, then, $e_0 \notin \psi_I^\infty$.*

Proof. Let M be any hub-stable matching over I; then, by proposition 5.82, $\gamma(P_0)$ is also hub-stable, and so M and $\gamma(P_0)$ are costable. Since $\gamma(P_0)$ matches $m_{\sigma'}$ with their top choice $d_{\sigma'}$, $m_{\sigma'}$ prefers $d_{\sigma'}$, their partner in $\gamma(P_0)$, to that in M. Since M and $\gamma(P_0)$ are costable, $d_{\sigma'}$ must prefer their partner in M to $m_{\sigma'}$ - which implies that $p_M(d_{\sigma'}) = m_{\sigma'}$ or $m_{\sigma''}$.

Now, assume for the sake of contradiction that $p_M(m_{\sigma''}) = d_e$ for some $e \in E_0$. Let $M^* = M \wedge \gamma(P_0)$; since $m_{\sigma''}$ prefers d_e to $d_{\sigma''}$, $p_{M^*}(m_{\sigma''}) = d_e$. As a result, there exists some $e' \in E_0$ such that $p_{M^*}(m_{e'}) \notin \{d_{e''} : e'' \in E_0\}$. However, this implies that $m_{e'}$ prefers their partner in $\gamma(P_0)$ to their partner in M^* - creating a contradiction.

Since $m_{e_0} = m_{\sigma''}$ or $d_{e_0} = d_{\sigma'}$ for any $e_0 \in S_c$, we see that no such e_0 can appear in a hub-stable matching, and so $e_0 \notin \psi_I^\infty$. $\square$

Lemma 5.88. *If $e_0 \in S_f$, then, $e_0 \notin \psi_I^\infty$.*

Proof. Let M be any hub-stable matching over I; then, by proposition 5.82, $\gamma(P_0)$ is also hub-stable, and so M and $\gamma(P - P_1)$ are costable. Since $\gamma(P-P_1)$ matches $d_{\upsilon'}$ with their top choice m_υ, $d_{\upsilon'}$ prefers m_υ, their partner in $\gamma(P - P_1)$, to that in M. Since M and $\gamma(P_0)$ are costable, m_υ must prefer their partner in M to $d_{\upsilon'}$ - which implies that $p_M(m_\upsilon) = d_\upsilon$ or $d_{\upsilon'}$.

Now, assume for the sake of contradiction that $p_M(d_{\upsilon''}) = m_e$ for some $e \in E_1$. Let $M^* = M \vee \gamma(P - P_1)$; since $d_{\upsilon''}$ prefers m_e to $m_{\upsilon'}$, $p_{M^*}(m_{\sigma''}) = d_e$. As a result, there exists some $e' \in E_1'$ such that $p_{M^*}(d_{e'}) \notin \{m_{e''} : e'' \in E_1\}$. However, this implies that $m_{e'}$ prefers their partner in $\gamma(P_0)$ to their partner in M^* - creating a contradiction.

Since $m_{e_0} = m_\upsilon$ or $d_{e_0} = d_{\upsilon''}$ for any $e_0 \in S_f$, we see that no such e_0 can appear in a hub-stable matching, and so $e_0 \notin \psi_I^\infty$. $\square$

Lemma 5.89. *$\psi_I^\infty = G_h$.*

Proof. By Theorem 5.69, every edge in G_h appears in a stable matching over $I[G_h]$. Such a matching is still G_h-stable over I, so every edge in G_h is in $\psi_I(G_h)$ - implying that $\psi_I(G_h) \supseteq G_h$.

Conversely, we note that by Lemma 5.83, Lemma 5.84, Lemma 5.86, Lemma 5.87, and Lemma 5.88, every edge in $E(G(I)) - G_h$ is also $\notin \psi_I^\infty$, so $\psi_I^\infty \subseteq G_h$. By Proposition 4.3, this implies that $\psi_I(G_h) \subseteq \psi_I(\psi_I^\infty) = \psi_I^\infty$, so $\psi_I(G_h) \subseteq G_h$. As a result, $\psi_I(G_h) = G_h$, so G_h is the unique hub of I. $\square$

Corollary 5.90. *The set of hub-stable matchings over I is the set of stable matchings over $I[G_h]$.*

Since every stable matching is hub-stable, every stable matchings over I appears in $\mathcal{S}$, as defined in Lemma 5.79. Consequentially, $\mathcal{L}_h(I)$ and $\mathcal{L}_s(I)$ have the desired structure. We only have to show that $\mathcal{L}_c(I)$ has the desired structure.

Lemma 5.91. *The set of* 3*-stable matchings over I is $\gamma(\mathcal{D}(P, \leq^c))$.*

Proof. By Lemma 5.79, the set of stable matchings over I is $\gamma(\mathcal{D}(P, \leq^s))$. As seen in the proof of Lemma 4.40, $\psi_I^3(\emptyset)$ is the union of every perfect matching that is costable with every stable matching. By Lemma 5.83 and Lemma 5.84, no edge in S_b, S_δ, or S_ϵ appears in a matching that is costable with every (or any) stable matching. However, every edge in S_2 appears in some such matching, so $S_2 = \psi_I^3(\emptyset)$. (For the following proof, for any $e \in E$, $d_e^* = p_{\gamma(P)}(m_e)$.)

Lemma 5.92. *Every edge in $e \in S_b'$, S_b'', S_c, or S_f appears in some matching M_e that is costable with every stable matching.*

Proof. If $e \in S_b'$, then there exists some $e' \in E_0$ such that $e = (m_{\sigma''}, d_{e'})$ or $(m_{e'}, d_{\sigma'})$. Either way, let $M_e = \{(m_{\sigma''}, d_{e'}), (m_{e'}, d_{\sigma'}), (m_{\sigma'}, d_{\sigma''}) \cup \{(m_a, d_a) : a \notin \{e', \sigma', \sigma''\}$. We note that every mentor prefers their partner in M_e to their partner in $\gamma(P_0^*)$, the mentor-optimal stable matching, and every student prefers their partner in $\gamma(P_0^*)$ to their partner in M_e. Therefore, M_e is costable with every stable matching.

If $e \in S_c$, then $e = (m_{e'}, d_{e''})$ for some $e', e'' \in E_0$. We can set:

$$M_e = \{(m_{\sigma''}, d_{e'}), (m_{e'}, d_{e''}), (m_{e''}, d_{\sigma'}), (m_{\sigma'}, d_{\sigma''})\} \cup \{(m_a, d_a) : a \notin \{e', e'', \sigma', \sigma''\}\}$$

We note that every mentor prefers their partner in M_e to their partner in $\gamma(P_0^*)$, the mentor-optimal stable matching, and every student prefers their partner in $\gamma(P_0^*)$ to their partner in M_e. Therefore, M_e is costable with every stable matching.

If $e \in S_b''$, then there exists some $e' \in E_1$ such that $e = (m_\upsilon, d_{e'}^*)$ or $(m_{e'}, d_{\upsilon''})$. Either way, let $M_e = \{(m_\upsilon, d_{e'}^*), (m_{e'}, d_{\upsilon''}), (m_{\upsilon'}, d_{\upsilon'}) \cup \{(m_a, d_a^*) : a \notin \{e', \upsilon, \upsilon'\}$. (This is a perfect matching, because $d_{\upsilon'} = d_\upsilon^*$ and $d_{\upsilon''} = d_{\upsilon'}^*$.) We note that every student prefers their partner in M_e to their partner in $\gamma(P - P_1^*)$, the student-optimal stable matching, and every mentor prefers their partner in $\gamma(P - P_1^*)$ to their partner in M_e. Therefore, M_e is costable with every stable matching.

If $e \in S_f$, then $e = (m_{e'}, d^*_{e''})$ for some $e', e'' \in E_1$. We can set:

$$M_e = \{(m_\upsilon, d^*_{e'}), (m_{e'}, d^*_{e''}), (m_{e''}, d_{\upsilon''}), (m_{\upsilon'}, d_{\upsilon'})\} \cup \{(m_a, d^*_a) : a \notin \{e', e'', \upsilon, \upsilon'\}\}$$

(Note that M_e is a perfect matching, because $d_{\upsilon'} = d^*_\upsilon$ and $d_{\upsilon''} = d^*_{\upsilon'}$.) We note that every student prefers their partner in M_e to their partner in $\gamma(P - P^*_1)$, the student-optimal stable matching, and every mentor prefers their partner in $\gamma(P - P^*_1)$ to their partner in M_e. Therefore, M_e is costable with every stable matching. $\square$

By Lemma 5.85, the set of $\psi^3_I(\emptyset)$-stable matchings that are $\subseteq G_h$ is $\gamma(\mathcal{D}(P, \leq^c))$. However, since the union of all $\psi^3_I(\emptyset)$-stable matchings is $\psi^4_I(\emptyset) \subseteq \psi^\infty_I$, this implies that the set of all $\psi^3_I(\emptyset)$-stable matchings is $\gamma(\mathcal{D}(P, \leq^c))$. $\square$

We are now ready to finish proving Theorem 5.65.

Proof. By Theorem 5.6, we may find a pointed order $(P, \leq^h)$ and nested separated extensions $(P, \leq^c)$ and $(P, \leq^s)$ such that $(\mathcal{D}(P, \leq^s), \mathcal{D}(P, \leq^c), \mathcal{D}(P, \leq^h))$ is isomorphic to $(\mathcal{L}_s, \mathcal{L}_c, \mathcal{L}_h)$. Let I be the instance created by Algorithm 5.66 given $(P, \leq^h)$, $(P, \leq^c)$, and $(P, \leq^s)$. By Corollary 5.90 and Lemma 5.78, the lattice of hub-stable matchings over I is isomorphic to $\mathcal{L}_h$, and the bijection γ maps $\mathcal{D}(P, \leq^h)$ to the set of all hub-stable matchings over I. Furthermore, by Lemma 5.79, γ also maps $\mathcal{D}(P, \leq^s)$ to the set of all stable matchings over I that are $\subseteq G_h$; however, every stable matching is hub-stable, and every hub-stable matching is $\subseteq G_h$ by Lemma 5.89, so γ maps $\mathcal{D}(P, \leq^s)$ to the set of all stable matchings over I. Finally, by Lemma 5.91, γ also maps $\mathcal{D}(P, \leq^c)$ to the set of all 3-stable matchings over I. Therefore, $(\mathcal{L}_s(I), \mathcal{L}_c(I), \mathcal{L}_h(I))$ is isomorphic to $(\mathcal{D}(P, \leq^s), \mathcal{D}(P, \leq^c), \mathcal{D}(P, \leq^h))$ - which is isomorphic to $(\mathcal{L}_s, \mathcal{L}_c, \mathcal{L}_h)$. $\square$

5.6.6 Generalizing Algorithm 5.66 to Prove Conjecture 5.64

We speculate that the procedure described in Algorithm 5.66 can be extended in order to prove Conjecture 5.64. In particular, for all $r \in [z-1] \cup \{0\}$, we define $P^{(r)}_0 = \{p \in P : p \leq^r \hat{0}\}$ and $P^{(r)}_1 = \{p \in P : p \geq^r \hat{1}\}$; for all $r \in [z-1]$, we set $i^(_0 r)$ to be the least index among the elements of $P^{(r)}_0 - P^{(r-1)}_0$, and $i^(_1 r)$ to be the greatest index among the elements of $P^{(r)}_1 - P^{(r-1)}_1$. We then alter Algorithm 5.66 such that the indices of our vertices include $\{\sigma_{(r)}, \sigma'_{(r)}, \sigma''_{(r)}, \upsilon_{(r)}, \upsilon'_{(r)}, \upsilon''_{(r)}\}$ for all $r \in [z-1]$. Under this, we alter step 4c as follows:

- $i_0 - 1$ is replaced with $i_0^{(r)} - 1$ (for some value of r)
- σ' (resp. σ'') is replaced with $\sigma'_{(r)}$ (resp. $\sigma''_{(r)}$)

Similarly, we alter step 4f as follows:

- i_1 is replaced with $i_1^{(r)}$ (for some value of r)
- D_1 is replaced with $D_1^{(r)}$
- υ (resp. υ'') is replaced with $\upsilon_{(r)}$ (resp. $\upsilon''_{(r)}$)

Lastly, step 5 becomes:

- For all $r \in [z-1]$, do the following: for every $d_{e'} \in D_1^{(r)}$, place $m_{\upsilon_{(r)}}$ at the top of $d_{e'}$'s preference list.

Chapter 6

The Structure of Fractional S-Stable Matchings

Given a bipartite graph G, we consider the set of all non-negative valued functions on the edge set of G. Any edge subset - and in particular any perfect matching - over G can be identified with $0-1$ valued functions in the natural way. The perfect matching polytope is the convex hull of all perfect matchings. It is well known that the perfect matching polytope has a nice description as the set of solutions to a finite set of linear constraints (see theorem 6.1). For an arbitrary stable matching instance I over the complete graph, the stable matching polytope is the convex hull of the set of stable matchings. This polytope also has a nice description as the solution set of a set of linear constraints (see theorem 6.2).

In this chapter, we consider the problem of extending these results to the class of S-stable matchings of an instance I, where S is an arbitrary set of edges. We make partial progress in this direction.

6.1 Known Results on the Fractional S-Stable Matchings

For this chapter, we will work under the assumption that the $n \times n$ instance I is complete (and thereby satisfactory as well). For any given $S \subseteq E(G(I))$, we define $P_S(I)$, the **polytope of S-stable matchings**, to be the convex hull of the perfect S-stable matchings; we refer to any $wt : E(G(I)) \to \mathbb{R}$ in $P_S(I)$ as a **fractional S-stable matching**. (In cases where I is implied, we shorten this notation to P_S.) The overaching goal of this chapter is to find, for general S, a family of constraints that define the polytope of fractional S-stable matchings - in particular, we would like to identify and eliminate redundant constraints when possible, and we would like each constraint to be verifiable in polynomial time. (As an example of constraints that are not verifiable in polynomial time, we know that $wt(e) = 0$ for all $e \notin \psi(S)$. However, verifying this constraint in polynomial time would require us to compute $\psi(S)$ in polynomial time - something that we do not know how to do in general presently.)

Two particularly noteworthy subcases of this question are when $S = \emptyset$, and when $S = E(G(I))$. In the former case, the polytope of S-stable matchings is the convex hull of all perfect matchings over $K_{n,n}$; in the latter, the polytope of S-stable matchings is the convex hull of the stable matchings over I. For these theorems, we define four families of linear constraints that we may use on a polytope:

Q_1: For all $(m, d) \in E(I)$, $wt(m, d) \geq 0$.

Q_2: For all $m \in V_m(I)$, $\sum_{d \in V_d(I)} wt(m, d) = 1$.

Q_3: For all $d \in V_d(I)$, $\sum_{m \in V_m(I)} wt(m, d) = 1$.

$Q_4(S)$: For all $(m, d) \in S$, $wt(m, d) + \sum_{d' <_m d} wt(m, d') + \sum_{m' <_d m} wt(m', d) \leq 1$.

Theorem 6.1. *The convex hull of all perfect matchings over $K_{n,n}$ is the polytope on the domain of functions $wt : E(G(I)) \to \mathbb{R}$ with the constraints $\{Q_1, Q_2, Q_3\}$. ([Dan63], Chapter 15, Theorem 1)*[1]

Theorem 6.2. *Given an $n \times n$ complete instance I, the convex hull of all stable matchings over I is th polytope on the domain of functions $wt : E(G(I)) \to \mathbb{R}$ with the constraints $\{Q_1, Q_2, Q_3, Q_4(E(G(I))\}$. ([VV89], Theorem 1)*

It is straightforward to see that for general $S \subseteq E(G(I))$, P_S must be constrained by every constraint in $\{Q_1, Q_2, Q_3, Q_4(S)\}$ - the vertices of this polytope, which are the S-stable matchings, are. Consequentially, it is natural to ask whether, for arbitrary S, the conditions $\{Q_1, Q_2, Q_3, Q_4(S)\}$ are sufficient to constrain P_S. This is not the case (see Example 6.11). However, we can show that this does hold for $S = \psi_I^k(\emptyset)$ for some $k \in \mathbb{N}$. (We recall that the k-stable matchings over I are the $\psi_I^k(\emptyset)$-stable matchings.)

---

[1] *Geore Dantzig attributed this theorem to Garrett Birkhoff ([Bir46]).*

Theorem 6.3. *Let I be an $n \times n$ stable matching instance, and $k \in \mathbb{N}$. Then, the polytope P_k of fractional k-stable matchings is the set of all $wt : E(G(I)) \to \mathbb{R}$ that uphold the constraints $\{Q_1, Q_2, Q_3, Q_4(\psi_I^k(\emptyset))\}$.*

The proof of Theorem 6.3 is given in Section 6.2. In Section 6.3, we conjecture at how we could find a list of sufficient conditions to constrain the S-stable matchings over I, using experimental observations.

6.2 Proof of Theorem 6.3

As seen in Chapter 4, finding a compact representation of all of the S-stable matchings over I appears to be very difficult for a general instance I and $S \subseteq E(G(I))$, and the same truth appears to hold for representing the polytope of fractional S-stable matchings. However, just as we can find the sequence $(\emptyset, \psi(\emptyset), \psi^2(\emptyset), \ldots\}$ efficiently, we can also find a compact set of necessary and sufficient constraints for the polytope of S-stable matchings when $S = \psi_I^k(\emptyset)$ - i.e. the k-stable matchings - for an arbitrary value of k. In this section, we prove Theorem 6.3. We do this by first showing that the theorem holds for even k, then showing that it holds for odd k.

Theorem 6.4. *Let S be a union of stable matchings. Then, the polytope P_S of fractional S-stable matchings for an $n \times n$ instance I is the set of all $wt : E(G(I)) \to \mathbb{R}$ that uphold the constraints $\{Q_1, Q_2, Q_3, Q_4(S)\}$.*

Proof. It is obvious that every $wt \in P_S$ obeys all of the above constraints, since its vertices do; therefore, we only need to show that any wt that obeys the above constraints is in P_S. Consider any $wt \in P_S$. In particular, since $wt \in P_\emptyset$, it upholds $\{Q_1, Q_2, Q_3\}$ by Theorem 6.1, and we can express it as a weighted average of perfect matchings $wt = \sum a_i M_i$, where $\sum a_i = 1$. Now, we can consider any stable matching $M \subseteq S$. By summing 4 for every $(m, d) \in M$, we see that $\sum_{(m,d)\in M} wt(m, d) + \sum_{(m,d)\in E_1} wt(m, d) + \sum_{(m,d)\in E_2} wt(m, d) \leq n$, where $E_1 = \{(m, d) \in E(G(I)) : d <_m p_M(m)\}$ and $E_2 = \{(m, d) \in E(G(I)) : m <_d p_M(d)\}$.

However, since M is stable, every edge in $E(G(I))$ is in at least one of M, E_1, and E_2; this implies that, for every perfect matching $M_i : E(G(I)) \to \{0, 1\}$ in $P_\emptyset$,

$\sum_{(m,d)\in M} M_i(m,d) + \sum_{(m,d)\in E_1} M_i(m,d) + \sum_{(m,d)\in E_2} M_i(m,d) \geq n$. In addition, this inequality holds with equality iff M, E_1, and E_2 do not share any element of M_i. Since M is obviously disjoint from E_1 and E_2, this occurs iff there is no edge $(m,d) \in M$ such that m and d prefer each other to their respective partners in M_i - which occurs iff M_i is M-stable.

Since $wt = \sum a_i M_i$, we have $\sum_{E_1} wt(m,d) + \sum_{E_2} wt(m,d) \geq n$; this means that $\sum_{E_1} wt(m,d) + \sum_{E_2} wt(m,d) = n$. Furthermore, the equality can only hold if $a_i = 0$ for every M_i that is not M-stable - i.e. every M_i that is not $\{e\}$-stable for some $e \in M$.

However, our choice of $M \subseteq S$ was arbitrary; since every edge in S appears in some stable matching $\subseteq S$, we see that $a_i = 0$ for every M_i that is not $\{e\}$-stable for some $e \in S$. Consequentially, the representation $wt = \sum a_i M_i$ has wt expressed as a weighted average of S-stable matchings, so $wt \in P_S$ and we are done. □

Corollary 6.5. *The polytope P_2 of fractional 2-stable matchings for an $n \times n$ instance I is the set of all $wt : E(G(I)) \to \mathbb{R}$ that uphold the constraints $\{Q_1, Q_2, Q_3, Q_4(\psi_I^2(\emptyset))\}$.*

Proof. By the definition of ψ_I, $\psi_I^2(\emptyset) = \psi_I(E(G(I)))$ is the union of all stable matchings over I. By substituting $S = \psi_I^2(\emptyset)$ in Theorem 6.4, we see that the above constraints are necessary and sufficient for P_2. □

Given that the polytope of the fractional stable matchings and the polytope of the fractional 2-stable matchings have a similar structure, it is natural to ask if the polytope of the fractional k-stable matchings has an analogous structure for all positive $k \in \mathbb{N}$. We will show a necessary and sufficient list of conditions for the polytope of the fractional k-stable matchings, when k is even. For the following, we use $Q_5(S)$ to represent the family of constraints that for all $(m,d) \notin S$, $wt(m,d) = 0$. We recall from Section 2.3 that $I[S]$ is the restriction of I to the set of edges $S \subseteq E(G(I))$.

Theorem 6.6. *For all $k \in \mathbb{N}$, the polytope P_{2k} of fractional $2k$-stable matchings for an $n \times n$ instance I is the set of all $wt : E(G(I)) \to \mathbb{R}$ that uphold the constraints $\{Q_1, Q_2, Q_3, Q_4(\psi_I^{2k}(\emptyset))\}$.*

Proof. We prove this result by induction on k. For the base case $k = 0$, the statement reduces to Theorem 6.1.

Now for the inductive step, assume that the polytope P_{2k} of fractional $2k$-stable matchings is the set of all $wt : E(G(I)) \to \mathbb{R}$ that uphold the constraints $\{Q_1, Q_2, Q_3, Q_4(\psi_I^{2k}(\emptyset))\}$. Since P_{2k} is the convex hull of $\psi_I^{2k}(\emptyset)$-stable matchings, the set of all edges $e \in E(G(I))$ such that $wt(e)$ is not identically 0 for all $wt \in P_{2k}$ is $\psi_I^{2k+1}(\emptyset)$. Consequentially, for all $wt \in P_{2k}$, $wt(m, d) = 0$ if $(m, d) \notin \psi_I^{2k+1}(\emptyset)$.

Now, consider the restriction $I' = I[\psi_I^{2k+1}(\emptyset)]$. By Corollary 6.5, we note that the set of fractional $\psi_{I'}^2(\emptyset)$-stable matchings over I' is the set of all $wt : E(G(I)) \to \mathbb{R}$ that uphold the constraints $\{Q_1, Q_2, Q_3, Q_4(\psi_{I'}^2(\emptyset)), Q_5(E(G(I')))\}$. For the next step, we need the following proposition, which is a generalization of Proposition 4.59.

Proposition 6.7. *Let I be any instance, and $I_b = I[\psi_I^{2b+1}(\emptyset)]$ for some positive $b \in \mathbb{N}$. Then, for all positive $k \in \mathbb{N}$, $\psi_{I_b}^k(\emptyset) = \psi_I^{k+2b}(\emptyset)$.*

Proof. We prove this result by induction on b. For the base case, when $b = 1$, the statement is a restatement of Proposition 4.59.

For our inductive step, suppose that for all positive $k \in \mathbb{N}$, $\psi_{I_b}^k(\emptyset) = \psi_I^{k+2b}(\emptyset)$. We now consider $I_{b+1} = I[\psi_I^{2b+3}(\emptyset)]$. By our inductive assumption, $\psi_I^{2b+3}(\emptyset) = \psi_{I_b}^2(\emptyset)$; since I_b is a restriction of I and $\psi_{I_b}^2(\emptyset) \subseteq E(G(I))$, $I_b[\psi_{I_b}^2(\emptyset)] = I[\psi_{I_b}^2(\emptyset)] = I_{b+1}$. As a result, by Proposition 4.59, for all positive $k \in \mathbb{N}$, $\psi_{I_{b+1}}^k(\emptyset) = \psi_{I_b}^{k+2}(\emptyset)$, which equals $\psi_I^{k+2b+2}(\emptyset) = \psi_I^{k+2(b+1)}(\emptyset)$. Thus, we have proven our inductive step, and by induction, we are done. □

Now, by Proposition 6.7, the set of all $\psi_{I'}^2(\emptyset)$-stable matchings is the set of all $\psi_I^{2k+2}(\emptyset)$-stable matchings, so $\{Q_1, Q_2, Q_3, , Q_4(\psi_{I'}^2(\emptyset)), Q_5(E(G(I')))\}$ is the set of constraints for the polytope P_{2k+2}. However, as noted above, a subset of these constraints are sufficient to enforce that $wt(m, d) = 0$ for all $(m, d) \notin E(G(I'))$, so the condition that $wt(m, d) = 0$ for all $(m, d) \notin E(G(I'))$ is redundant. In addition, by Proposition 6.7, $\psi_{I'}^2(\emptyset) = \psi_I^{2k+2}(\emptyset)$, so we see that the necessary and sufficient conditions of P_{2k+2} are $\{Q_1, Q_2, Q_3, Q_4(\psi_I^{2k+2}(\emptyset))\}$.

Thus, we have shown the necessary inductive step, and by induction, we are done.

□

Corollary 6.8. *The polytope P_h of fractional hub-stable matchings for an $n \times n$ instance I is the set of all $wt : E(G(I)) \to \mathbb{R}$ that uphold the constraints $\{Q_1, Q_2, Q_3, Q_4(\psi_I^\infty)\}$.*

Proof. We note that by Theorem 4.51, $\psi_I^{2n}(\emptyset) = \psi_I^\infty$. Since the hub-stable matchings are the ψ_I^∞-stable matchings over I, the above set of constraints is necessary and sufficient to describe the polytope of fractional hub-stable matchings by Theorem 6.6.

□

We also note that the result of Corollary 6.8 can be extended to the convex hulls of S-stable matchings for any $S \supseteq \psi_I^\infty$.

Theorem 6.9. *Let S be set of edges such that $\psi_I^\infty \subseteq S \subseteq E(G(I))$, and P_S be the polytope of fractional S-stable matchings. Then, P_s is the set of all $wt : E(G(I)) \to \mathbb{R}$ that uphold the constraints $\{Q_1, Q_2, Q_3, Q_4(S)\}$.*

Proof. We note that the above list of conditions is a superset of the conditions on P_h from Corollary 6.8; consequentially, $P_S \subseteq P_h$. As a result, for every $wt \in P_S$ and $(m, d) \in E(G(I)) - \psi_I^\infty$, $wt(m, d) = 0$. In particular this is true for all $(m, d) \notin S$, so for every $wt \in P_S$, $wt(m, d) \equiv 0$ for every $e \notin S$. As a result, by Theorem 6.2, the above constraints describe the polytope of the fractional stable matchings of $I[S]$. However, since $S \supseteq \psi_I^\infty \supseteq \psi(S)$, the S-stable matchings over I are precisely the stable matchings over $I[S]$ by Theorem 3.10; consequentially, their convex hulls are the same, and so the polytope of fractional S-stable matchings is also the set of all wt that uphold the above constraints. □

We note that Theorem 6.9 is sufficient to show that the polytope of fractional k-stable matchings has the expected structure when k is odd as well.

Corollary 6.10. *For all $k \in \mathbb{N}$, the polytope P_{2k+1} of fractional $2k+1$-stable matchings for an $n \times n$ instance I is the set of all $wt : E(G(I)) \to \mathbb{R}$ that uphold the constraints $\{Q_1, Q_2, Q_3, Q_4(\psi_I^{2k+1}(\emptyset))\}$.*

Proof. By Theorem 4.9, $\psi_I^\infty \subseteq \psi_I^{2k+1}(\emptyset) \subseteq E(G(I))$; therefore, by Theorem 6.9, the desired result holds. $\square$

Taking together Theorem 6.6 and Corollary 6.10, we see that Theorem 6.3 holds.

6.3 Counterexamples on Characterizations of the S-Stable Polytopes

It is easy to see that for any instance I and $S \subseteq E(G(I))$, the polytope $P_S(I)$ of fractional S-stable matchings is constrainted by all of the elements in $\{Q_1, Q_2, Q_3, Q_4(S)\}$. Furthermore, by Theorem 6.3, we showed that these constraints are also sufficient to define $P_S(I)$ when $S = \psi_I^k(\emptyset)$ for some $k \in \mathbb{N}$. In general, however, it is not true that these constraints suffice to define $P_S(I)$. In this section, we present an example that shows this. Motivated by this example, we define a new family of constraints that are satisfied by all S-stable matchings, and show that when added to $\{Q_1, Q_2, Q_3, Q_4(S)\}$, they are sufficient to define P_S for the given example. However, we then provide another example for which our expanded family of constraints is once again insufficient.

For the following examples, we let I_0 be the instance such that the following holds:

- $V_m = \{m_1, m_2, m_3, m_4\}$ and $V_d = \{d_1, d_2, d_3, d_4\}$.
- For all $i \in \{1, 2, 3, 4\}$, m_i's preference list is $[d_1, d_2, d_3, d_4]$ and d_i's preference list is $[m_1, m_2, m_3, m_4]$.

Example 6.11. *Let $S = \{(m_2, d_2)\}$, and $wt_0 : E(G(I)) \to \mathbb{R}$ be the fractional matching such that $wt_0(m_i, d_j) = 1$ if $(i, j) = (4, 4)$, $\frac{1}{2}$ if $(i, j) \in \{(1, 1), (1, 2), (2, 1), (2, 3), (3, 2), (3, 3)\}$, and 0 otherwise. Then, wt_0 satisfies all of the conditions in $\{Q_1, Q_2, Q_3, Q_4(S)\}$, but is not in the polytope of S-stable matchings.*

The only way to represent wt_0 from Example 6.11 as a linear combination of perfect stable matchings is as $\frac{1}{2}(M' + M^*)$, where $M' = \{(m_1, d_2), (m_2, d_1), (m_3, d_3), (m_4, d_4)\}$ and $M^* = \{(m_1, d_1), (m_2, d_3), (m_3, d_2), (m_4, d_4)\}$; in particular, we note that M^* is not (m_2, d_2)-stable. This example shows us that we need to consider another family of potential constraints for the polytope of S-stable matchings for arbitrary $S \subseteq E(G(I))$.

We note that $mn(T)$, the **matching number** of a graph T, is the maximum number of edges in any matching that is a subgraph of T. Given any instance I, any edge $e_0 = (m_0, d_0) \in E(G(I))$, and any $T \subseteq G(I)$, we define $\zeta_m(e_0, T)$ as follows:

- If T contains some edge $e_1 = (m_0, d_1)$ such that m_0 prefers d_1 to d_0, then $\zeta_m(e_0, T) = mn(T')$, where $T' = T - \{(m_0, d) : m_0 \text{ strictly prefers } d_0 \text{ to } d\}$.

- Otherwise, $\zeta_m(e_0, T) = mn(T'') - 1$, where $T'' = T \cup \{(m_0, d) : m_0 \text{ prefers } d \text{ to } d_0\} - \{(m_0, d) : m_0 \text{ strictly prefers } d_0 \text{ to } d\}$.

Similarly, we define $\zeta_d(e_0, T)$ as follows:

- If T contains some edge $e_1 = (m_1, d_0)$ such that d_0 prefers m_1 to m_0, then $\zeta_d(e_0, T) = mn(T')$, where $T' = T - \{(m, d_0) : d_0 \text{ strictly prefers } m_0 \text{ to } m\}$.

- Otherwise, $\zeta_d(e_0, T) = mn(T'') - 1$, where $T'' = T \cup \{(m, d_0) : d_0 \text{ prefers } m \text{ to } m_0\} - \{(m, d_0) : d_0 \text{ strictly prefers } m_0 \text{ to } m\}$.

We then define $\zeta(e_0, T) = max(\zeta_m(e_0, T), \zeta_d(e_0, T))$.

Theorem 6.12. *Let M be any S-stable matching, e_0 be any edge in S, and $T \subseteq G(I)$. Then, $|M \cap T| \leq \zeta(e_0, T)$.*

Proof. Let $e_0 = (m_0, d_0)$. By the fact that M is $\{e_0\}$-stable, we see that m_0 prefers $p_M(m_0)$ to d_0 or d_0 prefers $p_M(d_0)$ to m_0 (or both).

Suppose m_0 prefers $p_M(m_0)$ to d_0. Then, every edge of T in M must also be in T', so $|M \cap T| = |M \cap T'| \leq mn(T')$. In addition, if no edge in T is of the form (m_0, d_1), where m_0 prefers d_1 to d_0, then M must include exactly one edge in $T'' - T$, and so $|M \cap T| = |M \cap T''| - 1 \leq mn(T'') - 1$. In either case, $|M \cap T| \leq \zeta_m(e_0, T) \leq \zeta(e_0, T)$.

Similarly, if d_0 prefers $p_M(d_0)$ to m_0, then $|M \cap T| \leq \zeta_d(e_0, T) \leq \zeta(e_0, T)$. $\square$

Corollary 6.13. *The polytope of S-stable matchings over I is constrained by $\sum_{e \in T} wt(e) \leq \zeta(e_0, T)$ for all $e_0 \in S, T \subseteq E(G(I))$.*

In particular, when we look at the polytope of $\{m_2, d_2)\}$-stable matchings in I_0 (using the PolyhedralSet function in Maple), we find that it has the following constraints:

1. For all $i \in \{1,2,3,4\}$, $\sum_{j=1}^{4} wt(m_i, d_j) = \sum_{j=1}^{4} wt(m_j, d_i) = 1$.

2. For all $i, j \in \{1,2,3,4\}$, $wt(m_i, d_j) \geq 0$.

3. $\sum_{i=2}^{4} wt(m_2, d_i) + \sum_{i=3}^{4} wt(m_i, d_2) \leq 1$.

4. For all $i, j \in \{3,4\}$, $wt(m_2, d_j) + wt(m_i, d_j) + wt(m_i, d_2) \leq 1$.

5. $\sum_{i,j=2}^{4}(wt(m_i, d_j)) - (m_2, d_2) \leq 2$.

In particular, we note that the constraints listed in items 1-3 consist of the elements of $\{Q_1, Q_2, Q_3, Q_4(\{(m_2, d_2)\})\}$, while the constraints in items 4 and 5 are constraints described by Corollary 6.13 (with $T = \{(m_2, d_j), (m_i, d_2), (m_i, d_j)\}$ for each $i, j \in \{3,4\}$ for item 4, and $T = \{(m_i, d_j) : i, j \in \{2,3,4\}\} - \{(m_2, d_2\}$ for item 5). However, there exist examples where this family is insufficient.

Example 6.14. *Let $S = \{(m_2, d_2), (m_3, d_3)\}$. Then, the polytope of S-stable matchings over I_0 is constrained by the following:*

1. *For all $i \in \{1,2,3,4\}$, $\sum_{j=1}^{4} wt(m_i, d_j) = \sum_{j=1}^{4} wt(m_j, d_i) = 1$.*

2. *For all $i, j \in \{1,2,3,4\}$, $wt(m_i, d_j) \geq 0$.*

3. *$\sum_{i=2}^{4} wt(m_2, d_i) + \sum_{i=3}^{4} wt(m_i, d_2) \leq 1$.*

4. *For all $i \in \{3,4\}$, $wt(m_2, d_i) + wt(m_i, d_2) + wt(m_i, d_i) \leq 1$.*

5. *$wt(m_3, d_3)) + wt(m_3, d_4) + wt(m_4, d_3) \leq 1$.*

6. *$wt(m_3, d_4) + wt(m_4, d_3) + wt(m_4, d_4) \leq 1$.*

7. *$wt(m_1, d_1) + wt(m_2, d_1) + wt(m_2, d_3) + wt(m_2, d_4) \leq 1$.*

8. *$wt(m_1, d_4) + wt(m_3, d_1) - wt(m_4, d_3) \geq 0$.*

9. *$wt(m_1, d_3) + wt(m_4, d_1) - wt(m_3, d_4) \geq 0$.*

We note that the constraints in item 1 are $Q_2 \cup Q_3$, the constraints in item 2 are Q_1, and the constraints in items 3 and 5 are $Q_4(S)$. The constraints in items 4, 6, and 7 are ones described by Corollary 6.13 (with $e_0 = (m_2, d_2)$ for the constraints in items

4 and 7, and $e_0 = (m_3, d_3)$ for the constraint in item 6). However, the constraints in items 8 and 9 are not part of any known family of constraints, and at the present time, we do not know the best way of identifying a general family of constraints that they belong to.

Chapter 7

Achieveable Graphs

In many of the previous sections, we have explored how an arbitrary instance I is influenced by the underlying graph $G(I)$. It is trivial to see that for any bipartite graph G, there exists an instance I such that $G(I) = G$; however, as our analyses in the previous sections have shown, some number of these edges may ultimately be irrelevant to the overall structure of the stable matchings over I. (As an example, an incomplete instance I and its completion are very similar instances, but their underlying graphs are very different.) We therefore constrain our question further, and restrict our instances to those where every edge appears in a stable matching.

We define an instance I to be **concise** if every edge appears in a stable matching over I. (We note that this is equivalent to $\psi_I^\infty = G(I)$.) Given a bipartite graph G with vertex parts V_m and V_d, we say that a concise instance I **achieves** G if $G = G(I)$, and that G is **achieveable** if there exists a concise instance that achieves G.

In this chapter, we will look at various questions related to whether a given graph is achieveable, and consider what types of instances can achieve it. In Section 7.1, we consider what instances can achieve a complete graph. Section 7.2 looks at a number of necessary conditions that we identified for a graph to be achieveable. We look at a number of conditions that we originally conjectured to be sufficient to show that a graph is achieveable - along with relevant counterexamples - in the last part of Section 7.2 as well as Section 7.3.

7.1 Achieving the Complete Bipartite Graph

Given the significance of complete instances, it is natural to ask about the structure of instances that are both concise and complete.

Proposition 7.1. *Suppose that an instance I with n mentors and n' students is concise and complete. Then, $n = n'$.*

Proof. If I is both concise and complete, then every vertex appears in an edge for some stable matching. However, by Theorem 2.4, every stable matching over I covers the same vertices, so every stable matching must be perfect - and a perfect matching over I can only exist if I has the same number of mentors and students. □

We would like to obtain simple necessary and sufficient conditions on an instance with n mentors and n students so that it achieves $K_{n,n}$ - i.e. so that every mentor-student pair belongs to some stable matching. There is a simple general construction that achieves this: recall that an $n \times n$ **Latin square** is a matrix in which each row and column is a permutation of $\{1, \ldots, n\}$. Given any $n \times n$ Latin square C, we construct the stable matching instance $I(C)$ associated with C as follows:

1. Set $V_m(I(C)) = \{m_i : i \in [n]\}$ and $V_d(I(C)) = \{d_i : i \in [n]\}$. Initialize each vertex's preference list as empty.

2. For k from 1 to n, do the following: for all (i, j) such that $C_{(i,j)} = k$, put d_j at the bottom of m_i's preference list and m_i at the top of d_j's preference list.

We note that for any $n \times n$ Latin square C, $I(C)$ is a concise and complete $n \times n$ instance. (To see that every mentor-student pair appears in some stable matching, we note that for all $k \in [n]$, $M_k = \{(m_i, d_j) : C_{(i,j)} = k\}$ is a stable matching over $I(C)$, and that $\cup_{k \in [n]} M_k = K_{n,n}$.) As a result, every instance associated to an $n \times n$ Latin square achieves $K_{n,n}$. It is natural to ask whether every instance that achieves $K_{n,n}$ is associated to some $n \times n$ Latin squares. With the aid of Maple, we established this for all $n \geq 4$, but obtained a 5×5 instance that achieves $K_{5,5}$ but does not come from a Latin square.

Example 7.2. *Consider the instance with $V_m = \{m_1, m_2, m_3, m_4, m_5\}$, $V_d = \{d_1, d_2, d_3, d_4, d_5\}$, and the following preference lists:*

$$m_1 : (d_1, d_4, d_2, d_5, d_3) \qquad d_1 : (m_3, m_2, m_5, m_4, m_1)$$

$$m_2 : (d_2, d_3, d_5, d_1, d_4) \qquad d_2 : (m_5, m_3, m_1, m_4, m_2)$$

$$m_3 : (d_3, d_5, d_4, d_2, d_1) \qquad d_3 : (m_1, m_4, m_2, m_5, m_3)$$

$$m_4 : (d_4, d_1, d_2, d_3, d_5) \qquad d_4 : (m_2, m_5, m_3, m_1, m_4)$$

$$m_5 : (d_5, d_3, d_1, d_4, d_2) \qquad d_5 : (m_4, m_1, m_2, m_3, m_5)$$

This instance achieves $K_{5,5}$, but there does not exist a perfect matching where each mentor is partnered with their second choice (since m_2 and m_5 have the same second choice).

7.2 Properties of Achieveable Graphs

In our investigations of achieveable graphs, we have focused on proving (or disproving) the following conjecture:

Conjecture 7.3. *Given a bipartite graph $G \subseteq K_{n,n}$, there exists an algorithm to determine whether G is achieveable in $O(n^k)$ time, for some $k \in \mathbb{N}$.*

While we were ultimately unsuccessful in finding a conclusive answer to this question, we did uncover a number of interesting results.

Proposition 7.4. *A bipartite graph G is achieveable iff there exists an instance I' such that G is the union of all stable matchings over I'.*

Proof. If G is achieveable, then there exists a concise instance I that achieves G. Since I is concise, the union of all stable matchings over I is $G(I) = G$.

Conversely, suppose there exists an instance I' such that G is the union of all stable matchings over I'; then, we can set $I = I'[G]$. Since every stable matching M over I' is $\subseteq G$, every such M is also stable over I'. Consequentially, every edge in $G(I) = G$ appears in a matching that is stable over I and I', so I is concise and achieves G - thereby proving G is achieveable. □

Recall that ψ_I^∞ is the unique hub of I.

Proposition 7.5. *A bipartite graph G is achieveable iff there exists an instance I' such that $G = \psi_{I'}^\infty$.*

Proof. If G is achieveable, then there exists a concise instance I that achieves G. Since I is concise, $\psi(G(I)) = G(I)$, so $\psi_I^\infty = G(I) = G$.

Conversely, suppose there exists an instance I' such that $\psi_{I'}^\infty = G$. We can set $I = I'[G]$ - by Proposition 4.14, the set of stable matchings over I is the set of all hub-stable matchings over I', and so their union equals $G = G(I)$. Therefore, I is concise and achieves G - thereby proving G is achieveable. □

Note that stable matchings are not necessarily perfect matchings, and G may have isolated vertices; however, such vertices are ultimately irrelevant in determining whether G is achieveable.

Proposition 7.6. *Suppose G has an isolated vertex v_0, and $G' = G - \{v_0\}$. Then, G is achieveable iff G' is achieveable.*

Proof. If G' is achieveable, then there exists some concise instance I' such that $G(I') = G'$. Then, the instance I with $V_m(I) = V_m(I') \cup \{v_0\}$ and $V_d(I) = V_d(I')$ such that v_0's preference list is empty, and every other vertex has the same preference list as in I', achieves G.

Conversely, suppose G is achieved by the instance I. WLOG, assume that $v_0 \in V_m(I)$. Then, the instance I' with $V_m(I') = V_m(I) - \{v_0\}$ and $V_d(I') = V_d(I)$, such that every vertex has the same preference list as in I, achieves G'. □

In a similar way, we can show that a bipartite graph is achieveable iff all of its connected components are achieveable. However, this still leaves the question of whether a given connected graph is achieveable.

One necessary condition for an achieveable graph with no isolated vertices is that it is the union of perfect matchings, since every stable matching must be a perfect matching. We can find necessary and sufficient conditions for this property to hold. We define a bipartite graph to uphold the **extended Hall's condition** if it upholds Hall's condition and, for any set of vertices X such that its neighborhood $N(X)$ has $|N(X)| = |X|$, that $N(N(X)) = X$.

Theorem 7.7. *A nonempty graph with no isolated vertices is a union of perfect matchings iff it upholds the extended Hall's condition.*

Proof. Suppose that a bipartite graph G is a union of the elements of $\mathcal{M}$, a (nonempty) set of perfect matchings. Since G contains at least one perfect matching - namely, any $M \in \mathcal{M}$ - it must uphold Hall's condition. Now, consider any set of vertices X such that its neighborhood $N(X)$ in G upholds $|N(X)| = |X|$. For any matching $M \in \mathcal{M}$, every element of X is matched by M with an element of $N(X)$, the neighborhood of X in G; however, since $|N(X)| = |X|$, each such element in $N(X)$ must be matched with an element of X by the pigeonhole principle. Since M is an arbitrary matching in $\mathcal{M}$, and G is the union of all such M, any element of $N(X)$ can only have elements of X in its neighborhood. Since $X \subseteq N(N(X))$, $N(N(X)) = X$, and G upholds the extended Hall's condition.

Conversely, suppose that a bipartite graph G upholds the extended Hall's condition. Consider any edge $e \in G$; if we remove the vertices in e from G, the resulting graph upholds Hall's condition, and so contains a perfect matching. If we add e to that matching, we produce a perfect matching $M_e \subseteq G$ that contains e. Taking the union of M_e for all $e \in G$ produces G, so G is the union of a set of perfect matchings. $\square$

However, not every graph which can be expressed as the union of perfect matchings is achieveable. For example, consider the complete 3×3 bipartite graph G_0 with a single edge removed.

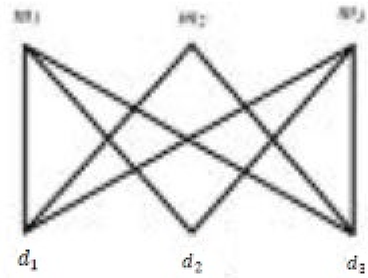

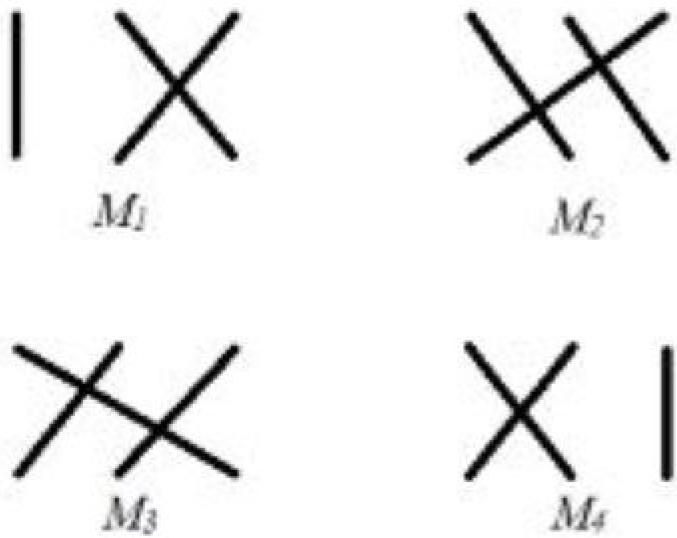

Proposition 7.8. *G_0 can be expressed as the union of perfect matchings, but it is not achieveable.*

Proof. We note that the perfect matchings M_1, M_2, M_3, and M_4 in the above figure are the only perfect matchings contained in G_0; these four matchings have their union equal G_0. In addition, since each contains an edge not in the others, G_0 can only be expressed as a union of perfect matchings in this way, so any instance that achieves G_0 must have $\{M_1, M_2, M_3, M_4\}$ as its set of stable matchings.

Assume we have such an instance I_0, with M_1 as the mentor-optimal stable matching WLOG. Since each matching contains an edge not in any of the others, the lattice of stable matchings of I_0 must be totally ordered. (Otherwise, there would exist some i such that $\{M_i\}$ is the sublattice of stable matchings with a specific edge, and a j such that M_j is neither above nor below this sublattice, creating a contradiction). Now, look at $e_2 = (m_2, d_3)$ and $e_3 = (m_3, d_2)$. Since e_2 is only present in M_1 and M_2, and M_1 is the mentor-optimal stable matching, M_2 must cover M_1 in the lattice of stable matchings. However, since e_3 is only present in M_1 and M_3, M_3 must similarly cover M_1 in the lattice. This creates a contradiction - since the lattice is totally ordered, only one stable matching can cover M_1 - so no such instance can exist, and G_0 is not achieveable. □

7.3 More Counterexamples in Achieveability

While our investigations did not lead to an efficient algorithm that would determine if a graph is achieveable, we did find a number of noteworthy examples that expanded our understanding of achieveable graphs. During our inquiries, there were a number of statements that we initially thought might be true, but we ultimately disproved. The most noteworthy ones are listed here:

1. Any graph that can be expressed as the union of perfect matchings is achieveable.
2. If G is an achieveable graph and $e \in E(G)$, then there exists an instance I such that I achieves G and e is an edge in the mentor-optimal stable matching over I.
3. Given an instance I, every minimal set of stable matchings that covers $G(I)$ has the same number of members.
4. If G and H are achieveable graphs with $G \cap H$ isomorphic to the uniform degree 1 graph, then $G \cup H$ is achieveable.

Statement 1 is shown to be false by Proposition 7.8; in this section, we will show that remaining three statements are false, and provide relevant counterexamples.

Example 7.9. *There exists an achieveable graph G and an edge $e \in E(G)$ such that no instance which achieves G includes e in the mentor-optimal stable matching.*

The example we use is the seven-edge graph G shown below. It can only be expressed as a union of perfect matchings with the shown three matchings, so any instance that achieves G has exact that set of stable matchings. Furthermore, since every vertex has degree > 2, the mentor-optimal and student-optimal stable matchings cannot share an edge (as otherwise, every stable matching would have that edge); therefore, M_2 cannot be the mentor-optimal stable matching. The edge (m_2, d_2) only appears in M_2, so it cannot be in the mentor-optimal stable matching.

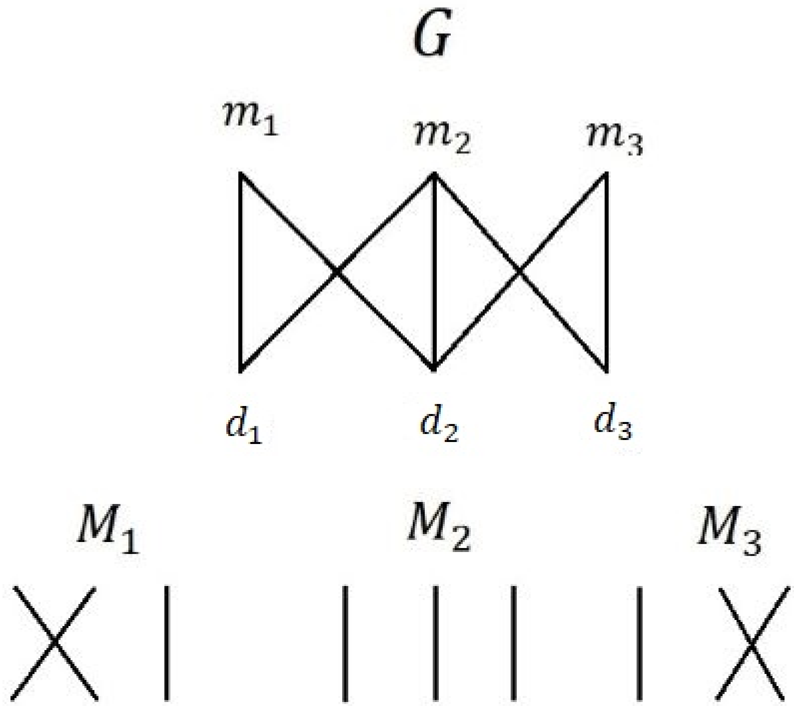

Example 7.10. *There exists an instance I such that not every minimal set of stable matchings that covers $G(I)$ has the same number of members.*

The example instance I is shown below. In particular, I achieves the vertex-disjoint union of two copies of $K_{3,3}$; two different minimal sets T and T' that cover the achieved graph are shown with it. (Note that $|T| = 3 \neq 4 = |T'|$.)

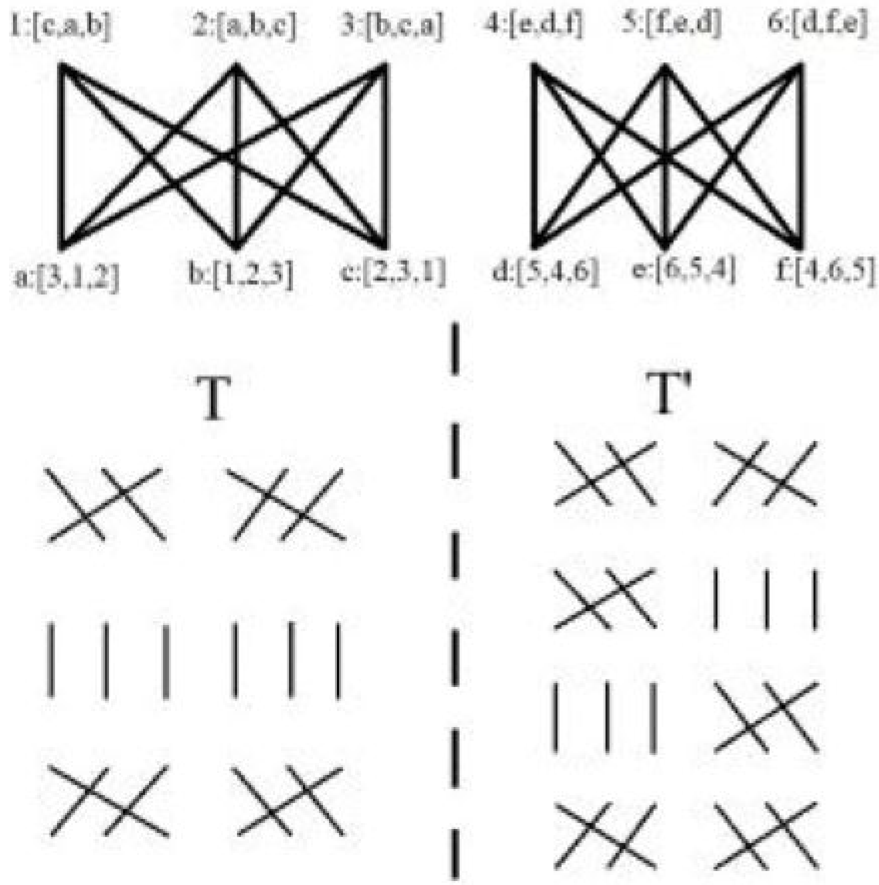

Example 7.11. *There exist achieveable graphs with G, H such that $G \cap H$ isomorphic to the uniform degree 1 graph and $G \cup H$ is not achieveable.*

An example of such a pair $\{G, H\}$ is shown below. (Notice that the G used here is the same G used as our example in Example 7.9, with an additional disjoint edge added - in fact, the reason that this serves as a counterexample is because the edge of G that cannot be part of the mentor-optimal matching is shared with H.) In particular, this example gives us reason to think that determining whether a graph is achieveable is very difficult without information on what the mentor-optimal (or student-optimal) stable matching is.

G

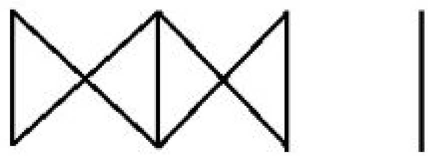

H

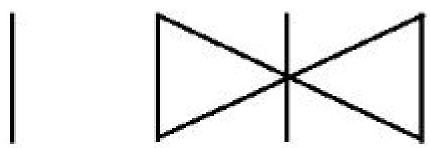

Chapter 8

Bounding the Number of Variables in a Low Degree Boolean Function

In this chapter, we prove that there is a constant $C \leq 6.614$ such that every Boolean function of degree at most d (as a polynomial over $\mathbb{R}$) is a $C \cdot 2^d$-junta, i.e. it depends on at most $C \cdot 2^d$ variables. This improves the $d \cdot 2^{d-1}$ upper bound of Nisan and Szegedy [Computational Complexity 4 (1994)].

The bound of $C \cdot 2^d$ is tight up to the constant C, since a read-once decision tree of depth d depends on all $2^d - 1$ variables. We slightly improve this lower bound by constructing, for each positive integer d, a function of degree d with $3 \cdot 2^{d-1} - 2$ relevant variables. A similar construction was independently observed by Shinkar and Tal. [1]

8.1 Introduction to the Degree of a Boolean Function

The **degree** of a Boolean function $f : \{0,1\}^n \to \{0,1\}$, denoted $\deg(f)$, is the degree of the unique multilinear polynomial in $\mathbb{R}[x_1, ..., x_n]$ that agrees with f on all inputs from $\{0,1\}^n$. Minsky and Papert [MP88] initiated the study of combinatorial and computational properties of Boolean functions based on their representation by polynomials. We refer the reader to the excellent book of Ryan O'Donnell [O'D14] on analysis of Boolean functions, and surveys by Harry Buhrman and Ron DeWolf [BDW02], and Pooya Hatami [HKP11] discussing relations between various complexity measures of Boolean functions.

An input variable x_i is **relevant to** f if x_i appears in a monomial having nonzero coefficient in the multilinear representation of f. Let $R(f)$ denote the number of relevant

[1] This section was previously published on ArXiv in 2018, and was published in Combinatorica earlier this year [CHS20].

variables of f. Nisan and Szegedy ([NS94], Theorem 1.2) proved that $R(f) \leq \deg(f) \cdot 2^{\deg(f)-1}$.

Let R_d denote the maximum of $R(f)$ over Boolean functions f of degree at most d, and let $C_d = R_d 2^{-d}$. By the result of Nisan and Szegedy, $C_d \leq d/2$. On the other hand, $R_d \geq 2R_{d-1} + 1$, since if f is a degree $d-1$ Boolean function with R_{d-1} relevant variables, and g is a copy of f on disjoint variables, and z is a new variable then $zf + (1-z)g$ is a degree d Boolean function with exactly $2R_{d-1} + 1$ relevant variables. Thus $C_d \geq C_{d-1} + 2^{-d}$, and so $C_d \geq 1 - 2^{-d}$. Since C_d is an increasing function of d it approaches a (possibly infinite) limit $C^* \geq 1$.

In this chapter we prove:

Theorem 8.1. *There is a positive constant C so that $R(f)2^{-\deg(f)} \leq C$ for all Boolean functions f, and thus $C_d \leq C$ for all $d \geq 0$. In particular C^* is finite.*

Throughout this paper we use $[n] = \{1, \ldots, n\}$ for the index set of the input variables for a Boolean function f. A **maxonomial** of f is a set $S \subseteq [n]$ of size $\deg(f)$ for which $\prod_{i \in S} x_i$ has a nonzero coefficient in the multilinear representation of f. A **maxonomial hitting set** is a subset $H \subseteq [n]$ that intersects every maxonomial. Let $h(f)$ denote the minimum size of a maxonomial hitting set for f, and let h_d denote the maximum of $h(f)$ over Boolean functions of degree d. In Section 8.2 we prove:

Lemma 8.2. *For every $d \geq 1$, $C_d - C_{d-1} \leq h_d 2^{-d}$.*

Through telescoping, this implies:

Corollary 8.3. *For every $d \geq 0$, $C_d \leq \sum_{i=1}^{d} h_i 2^{-i}$.*

The next lemma is a simple combination of previous results.

Lemma 8.4. *For any Boolean function f, $h(f) \leq \deg(f)^3$, and so for all $i \geq 1$, $h_i \leq i^3$.*

Proof. Nisan and Smolensky (see Lemma 6 of [BDW02]) proved $h_i \leq \deg(f) bs(f)$, where $bs(f)$ is the block sensitivity of f. Combining with $bs(f) \leq \deg(f)^2$ (proved by Avishay Tal ([**?**]), improving on $bs(f) \leq 2\deg(f)^2$ of Nisan and Szegedy ([NS94], Lemma 3.8) yields $h(f) \leq \deg(f)^3$. $\square$

Using Lemma 8.4, the infinite sum in Corollary 8.3 converges, and Theorem 8.1 follows.

Given that C^* is finite, it is interesting to obtain upper and lower bounds on C^*. The bounds that we will show in this paper are $3/2 \leq C^* \leq \frac{13545}{2048} \leq 6.614$; we discuss these bounds in Section 8.3. (Recently, Wellens ([Wel19],Theorem 3) refined our argument to obtain an improved upper bound of $C^* \leq 4.416$.)

Filmus and Ihringer ([FI18]) recently considered an analog of the parameter $R(f)$ for the family of **level k slice functions**, which are Boolean functions whose domain is restricted to the set of inputs of Hamming weight exactly k. They showed that, provided that $\min(k, n-k)$ is sufficiently large, every level k slice function on n variables of degree at most d depends on at most R_d variables. ([FI18], Theorem 1.1) As a result, our improved upper bound on R_d applies also to the number of relevant variables of slice functions.

Proof Overview

Similar to Nisan and Szegedy ([NS94], Section 2.3), we upper bound $R(f)$ by assigning a weight to each variable, and bounding the total weight of all variables. The weight assigned to a variable by Nisan and Szegedy was its **influence** on f; the novelty of our approach is to use a different weight function.

We assign to a variable x_i of a Boolean function f a weight $w_i(f)$ that is 0 if f does not depend on x_i and otherwise equals $2^{-\deg_i(f)}$ where $\deg_i(f)$ is the degree of the maximum degree monomial of f containing x_i. We then upper and lower bound the total weight $W(f)$ of a degree d Boolean function f. It follows from the definition that for a degree d Boolean function f, $W(f) \geq 2^{-d} \cdot R(f)$. Hence, to upper bound $R(f)$ it suffices to upper bound $W(f)$. Let W_d be the maximum of $W(f)$ among degree d Boolean functions f. We prove that

$$W_d \leq h_d 2^{-d} + W_{d-1}.$$

We show this by considering a minimum size maxonomial hitting set H for a $W(f)$ maximizing f. We argue that for such an f, all variables in H have maximum degree

d, and hence their total weight adds up to $2^{-d} \cdot |H|$. Additionally, we show that the remaining variables have total weight at most W_{d-1}, by considering degree $d-1$ restrictions of f that are achieved by fixing variables in H. See proof of Proposition 8.7 for more details.

Combining above with Lemma 8.4 we have shown that $R(f) \leq 2^d \cdot \sum_{i=1}^{d} i^3 2^{-i}$, which readily implies $R(f) \leq 26 \cdot 2^d$. However, the same argument as above also implies

$$R(f) \leq 2^d \cdot (W_k + \sum_{i=k+1}^{d} i^3 2^{-i}).$$

Finally, plugging a bound of $W_k \leq k/2$ which follows from previous works and optimizing the right hand side, we obtain an improved bound of $R(f) \leq 6.614 \cdot 2^d$.

8.2 Proof of Lemma 8.2

For a variable x_i, let $\deg_i(f)$ be the maximum degree among all monomials that contain x_i and have nonzero coefficient in the multilinear representation of f. Let $w_i(f) := 0$ if x_i is not relevant to f, and $w_i(f) := 2^{-\deg_i(f)}$ otherwise. Note that if x_i is a relevant variable of the degree d function f, then $w_i(f) = 2^{-\deg_i(f)} \geq 2^{-\deg(f)} = 2^{-d}$.

The weight of f, $W(f)$, is defined to be $\sum_i w_i(f)$, and W_d denotes the maximum of $W(f)$ over all Boolean functions f of degree at most d; this maximum is well defined since, by the Nisan-Szegedy upper bound of R_d, it is taken over a finite set of functions. A function f of degree at most d for which $W_d = W(f)$ is **W_d-maximizing**.

Lemma 8.2 will follow as an immediate consequence of $W_d = C_d$ (Corollary 8.6) and $W_d \leq W_{d-1} + h_d 2^{-d}$ (Proposition 8.7).

Proposition 8.5. *If f is W_d-maximizing then every relevant variable of f belongs to a degree d monomial.*

Proof. Let the relevant variables of f be $x_1, \ldots, x_n$. Assume for contradiction that there are $l \geq 1$ variables that do not belong to any degree d monomial, and that these variables are $x_1, \ldots, x_l$. We now construct a function g of degree at most d such that $W(g) > W(f)$, contradicting that f is W_d-maximizing. Let g be the $n+l+1$-variate function given by:

$$g(x_1, \ldots, x_{n+l+1}) := x_{n+l+1} f(x_1, \ldots, x_n) + (1 - x_{n+l+1}) f(x_{n+1}, \ldots, x_{n+l}, x_{l+1}, \ldots, x_n).$$

This function is boolean since it is equal to $f(x_{n+1}, \ldots, x_{n+l}, x_{l+1}, \ldots, x_n)$ if $x_{n+l+1} = 0$ and to $f(x_1, \ldots, x_n)$ if $x_{n+l+1} = 1$. It clearly has no monomials of degree larger than $d+1$. Since x_i appears in no degree d monomials of f for any $i \leq l$, $f(x_1, \ldots, x_n)$ and $f(x_{n+1}, \ldots, x_{n+l}, x_{l+1}, \ldots, x_n)$ have the same set of degree d monomials. Thus the degree $d+1$ monomials of $x_{n+l+1} f(x_1, \ldots, x_n)$ cancel out the degree $d+1$ monomials of $(1 - x_{n+l+1}) f(x_{n+1}, \ldots, x_{n+l}, x_{l+1}, \ldots, x_n)$, and g has degree at most d. Furthermore, all of the degree d monomials involving $x_{l+1}, \ldots, x_n$ appear with the same coefficient in g as in f so $w_i(g) = w_i(f) = 2^{-d}$ for all $i \in \{l+1, \ldots, n\}$. Also, for each $i \in \{1, \ldots, l\}$, any monomial $m = x_i m'$ containing x_i gives rise to monomials $x_{n+l+1} x_i m'$ and $-x_{n+1+i} x_{n+i} m'$ in g and so $w_i(g) = w_{n+i}(g) = \frac{1}{2} w_i(f)$. Thus we have:

$$\begin{aligned}
W(g) = \sum_{i=1}^{n+l+1} w_i(g) &= \sum_{i=1}^{l} (w_i(g) + w_{n+i}(g)) + \sum_{i=l+1}^{n} w_i(g) + w_{n+l+1}(g) \\
&= \sum_{i=1}^{l} w_i(f) + \sum_{i=l+1}^{n} w_i(f) + w_{n+l+1}(g) \\
&= W(f) + w_{n+l+1}(g) > W(f),
\end{aligned}$$

where the final inequality holds since x_{n+l+1} is a relevant variable of g (which is true since for any monomial m of f containing x_1, $m x_{n+l+1}$ is a monomial of g). Thus, g is a function of degree d with $W(g) > W(f)$, which gives us the desired contradiction to complete the proof. □

Corollary 8.6. *For all $d \geq 1$, $W_d = C_d$.*

Proof. For any function f of degree at most d, we have $W(f) \geq R(f) 2^{-d}$. Thus $W_d \geq C_d$. If f is W_d-maximizing then by Proposition 8.5, $W(f) = R(f) 2^{-d} \leq C_d$.[2] □

---

[2] In the first version of the paper we published with this result, our proof that $W_d \leq C_d$ was erroneous; this has been amended to its present form in this version. We thank Jake Lee Wellens for pointing out the error in the previous version.

Therefore, to prove Lemma 8.2 it suffices to prove:

Proposition 8.7. $W_d - h_d 2^{-d} \leq W_{d-1}$.

Proof. Again, let f be W_d-maximizing. Let H be a maxonomial hitting set for f of minimum size. Note that $\deg_i(f) = d$ for all $i \in H$, as otherwise $H - \{i\}$ would be a smaller maxonomial hitting set. Thus:

$$W(f) = \sum_i w_i(f) = 2^{-d}|H| + \sum_{i \notin H} w_i(f). \tag{8.1}$$

We will now show:

$$\sum_{i \notin H} w_i(f) \leq W_{d-1}, \tag{8.2}$$

which, combined with Equation (8.1), yields the desired conclusion $W_d \leq 2^{-d} h_d + W_{d-1}$. We deduce Equation (8.2) by bounding $w_i(f)$ by the average of $w_i(f')$ over a collection of **restrictions** f' of f (which we will define later). We recall some definitions. A **partial assignment** is a mapping $\alpha : [n] \longrightarrow \{0, 1, *\}$, and Fixed$(\alpha)$ is the set $\{i : \alpha(i) \in \{0, 1\}\}$. For $J \subseteq [n]$, PA(J) is the set of partial assignments α with Fixed$(\alpha) = J$. The **restriction** of f by α, f_α, is the function on variable set $\{x_i : i \in [n] - \text{Fixed}(\alpha)\}$ obtained by setting $x_i = \alpha_i$ for each $i \in \text{Fixed}(\alpha)$.

Lemma 8.8. *For every $J \subseteq [n]$ and $i \notin J$,*

$$w_i(f) \leq 2^{-|J|} \sum_{\alpha \in \text{PA}(J)} w_i(f_\alpha).$$

Proof. Fix $j \in J$ and write $f = (1 - x_j) f_0 + x_j f_1$ where f_0 is the restriction of f to $x_j = 0$ and f_1 is the restriction of f to $x_j = 1$.

We proceed by induction on $|J|$. We consider the base cases of $|J| \leq 1$. The $|J| = 0$ case is trivial. Let us now consider the $|J| = 1$ case where we have $J = \{j\}$.

- If f_0 does not depend on x_i, then $w_i(f) = w_i(f_1)/2 \leq (w_i(f_0) + w_i(f_1))/2$.
- If f_1 does not depend on x_i, then $w_i(f) = w_i(f_0)/2 \leq (w_i(f_0) + w_i(f_1))/2$.
- Suppose f_1 and f_0 both depend on x_i.

- If $\deg_i(f_0) < \deg_i(f_1)$, let m be a monomial containing x_i of degree $\deg_i(f_1)$ that appears in f_1. Then $x_j m$ is a maxonomial of $f = x_j(f_0 - f_1) + f_0$. Therefore $\deg_i(f) = 1 + \deg_i(f_1)$. Thus $w_i(f) = \frac{1}{2}w_i(f_1) \leq \frac{1}{2}(w_i(f_0) + w_i(f_1))$.
- If $\deg_i(f_0) \geq \deg_i(f_1)$ then $w_i(f_0) \leq w_i(f_1)$. It suffices that $w_i(f) \leq w_i(f_0)$, and this holds because each monomial that appears in f_0 appears with the same coefficient in $f = x_j(f_1 - f_0) + f_0$.

In every case, we have $w_i(f) \leq \frac{1}{2}(w_i(f_0) + w_i(f_1))$, as desired.

For the induction step, assume $|J| \geq 2$. We start with $w_i(f) \leq \frac{1}{2}(w_i(f_0) + w_i(f_1))$, and apply the induction hypothesis separately to f_0 and f_1 with the set of variables $J - \{j\}$:

$$\begin{aligned} w_i(f) &\leq \frac{1}{2}(w_i(f_0) + w_i(f_1)) \\ &\leq \frac{1}{2}\left(2^{1-|J|}\left(\sum_{\beta \in \mathrm{PA}(J-\{j\})} w_i(f_{0,\beta})\right) + 2^{1-|J|}\left(\sum_{\beta \in \mathrm{PA}(J-\{j\})} w_i(f_{1,\beta})\right)\right) \\ &\leq 2^{-|J|} \sum_{\alpha \in \mathrm{PA}(J)} w_i(f_\alpha). \end{aligned}$$

□

To complete the proofs of Equation (8.2) and Proposition 8.7 apply Lemma 8.8 with J being a hitting set H of minimum size, and sum over $i \in [n] - H$ to get:

$$\sum_{i \in [n]-H} w_i(f) \leq 2^{-|H|} \sum_{i \in [n]-H} \sum_{\alpha \in \mathrm{PA}(H)} w_i(f_\alpha) = 2^{-|H|} \sum_{\alpha \in \mathrm{PA}(H)} W(f_\alpha) \leq W_{d-1},$$

where the last inequality follows since $\deg(f_\alpha) \leq d - 1$ for all $\alpha \in \mathrm{PA}(H)$. □

As noted earlier Corollary 8.6 and Proposition 8.7 combine to prove Lemma 8.2.

8.3 Bounds on C^*

Lemma 8.2 implies $C_d \leq \sum_{i=1}^{d} 2^{-i} h_i$. Combining with Lemma 8.4 yields $C_d \leq \sum_{i=j}^{d} i^3 2^{-i}$, and thus $C^* \leq \sum_{i=1}^{\infty} i^3 2^{-i}$, which equals 26 (since $\sum_{i \geq 0} \binom{i}{j} 2^{-i} = 2$ for all $j \geq 0$, and $i^3 = 6\binom{i}{3} + 6\binom{i}{2} + i$). As noted in the introduction, $R_d \geq 2^d - 1$, and so $C^* \geq 1$. We improve these bounds to:

Theorem 8.9. $\frac{3}{2} \leq C^* \leq \frac{13545}{2048}$.

Proof. For the upper bound, Lemma 8.2 implies that for any positive integer d,

$$C^* \leq C_d + \sum_{i=d+1}^{\infty} 2^{-i} h_i.$$

Using $C_d \leq d/2$ as proved by Nisan and Szegedy, we have

$$C^* \leq \min_d \left(\frac{d}{2} + \sum_{i=d+1}^{\infty} i^3 2^{-i} \right).$$

The minimum occurs at the largest d for which $d^3 2^{-d} > 1/2$, which is 11. Evaluating the right hand side for $d = 11$ gives $C^* \leq \frac{13545}{2048} \leq 6.614$.

We lower bound C^* by exhibiting, for each d,a function Ξ_d of degree d with $l(d) = \frac{3}{2}2^d - 2$ relevant variables. (A similar construction was found independently by Shinkar and [ST17].) It is more convenient to switch our Boolean set to be $\{-1, 1\}$.

We define $\Xi_d : \{-1, 1\}^{l(d)} \to \{-1, 1\}$ as follows. $\Xi_1 : \{-1, 1\} \to \{-1, 1\}$ is the identity function, and for all $d > 1$, Ξ_d on $l(d) = 2l(d-1) + 2$ variables is defined recursively by:

$$\Xi_d(s, t, \vec{x}, \vec{y}) = \frac{s+t}{2} \Xi_{d-1}(\vec{x}) + \frac{s-t}{2} \Xi_{d-1}(\vec{y})$$

for all $s, t \in \{-1, 1\}$ and $\vec{x}, \vec{y} \in \{-1, 1\}^{l(d-1)}$. It is evident from the definition that $\deg(\Xi_d) = 1 + \deg(\Xi_{d-1})$, which is d by induction (as for the base case $d = 1$, Ξ_1 is linear). It is easily checked that Ξ_d depends on all of its variables, and that $\Xi_d(s, t, \vec{x}, \vec{y})$ equals $s \cdot \Xi_{d-1}(\vec{x})$ if $s = t$ and equals $s \cdot \Xi_{d-1}(\vec{(y)})$ if $s \neq t$, and is therefore Boolean. $\square$

[Wel19] recently refined the arguments of this paper to improve the upper bound to $C^* \leq 4.416$.

Chapter 9

A Lower Bound on $H(d)$

In the previous chapter, we showed that we can bound the maximum number of relevant variables in a degree d boolean function by a weighted sum of $H(i)$, for $i \in [d]$. A natural question that arises is how strong of a bound we can achieve via this method - specifically, how small of an upper bound on $H(d)$ we can find. Currently, our best known upper bound on $H(d)$ is d^3, but we conjecture that $H(d)$ is significantly smaller (for example, we know that $H(2) = 2$, which is significantly smaller than $2^3 = 8$). In this chapter, we will discuss the best lower bounds that we have found on $H(d)$, thereby putting a limit on how strong a result our argument can produce without strategic adaptation.

9.1 Maxinomial Hitting Set Size of Compositions

In order to find lower bounds on $H(d)$, we will leverage the behavior of Boolean functions under composition. Recall that for Boolean functions $f : \{0,1\}^n \to \{0,1\}$ and $g : \{0,1\}^m \to \{0,1\}$, their **composition**

$$f \circ g = f\left(g(t_{1,1}, \ldots, t_{1,m}), \ldots, g(t_{n,1}, \ldots, t_{n,m})\right)$$

is a Boolean function in mn variables with variable set $\{t_{i,j} : i \in [n], j \in [m]\}$. It is well known that $\deg(f \circ g) = \deg(f) \cdot \deg(g)$: the set of monomials of $f \circ g$ is the set of all monomials of the form $c_M \prod_{x_i \in M} m_i$, where $M = c_M \prod_{x_i \in M} x_i$ is a monomial of $f(x_1, \ldots, x_n)$ and, for all relevant i, m_i is a monomial of $g(t_{i,1}, \ldots, t_{i,m})$. The degree of such a monomial is maximized when M and all corresponding m_i's are maxonomials, in which case its degree is $\sum_{x_i \in M} \deg(g) = \deg(f) \cdot \deg(g)$. However, we still must show that hitting set size is also multiplicative under composition.

Proposition 9.1. *Let $f : \{0,1\}^n \to \{0,1\}$ and $g : \{0,1\}^m \to \{0,1\}$ be Boolean functions. Then,*

$$h(f \circ g) = h(f) \cdot h(g).$$

Proof. It is easy to check that $S_0 = \{(i,j) : i \in S_1, j \in S_2\}$ is a maxonomial hitting set of $f \circ g$, where S_1 is any maxonomial hitting set of $f(x_1, \ldots, x_n)$ and S_2 is any maxonomial hitting set of $g(t_{1,1}, \ldots, t_{1,m})$. Therefore, $h(f \circ g) \leq h(f) \cdot h(g)$.

We now show that $h(f \circ g) \geq h(f) \cdot h(g)$. Let $S \subseteq \{(i,j) : i \in [n], j \in [m]\}$ be a maxonomial hitting set of $f \circ g$. Let S_i be the set of pairs in S with first coordinate i, and let S' be the set of all $i \in [n]$ such that S_i is a maxonomial hitting set of $g(t_{i,1}, \ldots, t_{i,m})$. We claim that S' is a maxonomial hitting set of $f(x_1, x_2, \ldots)$. Assume to the contrary that there is a maxonomial M_f that S' does not cover. For each i such that $x_i \in M_f$, there is a maxonomial M_i of $g(t_{i,1}, \ldots, t_{i,m})$ that is not hit by S_i. Then, $\prod_{i:x_i \in M_f} M_i$ is a maxonomial of $f \circ g$ that is not hit by S, contradicting the fact that S was a maxonomial hitting set of $f \circ g$. This implies $|S'| \geq h(f)$. Since for every $i \in S'$, $|S_i| \geq h(g)$, we have $|S| \geq h(f)h(g)$. Therefore $h(f \circ g) \geq h(f)h(g)$, and so $h(f \circ g) = h(f)h(g)$. □

Theorem 9.2. *$H(d)$ is supermultiplicative - i.e. $H(d_1 \cdot d_2) \geq H(d_1) \cdot H(d_2)$ for all $d_1, d_2 \in \mathbb{N}$.*

Proof. By the definition of $H(d)$, we can find Boolean functions $f : \{0,1\}^n \to \{0,1\}$ and $g : \{0,1\}^m \to \{0,1\}$ such that $\deg(f) = d_1$, $h(f) = H(d_1)$, $\deg(g) = d_1$, and $h(g) = H(d_2)$. Then, $f \circ g$ is a Boolean function with degree $d_1 \cdot d_2$, and by Proposition 9.1, $h(f \circ g) = h(f) \cdot h(g) = H(d_1) \cdot H(d_2)$. However, by the fact that its degree is $d_1 \cdot d_2$, $h(f \circ g) \leq H(d_1 \cdot d_2)$, and we are done. □

9.2 Low Degree Functions with High Maxonomial Hitting Set Size

In order to find and confirm our lower bounds, we also need to show that $H(d)$ is an increasing function.

Theorem 9.3. *For all $d \in \mathbb{N}$, $H(d+1) \geq H(d)$.*

Proof. Let $f(x_1, \ldots, x_n)$ be a boolean function over with degree d such that $h(f) = H(d)$. We set $f_0(x_1, \ldots, x_n) = f(x_1, \ldots, x_n)$ and perform the following iterative process for each $i \in \mathbb{N} \cup \{0\}$:

- If every relevant variable of $f_i(x_1, \ldots, x_{n+i})$ belongs to a degree $d+1$ monomial, then we set $g(x_1, \ldots, x_{n+i}) = f_i(x_1, \ldots, x_{n+i})$.

- Otherwise, we select any relevant variable x_j of f_i that does not appear in any degree $d+1$ monomial, and define $f_{i+1}(x_1, \ldots, _{n+i+1})$ to be f_i with every occurrence of x_j replaced with $x_j * x_{n+i+1}$. We note that this operation preserves the number of variables required to hit every monomial that was originally degree d, and if f_i has degree at most $d+1$, so does f_{i+1}.

This process must terminate in at most $2^{d+1} * W(f)$ steps (since each f_{i+1} has one more relevant variable than f_i, and $W(f_{i+1}) \leq W(f_i)$, implying $W(f_i) \leq W(f)$ for all $i \in \mathbb{N}$. Thus, the resulting g is a boolean function of degree $d+1$, and $h(g) \geq h(f) = H(d)$. Consequentially, $H(d+1) \geq H(d)$. $\square$

Corollary 9.4. *Let $d_1, d_2 \in \mathbb{N}$. Then, $d_1 \geq d_2 \Rightarrow H(d_1) \geq H(d_2)$.*

With these preliminaries complete, we can now prove a theorem that will let us identify a lower bound on $H(d)$; it accomplishes this by using the iterated composition of a sample Boolean function.

Theorem 9.5. *Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function such that $\deg(f) = d_0$ and $h(f) = h_0$. Then, for all $d \in \mathbb{N}$, $H(d) \geq \frac{d^p}{h_0}$, where $p = \log_{d_0}(h_0)$.*

Proof. For $i \in \mathbb{N}$, we define $f_i : \{0,1\}^{n^i} \to \{0,1\}$ as follows: $f_1(x_1, \ldots, x_n) = f(x_1 \ldots, x_n)$ and $f_{i+1}(x_1, \ldots, x_{n^{i+1}}) = f \circ f_i(x_1, \ldots, x_{n^{i+1}})$ for all $i \in \mathbb{N}$. Then, for all i, $f_i(x_1, \ldots, x_{n^{i+1}})$ is a boolean function with $\deg(f_i) = d_0^i$ and $h(f_i) = h_0^i$, showing that $H(d_0^i) \geq h_0^i$.

Now, for any $d \in \mathbb{N}$, let i be the largest integer such that $d_0^i \leq d$. By Corollary 9.4,

$$H(d) \geq H(d_0^i) \geq h_0^i = \frac{h_0^{i+1}}{h_0} = \frac{(d_0^{i+1})^p}{h_0} > \frac{d^p}{h_0}$$

and we are done. $\square$

As an example, the Boolean function $R : \{0,1\}^4 \to \{0,1\}$ defined by $R(x_1, x_2, x_3, x_4) = x_1 + x_2 - x_1 * x_2 - x_1 * x_3 - x_2 * x_4 + x_3 * x_4$ has $\deg(R) = 2$ and $h(R) = 2$; therefore, by Theorem 9.5, $H(d) \geq \frac{d}{2}$, since $p = \log_2(2) = 1$. During our initial investigations, we conjectured that this lower bound was tight, up to a constant factor.

Conjecture 9.6. *For all $d \in \mathbb{N}$, $H(d) = d$.*

However, we were ultimately able to show that this conjecture was false, using the following function.

Example 9.7. *The function:*

$$contra(x) = x_1x_2x_6 - x_1x_2x_{10} + x_1x_3x_6 - x_1x_3x_9 - x_1x_6x_9 + x_1x_6x_{10}$$

$$-x_2x_3x_8 - x_2x_3x_{10} + x_2x_6x_{10} - x_2x_8x_9 + x_2x_9x_{10} + x_3x_6x_9 + x_3x_8x_{10} + x_8x_9x_{10}$$

$$-x_1x_6 + x_1x_9 + x_2x_3 - x_2x_6 + x_2x_8 - x_3x_6 - x_6x_{10} - x_8x_{10} - x_9x_{10} + x_6 + x_{10}$$

is a Boolean function such that $\deg(contra(x)) = 3$ *and* $h(contra(x)) = 4$.

Using a Maple program (described in greater depth in the next section), we were able to prove the following result.

Theorem 9.8. *$H(3) = 4$.*

While this does contradict our conjecture, the fact that $\deg(contra) = 3$ and $h(contra) = 4$ lets us improve our lower bound on $H(d)$ using Theorem 9.5.

Theorem 9.9. *$H(d) > \frac{d^p}{4}$, where $p = \log_3(4)$.*

9.3 The Computation of $H(3)$

It is easy to show that $H(1) = 1$ and $H(2) = 2$; previously, we have conjectured that $H(d) = d$ for all $d \in \mathbb{N}$. To that end, we look to find the value of $H(3)$, the first term that is not easily found by hand. In this section, we will describe a program that we wrote in order to find the value of $H(3)$, and the degree 3 boolean function f_4 with $h(f_4) = 4$ it found - thereby disproving Conjecture 9.6. For our explanation, the **maxonomial set** of a Boolean function is the sum of its maxonomials.

The basic principle that we use in our program is the following:

Proposition 9.10. *Let $f : \{0,1\}^n \to \{0,1\}$ be a boolean function such that $\deg(f) = d$ and $h(f) = k$. Then, there exists $j \in [n]$ such that for all $\alpha \in PA(\{j\})$:*

- *if $k = 1$, then $\deg(f_\alpha) < d$.*
- *if $k > 1$, then $\deg(f_\alpha) = d$ and $h(f_\alpha) = k - 1$.*

Proof. Set $j \in [n]$ to be any value such that j is in a minimum size maxonomial hitting set of f. For either $\alpha \in PA(\{j\})$, the set of degree d monomials in f_α is the set of all degree d monomials that do not contain x_j (since every degree d monomial that does contain x_j disappears or becomes a degree $d-1$ monomial respectively). If $k = 1$, then every maxonomial of f contains x_j, so $\deg(f_\alpha) < d$ for both $\alpha \in PA(\{j\})$.

If $k > 1$, then we note by the above observation that for either $\alpha \in PA(\{j\})$, $deg(f) = d$; furthermore, for any $S \subseteq [n] - \{j\}$, S is a maxonomial hitting set of f_α iff $S \cup \{j\}$ is a maxonomial hitting set of f. Since there exists an S of size $k-1$ such that $S \cup \{j\}$ is a maxonomial hitting set of f, $h(f_\alpha) \leq k - 1$. However, if $h(f_\alpha) < k - 1$, this would imply the existence of a maxonomial hitting set of f with $< k$ elements, so $h(f_\alpha) = k - 1$. ☐

As a result, we see that every boolean f with $\deg(f) = 3$ and $h(f) = 1$ can be expressed as $x_j * f_1 + (1 - x_j) * f_0$ for some $j \in [n]$ and f_0, f_1 of degree at most 2 that are independent of x_j; in addition, for $k > 1$, every boolean f with $\deg(f) = 3$ and $h(f) = k$ can be expressed as $x_j * f_1 + (1 - x_j) * f_0$ for some $j \in [n]$ and f_0, f_1 of degree at 3 and maxonomial hitting set size $k-1$ that are independent of x_j. Consequentially, if we know the set of all boolean functions of degree at most 2, we can find the set of all boolean functions of degree 3 - and easily find $H(3)$ by determining when the process terminates.

At first blush, this seems computationally infeasible; however, we note that there are many ways to express what is essentially the same boolean function. We define two functions $f, g : \{0,1\}^n \to \{0,1\}$ to be **isomorphic** if there exists some permutation $\Xi : [n] \to [n]$ and subset $A \subseteq [n]$ such that $g(x_1, \ldots, x_n) = f(\alpha_1(x_{\Xi(1)}), \ldots, \alpha_n(x_{\Xi(n)})$ or $1 - f(\alpha_1(x_{\Xi(1)}), \ldots, \alpha_n(x_{\Xi(n)}))$, where $\alpha_i(t) = 1 - t$ if $i \in A$ and $= t$ otherwise. Our program will use the schematic outlined above to inductively find the set of all boolean functions f with $\deg(f) = 3$ and $h(f) = k$ for all $k \in \mathbb{N}$.

9.3.1 Finding All Functions for $k = 1$

The set of all boolean functions of degree 2 or less, up to isomorphism, is as follows:

$$0, x_1, x_1x_2, x_1 + x_2 - 2x_1x_2, x_1 - x_1x_2 + x_2x_3,$$

$$x_1 - x_1x_2 - x_1x_3 + x_2x_3, x_1 + x_2 - x_1x_2 - x_1x_3 - x_2x_4 + x_3x_4$$

In order to find, up to isomorphism, all boolean functions f with $\deg(f) = 3$ and $h(f) = 1$, we note that every such f can be expressed in the form $x_6f_1(x)+(1-x_6)f_0(x)$, where f_1 and f_0 are each isomorphic to one of the above. There are a number of techniques that we use in order to save time in our computations. We begin by noting that if $f(x_1, \ldots, x_n) = x_1f_1(x_2, \ldots, x_n) + (1 - x_1)f_0(x_2, \ldots, x_n)$, where f_1 and f_2 have degree 2, then every maxonomial of f (considered as a degree 3 function) is of the form $(c_1 - c_0)x_1x_ax_b$, where c_1 and c_0 are the coefficients of x_ax_b in f_1 and f_0 respectively.

We can sort the resulting boolean functions into two categories: those with at least 7 relevant variables, and those with at most 6. The number of isomorphism classes for such f with at least 7 relevant variables is small, since having so many relevant variables means that f_1 and f_0 share at most 2 relevant variables - i.e. at most 1 maxonomial. A list of one member of each such isomorphism class for f appears in ExcepPool.

All of the functions with at most 6 relevant variables can be assumed to be of the form $f(x_1, x_2, x_3, x_4, x_5, x_6)$. For these functions, the set of all possible sets of maxonomials, up to isomorphism, is relatively small. We sort the corresponding possible sets of maxonomials by the maximum of the absolute values of the coefficients of the maxonomials; WLOG, this maxonomial is $c_{126}x_1x_2x_6$, and c is positive. (Since $f : \{0, 1\}^n \to \{0, 1\}$, every monomial has an integral coefficient.)

If $c = 1$, then every coefficient of a maxonomial is ± 1, so the number of possible maxonomials of f up to isomorphism is very small. The list of all such maxonomials is listed in SchemataOne(x) - in particular, we note that every such set of maxonomials has at most seven members. If $c \geq 2$, then we note that one of the following must be true:

- x_1x_2 has a coefficient of 1 in f_1 and a coefficient of -1 in f_2.
- x_1x_2 has a coefficient of 2 in f_1 or a coefficient of -2 in f_2.

In either case, this reduces the number of possibilities for f_1 and f_0 to a number that can reasonably be found by hand; a list of almost all possible sets of maxonomials up to isomorphism are listed in SchemataTwo(x) and SchemataThree(x). Furthermore, all but two possible maxinomial sets in these lists can be expressed so that $x_2x_3x_6$ and $x_3x_4x_6$ have a coefficient of 0. As such, in listing out the above sets of maxonomials, we ensure that each member does not include these monomials. (The two exceptions are listed in ExcepPool(x).) In the case that the maxonomials exclude some variable in $\{x_3, x_4, x_5\}$, we include x_4 over x_5 over x_3.

We now produce all f with $d(f) = 3$ and $h(f) = 1$ (excluding those in ExcepPool) by taking all pairs of (potentially degenerate) degree 2 Boolean functions $f_1, f_0 : \{x_1, \ldots, x_5\}$ with matching coefficients on x_2x_3, as well as x_3x_4, such that the coefficient of x_1x_2 is more greater in f_1 than f_0. (We make certain to group them by their maxonomials.)

9.3.2 Finding All Functions for $k \geq 2$

We recall that, when $k \geq 2$, every boolean f with $\deg(f) = 3$ and $h(f) = k$ can be expressed as $x_jf_1 + (1-x_j)f_0$ for some $j \in [n]$ and f_0, f_1 of degree at 3 and maxonomial hitting set size $k-1$ that are independent of x_j. Consequentially, to find all such f, we need to consider all pairs f_1, f_0 with $h(f_1) = h(f_0) = k-1$. However, we can immediately eliminate most such pairs by the following proposition:

Proposition 9.11. *Let f, f_1, f_0 be defined as above. Then, f_1 and f_0 have the same set of maxonomials.*

Proof. Assume for the sake of contradiction that they do not; then, there exists some $x_ax_bx_c$ such that, if the monomial's coefficient in f_1 and f_0 are c_1 and c_0 respectively, then $c_1 \neq c_0$. Since the coefficient of $x_a x_b x_c x_j$ in f is $c_1 - c_0$, it must be nonzero - implying that $deg(f) > 3$ and creating a contradiction with the fact that $deg(f) = 3$. $\square$

As such, to find all such f, we only need to find it for all f_1, f_0 with the same maxonomial set. WLOG, we may assume that $j = k + 6$; furthermore, since we are looking at all f up to isomorphism, the lists of all f_1, f_0 up to isomorphism are almost entirely sufficient. However, there is one pitfall we need to note for finding all f.

While we only need to consider one set of maxonomials from a collection of isomorphic sets, it is possible that there are different f_0 and f_1 that are isomorphic - so they must be considered as different functions for the purposes of f. This may occur if f_0 is dependent on a variable x_o that doesn't appear in any maxonomial - we refer to such an x_o as an **orphaned variable**.

Proposition 9.12. *Let $f\{0,1\}^n \to \{0,1\}$, and x_o be an orphaned variable of f. Then, x_o appears in a monomial of degree ≥ 2.*

Proof. Assume for the sake of contradiction that x_o only appears in a monomial of degree 1 - i.e. a monomial of the form $c * x_o$ with $c \neq 0$. Then, $f_{x_o=0}$ and $f_{x_o=1}$ are boolean functions such that $f_{x_o=1} = f_{x_o=0} + c$. However, this can only happen if one of $f_{x_o=0}$ and $f_{x_o=1}$ is identically 0 and the other is identically 1, so $f = x_o$ or $1 - x_o$ - contradicting the fact that x_o is an orphaned variable. $\square$

Theorem 9.13. *Let $f : \{0,1\}^n \to \{0,1\}$ be a function with $deg(f) = 3$. Then, f has at most one orphaned variable, and this variable appears in a monomial of degree 2.*

Proof. We prove this by induction on $k = h(f)$. If $k = 1$, then by Proposition 9.10, $f = x_j f_1 + (1 - x_j) f_0$ for some $j \in [n]$ with $\{x_j\}$ a maxonomial hitting set of f, where f_1 and f_0 have degree at most 2. x_j cannot be an orphaned variable, so any such x_o must be a variable in f_1 and/or f_0. Assume for the sake of contradiction that two such orphaned variables x_o, x_o' exist; then, one of the following must be true:

- If x_o or x_o' (WLOG x_o) is a relevant variable in only one of f_1 and f_0 (WLOG f_0), then $f_0 = x_o$ or $1 - x_o$, and f_1 is independent of x_o. $deg(f_1) = 2$ (otherwise,

$deg(f) < 3$), so no variable in f_1 can be an orphaned variable by Proposition 9.12; however, x'_o cannot be a variable in f_0, and so f is independent of x'_o, creating a contradiction.

- If x_o and x'_o are relevant variables in both f_1 and f_0, then f_1 and f_0 must have the same degree 2 monomials in x_o and x'_o; however, by looking at all possible f_0 and f_1, the only way this can happen is if $f_0 = f_1$, so $f = f_0 = f_1$ and $deg(f) \leq 2$ - creating a contradiction.

As a result, f can have only one orphaned variable.

Now, suppose that the statement is true when $h(f) = k$ for a given $k \in \mathbb{N}$; we will show it is true when $h(f) = k + 1$. By Proposition 9.10, $f = x_j f_1 + (1 - x_j) f_0$ for some $j \in [n]$ and f_0, f_1 of degree 3 with $h(f_1) = h(f_0) = k - 1$ that are independent of x_j. If some x_o is an orphaned variable in only one of f_0 and f_1 (WLOG f_1), then f_0 is independent of x_o and by Proposition 9.12, $cx_o x_a$ is a monomial in f_1 for some $c \neq 0, a \in [n]$; thus, $cx_o x_a x_j$ is a monomial in f, and x_o is not orphaned there. By our inductive assumption, f_0 and f_1 each have at most one orphaned variable, and any variable that isn't orphaned in either isn't orphaned in f (since its maxonomials include the maxonomials of f_0 and f_1). As a result, f contains x_o as an orphaned variable iff f_1 and f_0 do, so f can only have one orphaned variable. By induction, we are done. $\square$

Since every degree 3 boolean function has at most one orphaned variable, and f only has an orphaned variable if f_0 and f_1 do, it is sufficient for our family of degree 3 functions f with $h(f) = 1$ to allow two different variables to be the orphan variable for isomorphic functions. Furthermore, every such f has at most 4 variables in its maxonomials, and there are only two possible maxonomial sets that allow f to have more than 3. For those two, we note that x_3 is the only possible orphan variable, so we add each such f with x_3 replaced by x_7; for the rest, x_3 and x_5 are already present as potential orphan variables.

9.3.3 Managing Runtime

When we sort our Boolean functions f with degree 3 and maxonomial hitting set size 1 by their maxonomials, we see that the most common maxonomial set by far is $x_1x_2x_6$. Computing $x_8f_1+(1-x_8)f_0$ for all such f_1, f_0 would be very time-consuming; however, we can save most of that time with the following theorem.

Theorem 9.14. *Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function such that $deg(f) = 3$ and $h(f) = 2$, such that for any j, if $h(f_{x_j=0}(x)) = 1$, then $f_{x_j=0}(x)$ has a maxonomial hitting set that is isomorphic to $x_1x_2x_6$. Then, f's maxonomial hitting set is isomorphic to $x_1x_2x_6 + x_1x_2x_8 + x_1x_6x_8 + x_2x_6x_8$.*

Proof. WLOG, we may assume there exists an $a \in [n]$ such that $x_1x_2x_6$ is the maxonomial hitting set of $f_{x_a=0}$; since the maxonomial set of $f_{x_a=0}$ is the set of all maxonomials of f that don't include x_a, every maxonomial of $f(x)$ either is $x_1x_2x_6$ or contains x_a (so $\{i,a\}$ is a maxonomial hitting set of $f(x)$ for all $i \in \{1,2,6\}$). Furthermore, since $\{i\}$ is not a maxonomial hitting set of $f(x)$ for any $i \in [n]$, for each $i \in \{1,2,6\}$, there must exist a corresponding maxonomial of $f(x)$ that does not include x_i as a variable.

Suppose that there exist two distinct $b_1, b_2 \in [n] - \{1,2,6,a\}$ such that for each $b \in \{b_1,b_2\}$, there exists a $d \in [n]$ such that $cx_ax_bx_d$ appears a maxonomial in $f(x)$ (with $c \neq 0$). This implies that cx_bx_d appears in $f_{x_a=0}(x)$ for each such b; however, by our initial condition on x_a, x_{b_1} and x_{b_2} do not appear in any maxonomial of $f_{x_a=0}$, and so both are orphaned variables in $f_{x_a=0}$. However, by Theorem 9.13, $f_{x_a=0}$ can have at most one orphaned variable, so we have a contradiction, and two such b_1, b_2 cannot exist.

Now, suppose that there exists a unique $b \in [n] - \{1,2,6,a\}$ such that $cx_ax_bx_d$ appears as a maxonomial in $f(x)$ with $c \neq 0$ and $d \in \{1,2,6\}$ (WLOG $d = 6$). Now, the maxonomial set of f must contain another maxonomial (otherwise, $\{6\}$ is a maxonomial hitting set of f, contradicting $h(f) = 2$); furthermore, if any other maxonomial of $f(x)$ excludes x_i for any $i \in \{1,2\}$, then x_i has the property that $f'(x) = f_{x_i=0}(x)$ has $h(f') = 1$ (since $\{a\}$ is now a maxonomial hitting set of f'), and the maxonomial set of f' is not isomorphic to $x_1x_2x_6$. Consequentially, $f(x)$has $x_1x_2x_6+cx_ax_bx_6+c'x_1x_2x_a$ as its maxonomial hitting set for some $c' \neq 0$; however, this means that $f'(x) = f_{x_b=0}(x)$ has $h(f') = 1$ with the maxonomial hitting set not isomorphic to $x_1x_2x_6$, creating a contradiction, and so no such b can exist.

As a result, we note that the maxonomial hitting set of $f(x)$ is $x_1x_2x_6 + cx_1x_2x_a + c'x_1x_6x_a + c''x_2x_6x_a$ for some $c, c', c'' \neq 0$. As as result, $h(f_{x_i=0}(x)) = 1$ for all $i \in \{1, 2, 6, a\}$, and so $c, c', c'' = \pm 1$ by the condition on all such restrictions of f. Any such f has a maxonomial hitting set that is isomorphic to $x_1x_2x_6 + x_1x_2x_8 + x_1x_6x_8 + x_2x_6x_8$. □

This implies that when f_1 and f_0 both have $x_1x_2x_6$ as their set of maxonomials, we only need to consider pairs such that $x_8f_1 + (1 - x_8)f_0$ has $x_1x_2x_6 + x_1x_2x_8 + x_1x_6x_8 + x_2x_6x_8$ as its maxonomial hitting set - greatly reducing the runtime.

By the inductive process followed above, we find, up to isomorphism, every Boolean function f with $deg(f) = 3$ and $h(f) = k$ for $k = 2, 3, 4, 5$. The set of Boolean functions that we find when $k = 4$ is nonempty, and includes the function $contra(x)$ defined above; however, the set that we find for $k = 5$ is empty. This implies that no Boolean f with $deg(f) = 3$ and $h(f) > 5$ exists (by a simple induction argument, using that fact that any such f could be expressed as $x_jf_1(x) + (1 - x_j)f_0(x)$ for some f_1, f_0 with $deg(f_1) = deg(f_0) = 3$ and $h(f_1) = h(f_0) = k - 1$).

Appendix A

A Clarification of Gusfield

It is not always immediately obvious what the stable matchings that contain (m, d) are, or even if any do. Gusfield ([GI89], Section 2.2.2) states that "it is easy to test if there is a stable matching containing (m, d), and if so, to find $M(m, d)$. Simply modify the Gale-Shapley algorithm so that d rejects all proposals from anyone other than m, and such that no student other than d accepts a proposal from m." In this appendix, we disambiguate Gusfield's statement, and generalize it to not only determine whether (m, d) appears in a stable matching over I, but find a compact representation of every stable matching that contains (m, d).

We capture the structure of the stable matchings that contain (m, d) through the restriction $I^*_{(m,d)}$ of I, defined such that a given edge $(m', d') \in G(I^*_{(m,d)})$ iff either $(m', d') = (m, d)$, or all of the following conditions hold:

- $m' \neq m$ and $d' \neq d$.
- If d prefers m' to m, then m' prefers d' to d.
- If m prefers d' to d, then d' prefers m' to m.

We will typically shorten $I^*_{(m,d)}$ to I^* when (m, d) is implied.

In the case where $\mathcal{K}_e$ is nonempty, we note that this restriction is an example of a truncation $I_{(T_d, T_m)}$, where $T_m = \{(m, a(m)) : m \in V_m(I) \cap V\}$ and $T_d = \{(a(d), d) : d \in V_m(I) \cap V\}$. For the case of I^*, we note that $a(v)$ is as follows:

- $a(m) = d$ and $a(d) = m$.
- For all $m' \in V_m(I) - \{m\}$, if d prefers m' to m, then $a(m')$ is the element on m''s preference list directly above d; otherwise, $a(m')$ is the last element on m''s

preference list.

- For all $d' \in V_d(I) - \{d\}$, if m prefers d' to d, then $a(d')$ is the element on d''s preference list directly above m; otherwise, $a(d')$ is the last element on d''s preference list.

Theorem A.1. *For a given satisfactory instance I and edge $(m, d) \in G(I)$, let V_0 be the set of vertices covered by the stable matchings over I, and M be any matching such that $(m, d) \in M$ and the edges of M cover V_0. Then, M is a stable matching over I iff $M \subseteq G(I^*)$ and is a stable matching over I^*.*

Proof. If M is stable over I, then M cannot contain any edge not present in I^* - the presence of (m, d) in M tells us that there is no other edge in the matching containing either vertex, and if M contains some $(m', d') \notin G(I^*)$ with $m' \neq m$ and $d' \neq d$, then via the definition of I^*, we see that either (m', d) or (m, d') destabilizes M in I. Furthermore, M must be stable in I^* - if it wasn't, the edge (m', d') that destabilizes M over I^* would also destabilize M over I.

Now, suppose that $M \subseteq G(I^*)$ and is a stable matching over I^*; we assume for the sake of contradiction that M is not stable over I. As a result, there must exist an edge $(m_0, d_0) \in G(I)$ that destabilizes M over I.

- If $m_0 = m$ and $d_0 = d$, then (m_0, d_0) is in M, so it can't destabilize M.

- If $m_0 = m$ and $d_0 \neq d$, then m prefers d_0 to d and d_0 prefers m to $p_M(d_0)$. This means, by definition of I^*, that $(p_M(d_0), d_0) \notin G(I^*)$, so $M \subsetneq G(I^*)$, creating a contradiction.

- If $m_0 \neq m$ and $d_0 = d$, then d prefers m_0 to m and m_0 prefers d to $p_M(m_0)$. This means, by definition of I^*, that $(m_0, p_M(m_0)) \notin G(I^*)$, so $M \subsetneq G(I^*)$, creating a contradiction.

- If $m_0 \neq m$, $d_0 \neq d$, and $(m_0, d_0) \in G(I^*)$, then the fact that M is stable over I^* tells us that either m_0 prefers $p_M(m_0)$ to d_0 or d_0 prefers $p_M(d_0)$ to m_0; in either case, this tells us that no such (m_0, d_0) can destabilize M.

- If $m_0 \neq m$, $d_0 \neq d$, and $(m_0, d_0) \notin G(I^*)$, then either m_0 prefers d to d_0 and d prefers m_0 to m, or d_0 prefers m to m_0 and m prefers d_0 to d. In the former case, the fact that $(m_0, p_M(m_0)) \in G(I^*)$ and d prefers m_0 to m means that m_0 prefers $p_M(m_0)$ to d, so by the transitive property, m_0 prefers $p_M(m_0)$ to d_0. In the latter case, the fact that $(p_M(d_0), d_0) \in G(I^*)$ and m prefers d_0 to d means that d_0 prefers $p_M(d_0)$ to m, so by the transitive property, d' prefers $p_M(d_0)$ to m_0. Either way, we see that (m_0, d_0) cannot destabilize M.

Since we have a contradiction for every possible configuration of (m_0, d_0), there cannot be any such destabilizing edge. Therefore, M is stable over I. □

Corollary A.2. *Let V_0 be the set of vertices covered by the stable matchings over I. Then, the set of all stable matchings over I that include (m, d) is the set of all stable matchings over I^* that cover V_0.*

Proof. By Theorem 2.4, every stable matching over I covers V_0; therefore, by Theorem A.1, every stable matching over I that contains (m, d) is a stable matching over I^*, and continues to cover V_0. Similarly, every stable matching over I^* that covers V_0 is also a stable matching over I by Theorem A.1. Since every member of one set is part of the other, the two sets are the same. □

As such, we have reduced the problem of finding the poset $\mathcal{K}_e$ of all stable matchings that include a given edge to the problem of finding the set of all perfect stable matchings for a different instance. In particular, there exists a stable matching over I that includes (m, d) iff the stable matchings over I^* are perfect. We also note that the corollary of Theorem A.1 implies Theorem 5.25.

Appendix B

Proof of Lemma 4.10

As noted previously, Lemma 4.10 is not unique to this paper, and a lemma that uses the same reasoning appears in [Wak08]. However, we discovered it independently and only later discovered Wako's presentation. In this section, we will show that if J and K are any two subsets of E such that $J \subseteq K$, $\psi(J) = K$, and $\psi(K) = J$, then $J = K$.

B.1 The Association Partition

Our basic strategy to show that $J = K$ is by contradiction. We note that the K-stable matchings form a distributive lattice $\mathcal{L}_K$ by Theorem 3.10. If $K - J$ is nonempty, we can associate each edge of $K - J$ with an element of $P(\mathcal{L}_K)$ in such a way that, given an element $v \in P(\mathcal{L}_K)$ with at least one edge of $K - J$ associated with it, we can construct a K-stable matching using at least one edge associated with v; however, this creates a contradiction with the initial condition that $\psi(K) = J$, implying that every K-stable matching consists entirely of edges in J.

Proof. Since $\psi(K) = J \subseteq K$, by Theorem 3.10, the set of matchings $\mathcal{M}_K$ that are stable with respect to K can be placed under the distributive lattice structure $\mathcal{L}_K = (\mathcal{M}_K, \preceq)$. This in turn allows us to construct the poset of $P(\mathcal{L}_K)$ of join-irreducible elements of $\mathcal{L}_K$; by our previous observations, the elements of $P(\mathcal{L}_K)$ correspond to the rotations over $I[K]$. Let us define P' as the poset created by adding two additional elements to $P(\mathcal{L}_K)$ - $\hat{0}$, which is set to be less than all other elements in P', and $\hat{1}$, which is set to be greater than all other elements in P'. We also set $\hat{0} = M_m$, the mentor-optimal K-stable matching. (We note that the property from $\mathcal{L}_k$ that M_m dominates every element of $P(\mathcal{L}_K)$ is also preserved in P'.) We will construct a mapping $\nu : K - J \to P'$.

Now, consider any $e \in K - J$. Since $e \in K = \psi(J)$, there exists a matching M_e that J-stable and includes e. Now, consider any matching M' that is K-stable. In particular, since $E(M_e) \subseteq K$ and $E(M') \subseteq J$, M_e and M' are costable. By Theorem 3.1, $M_e \wedge_m M'$ and $M_e \wedge_d M'$ are the same matching, and so m_e prefers d_e to their partner in M' iff d_e prefers their partner in M' to m_e. (The order of preference in this case is always strict, because $e \notin M'$.)

We now consider the sublattice $\mathcal{L}_e^*$ of K-stable matchings M' such that d_e prefers their partner in M' to m_e. If this sublattice is empty, we define $\nu(e) = \hat{1}$. Otherwise, m_e prefers d_e to a nonempty subset of their possible partners in $\mathcal{L}_K$, and so the sublattice $\mathcal{L}_e^*$ of K-stable matchings M' such that d_e prefers their partner in M' to m_e is a nonempty sublattice of $\mathcal{L}_K$; as such, we may consider the mentor-optimal matching M_0 of $\mathcal{L}_e^*$ as the meet of every element of this sublattice. By Theorem 3.10, M_0 is also in $\mathcal{L}_e^*$, and either equals M_m or is a join-irreducible of $\mathcal{L}_K$. Either way, we see that M_0 is an element of P', and set $\nu(e) = M_0$. (Note that d_e prefers their partner in M_0 to m_e, and every K-stable matching M' with the same property is $\geq M_0$.)

We say that an edge $e \in K - J$ is **associated** with a vertex $v \in P'$ if $\nu(e) = v$ - in particular, every $e \in K - J$ is associated with some $v \in P'$.) However, we can show the following lemma:

Lemma B.1. *For any vertex $v \in P'$, $\nu^{-1}(v) = \emptyset$.*

Since any $e \in K - J$ must be associated with some vertex of P', no such e can exist. Therefore, $K \subseteq J$; since $J \subseteq K$ from our initial constraints on J and K, $J = K$. $\square$

B.2 Proof of Lemma B.1

In proving Lemma B.1, it is easiest to consider it as two separate sublemmas.

Lemma B.2. *Let v be any vertex of $P(\mathcal{L}_K)$ other than $\hat{1}$. Then, $\nu^{-1}(v) = \emptyset$.*

We hold off on the proof of this lemma for the time being.

Lemma B.3. $\nu^{-1}(\hat{1}) = \emptyset$.

Proof. Consider an instance I' created from I by reversing which vertices are mentors and which are students; J and K retain the property of mapping to each other via $\psi_{I'}$. By applying *Lemma B*.2 to I' with $v = \hat{0}$, no edge $e \in K - J$ can have the property that, for every stable matching M over I', d_e prefers m_e to $p_M(d_e)$ and m_e prefers $p_M(m_e)$ to d_e. (Recall that in I', d_e is a mentor and m_e is a student.) However, this property must continue to hold in I (since every vertex has the same preference list in I and I'). By the definition of ν, this means that no edge $e \in K - J$ can be in $\nu^{-1}(\hat{1})$. □

We now set out to prove Lemma B.2.

Proof. By the properties of P' stated in the proof of Lemma 4.10, if $v \neq \hat{1}$, $v \equiv M_0$ is a K-stable matching with the property that d_e prefers their partner in M_0 to m_e, and every K-stable matching M' with the same property is dominated by M_0. WLOG, let us assume that $M_0 = \{(m_1, d_1), (m_2, d_2), \ldots, (m_n, d_n)\}$, and for the sake of contradiction, $\nu^{-1}(M_0)$ is nonempty; for each such edge $e = (m_i, d_j) \in \nu^{-1}(M_0)$, m_i prefers d_j to d_i, and d_j prefers m_j to m_i. We seek to construct a K-stable matching M^* that dominates M_0 and includes at least one edge in $\nu^{-1}(M_0)$, by replacing some edge in M_0 with new edges. To this end, we create a directed graph D that represents the edges in K that we consider as candidates for M^*.

For each student d_j, if d_j appears as a vertex in some nonzero number of edges associated with v, we define $\chi(j)$ to be the mentor in these edges that appears first in d_j's preference list. If $M_0 = M_m$, this completes our definition of χ. For any other possible v, we note that, in $\mathcal{L}_K$, M_0 covers a unique matching M_1, and M_1 differs from M_0 by a rotation; WLOG, we may assume that:

$$M_1 = \{(m_1, d_r), (m_2, d_1), \ldots, (m_r, d_{r-1}), (m_{r+1}, d_{r+1}), \ldots, (m_n, d_n)\}$$

for some $2 \leq r \leq n$. (In addition, since $M_1 \not\geq M_0$, for every edge e associated with v, d_e prefers m_e to their partner in M_1, and m_e prefers their partner in M_1 to d_e.) For every $j \leq r$ that is otherwise undefined, we define $\chi(j)$ to be d_j's partner in M_1.

If we set α to be the set of all j such that $\chi(j)$ is defined, we can construct a directed graph D with vertex set $[n]$ and edge set $\{(j, \chi(j)) : j \in \alpha\}$. (The existence of an edge

$(j, i) \in D$ implies that $(m_i, d_j) \in K$.) We note that each vertex in D has outdegree at most 1; however, some vertices - corresponding to students that do not appear in any edge associated with v or in any edge that appears in the rotation between M_0 and M_1 - can have outdegree 0.

Proposition B.4. *Suppose $M_0 \neq M_m$. Then, for every vertex $i \in D$ such that $i > r$, i has outdegree and indegree 0.*

Proof. Let $e = (j, i)$ be any edge in D. By the definition of D, m_i prefers $p_{M_1}(m_i)$ to d_j, and strictly prefers d_j to $p_{M_0}(m_i)$; furthermore, d_j strictly prefers $p_{M_0}(d_j)$ to m_i, and prefers m_i to $p_{M_1}(d_j)$. This implies that m_i strictly prefers $p_{M_1}(m_i)$ to $p_{M_0}(m_i)$, and d_j strictly prefers $p_{M_0}(d_j)$ to $p_{M_1}(d_j)$. This only can occur if $i, j \leq r$; consequentially, if $i > r$, it has indegree and outdegree 0 in D. $\square$

Lemma B.5. *If a vertex $i \in D$ has indegree ≥ 1, then it has outdegree 1.*

Proof. Suppose $M_0 \neq M_m$, and the vertex $i \in D$ has indegree ≥ 1. By Proposition B.4, $i \leq r$; by the definition of χ, each such i has outdegree 1.

Now, suppose that $M_0 = M_m$, and the vertex $i \in D$ has indegree ≥ 1. This implies the existence of an edge $(m_i, d_j) \in K$ such that m_i prefers d_j to d_i. Since $K = \psi(J)$, there exists a J-stable matching M' that contains (m_i, d_j), and $M' \subseteq K$; since M_0 is K-stable, it is $\subseteq J$, and therefore, M_0 and M' are costable. By Proposition 3.6, the fact that m_i prefers $p_{M'}(m_i)$ to $d_i = p_{M_0}(m_i)$ implies that d_i prefers m_i to $p_{M'}(d_i) \equiv m_k$, which implies that m_k prefers d_i to d_k. As a result, there exists a mentor m_k that prefers d_i to d_k, so the vertex $i \in D$ has outdegree 1. $\square$

If we assume that there exists a vertex in D with outdegree 1, then we may create a sequence $\{i_1, i_2, \ldots\}$ where i_1 is a vertex $\in [n]$ with outdegree 1 and $i_{k+1} = \chi(i_k)$ for all $k \geq 1$. We know that $\chi(i_1)$ exists (since i_1 has outdegree 1), so i_2 is well defined. Meanwhile, for any $k > 1$, $i_k = \chi(i_{k-1})$, and so has indegree ≥ 1; by the contrapositive of the lemma above, this means that it has outdegree 1, and so i_k being well-defined implies that i_{k+1} is well-defined. By induction, we see that the entire sequence is well-defined.

Since this is an infinite sequence over a finite domain, there must be some term i_b that equals a previous term i_l. Now, consider the matching M^* such that d_{i_k} is matched with $m_{i_{k+1}} = m_{\chi(i_k)}$ for all $k \in \{l, l+1, \ldots, b-1\}$ and d_i is matched with m_i for all $i \notin \{i_l, i_{l+1}, \ldots, i_{b-1}\}$. Since every edge of the form $(m_{\chi(i)}, d_i)$ has the property that $m_{\chi(i)}$ prefers d_i to $d_{\chi(i)}$ and d_i prefers m_i to $m_{\chi(i)}$, M^* dominates M_0. Furthermore, if $M_0 \neq M_m$, then $m_{\chi(i)}$ prefers $p_{M_1}(m_{\chi(i)})$ to d_i and d_i prefers $m_{\chi(i)}$ to $p_{M_1}(d_i)$, so M_1 dominates M^*.

Lemma B.6. *M^* is K-stable.*

Proof. Assume for the sake of contradiction that M^* is not K-stable, so there exists an edge $\epsilon = (m_i, d_j) \in K$ such that M^* is not ϵ-stable - i.e. m_i and d_j prefer each other to their respective partners in M^*. Since M^* dominates M_0, m_i must still prefer d_j to their partner in M_0; however, since M_0 is K-stable, d_j must prefer their partner in M_0 to m_i.

If M_0 is the mentor-optimal K-stable matching, these two facts are sufficient to imply that ϵ is associated with M_0 (since the properties holding for the mentor-optimal K-stable matching imply that they hold for all K-stable matchings). Otherwise, M_1 dominates M^*, so d_j must still prefer m_i to their partner in M_1. However, since M_1 is K-stable, m_i must prefer their partner in M_1 to d_j. Consequentially, ϵ is associated with M_0, regardless of what M_0 is.

At least one edge associated with M_0 includes d_j (namely, ϵ), so $\chi(j)$ is the index of the mentor that is matched with d_j through an edge associated with v that appears first in d_j's preference list, and d_j weakly prefers $m_{\chi(j)}$ to m_i. By the definition of M^*, d_j is matched either with m_j or $m_{\chi(j)}$, and since $d_j = d_\epsilon$ prefers m_i to their partner in M^*, d_j is matched with m_j. However, d_j strictly prefers m_j to m_i, as (m_i, d_j) is associated with v, and thus d_j would prefer their partner in M_0. This creates a contradiction with the assumption that (m_i, d_j) destabilizes M^*, so our assumption must be false, and M^* is K-stable. $\square$

Since $\psi(K) = J$, this would imply that $M^* \subseteq J$; however, we can show that M^* contains at least one edge in $K - J$ - specifically, at least one such edge associated with

M_0.

Lemma B.7. *M^* contains at least one edge associated with M_0.*

Proof. M^* includes the edges $E^* := \{(m_{i_{k+1}}, d_{i_k}) : k \in \{l, l+1, \ldots, b-1\}\}$, none of which appear in M_0. If $M_0 = M_m$, then every edge in E^* is associated with v; otherwise, E^* consists of edges that are either associated with M_0 or in M_1.

For the sake of contradiction, assume that every edge in E^* is in M_1. As a result, every edge in D of the form (i_k, i_{k+1}) with $k \in \{l, l+1, \ldots, b-1\}$ corresponds to an edge from M_1, and so is in $\{(1,2),(2,3),\ldots,(r-1,r),(r,1)\}$. The only cycle that can be created from these edges requires every such edge; this can only exist as a cycle in D if $m_{\chi(i)} = p_{M_1}(d_i)$ for all $i \in [r]$. However, this implies that for every $i \leq r$, there is no edge associated with M_0 that includes d_i as a vertex. By Proposition B.4, for every $i > r$, there is no edge associated with M_0 that includes d_i as a vertex. These two observations together give us that no student can appear in an edge associated with M_0; this creates a contradiction with our assertion that at least one edge is associated with M_0, and so, by contradiction, M^* contains at least one edge associated with M_0. $\square$

We have thereby, given a vertex $v \neq \hat{1}$ with at least one edge $\in K - J$ associated with it, constructed a K-stable matching M^* that contains at least one edge in $K - J$. This creates a contradiction with $\psi(K) = J$, and so, by contradiction, Lemma B.2 must be true. $\square$

Appendix C

An Efficient Construction of ψ_I^∞

Previously, we proved that, for any given instance I, the hub-stable matchings over I form a distributive lattice $\mathcal{L}_K$ with $\vee$ and $\wedge$ as its join and meet functions respectively. This proof also provides a method to construct this lattice for a specific instance with n mentors and n students - generate ψ_I^∞ by computing the sequence $\{E(I), \psi(E(I)), \psi^2(E(I)), \ldots\}$, then finding the lattice of stable matchings over the limit of this sequence. This algorithm finds ψ_I^∞ in $O(n^3)$ time. However, as seen in Theorem 2.24 ([Wak10]), Jun Wako determined that there exists an algorithm that produces a description of the lattice of hub-stable matchings (and thereby the hub) in $O(n^2)$ time.

We independently discovered an algorithm that finds ψ_I^∞ in $O(n^3)$ time. This algorithm follows the following strategy:

1. Generate the mentor-optimal hub-stable matching M_0 and the student-optimal hub-stable matching M_1.

2. Consider the instance $I_{(M_0,M_1)}$. Then, the hub of I is the union of all stable matchings over $I_{(M_0,M_1)}$.

Theorem C.1. *Let M_0 and M_1 be the mentor-optimal and student-optimal hub-stable matchings respectively. Then, the hub of I is the union of all stable matchings over $I^* = I_{(M_0,M_1)}$.*

Proof. Over I^*, M_0 and M_1 are trivially the mentor-optimal and student-optimal hub-stable matchings (since M_0 matches each mentor with their top choice, and M_1 matches each student with their top choice); therefore, by Corollary 4.21, the hub of I^* is the union of all stable matchings over it.

By Corollary 4.31, $\psi_{I'}^{\infty} = \psi_I^{\infty} \cap G(I')$. By the definition of a truncation, $G(I')$ only excludes edges $e \in G(I)$ such that m_e strictly prefers $p_{M_1}(m_e)$ to d_e, or d_e strictly prefers $p_{M_0}(d_e)$ to m_e. If m_e strictly prefers $p_{M'}(m_e)$ to d_e, then $e \notin \psi_I^{\infty}$ - since M' is the student-optimal stable matching, every hub-stable matching has m_e partnered with a student they prefers to $p_{M_1}(m_e)$. Similarly, if d_e strictly prefers $p_{M_0}(d_e)$ to m_e, then $e \notin \psi_I^{\infty}$ - since M_0 is the mentor-optimal stable matching, every hub-stable matching has d_e partnered with a mentor they prefers to $p_{M'}(d_e)$. As a result, $G(I') \supseteq \psi_I^{\infty}$, and so $\psi_{I'}^{\infty}\psi_I^{\infty}$. □

Given $I_{(M_0,M_1)}$, we can generate the union of stable matchings over it in $O(n^2)$ time. Consequentially, the runtime of this algorithm is dependent on how efficiently we can find M_0 and M_1. We will present an algorithm that finds these matchings in $O(n^3)$ time; however, in [Wak10], Wako presents an algorithm that finds M_0 and M_1 in $O(n^2)$ time.

C.1 Generating the Mentor-Optimal Hub-Stable Matching

As an intermediate step in the generation of ψ^{∞}, we attempt to generate the mentor-optimal hub-stable matching without generating the sequence $\{\emptyset, \psi(\emptyset), \psi^2(\emptyset), \ldots\}$. One such algorithm is described in [Dig16]; we present the algorithm here, and prove that it produces the mentor-optimal hub-stable matching. (We note that while we did not discover the algorithm, our proof that it produces the mentor-optimal hub-stable matching is original. Digulescu also notes that this matching is the mentor-optimal hub-stable matching in the acknowledgments of [Dig19], which postdates our discovery of this fact.)

Algorithm C.2. *Given a satisfactory $n \times n$ instance I, we construct a perfect matching M_h over I as follows.*

1. *Set $t = n$ $I_n^* = I$, and $M_h = \emptyset$.*

2. *While $t > 0$, do the following:*

(a) *Let* $M \equiv M_{\{t\}}$ *be the mentor-optimal stable matching over* I_t^*. *Set* $I_t' = (I_t^*)_{(\emptyset,M)}$, *the faithful truncation of* I_t^* *restricted to edges* $(m_d) \in E(G(I_t^*))$ *such that* m *prefers* d *to* $p_M(m)$.

(b) *Let* $d_t \in V_d(I_t')$ *be a vertex in* $G(I_y')$ *with degree exactly* 1, *and* m_t *be the unique element of* $V_m(I_t')$ *such that* $(m_t, d_t) \in G(I_t')$. *(We note that such a* d_t *must exist - specifically, the last student proposed to in any operation of the Gale-Shapley algorithm on* I_t' *is such a* d_t.*) Set* $M_h = M_h \cup \{(m_t, d_t)\}$ *and* I_{t-1}^* *to be* I' *with the vertices* m_t *and* d_t *(and all edges incident to them) removed.*

(c) *Set* $t = t - 1$.

For $t \in [n]$, we define $M'_{\{t\}} = M_{\{t\}} \cup \{(m_k, d_k) : t < k \leq n\}$ and $I_t'' = I_{(\emptyset, M'_{\{t\}})}$.

Theorem C.3. *In Algorithm C.2,* $M_{\{t\}}$ *is a hub-stable matching for all* $t \in [n]$. *Furthermore, the perfect matching* M_h *constructed in Algorithm C.2 is the mentor-optimal hub-stable matching over* I.

Proof. We prove this result by strong induction on decreasing t - specifically, by showing, for all $2 \leq t \leq n$, if $M'_{\{t\}}$ is hub-stable, then $M'_{\{t-1\}}$ is hub-stable. For our base case, we note that $M'_{\{n\}} = M_{\{n\}}$ is the mentor-optimal stable matching over $I = I_n^*$, and so is hub-stable.

For our inductive step, since $M'_{\{t\}}$ is hub-stable, so by Theorem 4.35, we note that $\psi^\infty_{I_t''} = \psi^\infty_I \cap E(G(I_t''))$. In $G(I_t'')$, for all $i \geq t$, d_i has degree 1 and is incident with the edge (m_i, d_i). However, since $M'_{\{t\}}$ is a perfect stable matching over I_t'' (and thereby also hub-stable), every hub-stable matching over I_t'' is also perfect by Theorem 2.4. As a result, $\{(m_i, d_i) : t \leq i \leq n\}$ is a subset of every hub-stable matching over I_t'', and so $e \in$

$psi^\infty_{I_t''} \Rightarrow e \in S_t$, where $S_t = \{(m_i, d_j \in E(G(I_t'') : i = j \text{ or } i, j < t\}$.

As a result, $M'_{\{t-1\}}$ is thereby S_t-stable (since in any operation of the Gale-Shapley algorithm over $I_t''[S_t]$, m_i simply proposes to d_i for all $i \geq t$); this implies that $M'_{\{t-1\}}$ is hub-stable over I_t''. By Theorem 4.35, $M'_{\{t-1\}}$ is also hub-stable over I.

By induction, we see that if we define (m_1, d_i) to be the unique edge in $M_{\{1\}}$, $M'_{\{1\}} \equiv \{(m_i, d_i) : i \in [n]\}$ is hub-stable over I. To show that this is the mentor-optimal hub-stable matching, assume otherwise for the sake of contradiction; then, there exists a hub-stable matching over I that dominates M_h. By Theorem 4.35, this matching must also be hub-stable over $I_{(\emptyset, M_h)}$, and so $M_h \subset \psi^{\infty}_{I_{(\emptyset, M_h)}}$. However, by our inductive observations, $(m_i, d_j) \notin \psi^{\infty}_{I_{(\emptyset, M_h)}}$ if $i \neq j$ and $max(i, j) \geq 2$, so $\psi^{\infty}_{I_{(\emptyset, M_h)}} \subseteq M_h$. This creates a contradiction, so M_h is the mentor-optimal hub-stable matching over I. □

Theorem C.4. *We can run Algorithm C.2 in $O(n^3)$ time.*

Proof. Each iteration of step 2 can be run in $O(n^2)$ time. Given any satisfactory instance as I^*_t, we can find the mentor-optimal stable matching $M_{\{t\}}$, as well as m_t and d_t, in $O(n^2)$ time by using the Gale-Shapley algorithm. We also note that $E(I'_t)$ is the set of all (m, d) such that m proposes to d in the Gale-Shapley algorithm over I^*_t, and so can be found in $O(n^2)$ time as well; $E(I^*_{t-1})$ is just the set of all such edges where $m \neq m_t$.

Given that we run through step 2 n times, and the runtime of step 1 is trivial, we see that we can runAlgorithm C.2 in $O(n^3)$ time. □

We may also prove Theorem 2.25 at this juncture.

Proof. As noted in the proof of Theorem C.3, for all $i, j \in [n]$ such that $i < j$, m_i prefers $p_{M_{\{j\}}}(m_i)$ to d_j - otherwise, m_i would have proposed to d_j before $p_{M_{\{j\}}}(m_i)$. However, because $M_{\{j\}}$ is hub-stable over I and M_h is the mentor-optimal hub-stable matching, m_i prefers $p_{M_h}(m_i) = d_i$ to $p_{M_{\{j\}}}(m_i)$; therefore, m_i prefers d_i to d_j. □

Corollary C.5. *There exists an algorithm to construct the student-optimal hub-stable matching in $O(n^3)$ time.*

Proof. We may run Algorithm C.2, with the roles of the mentors and students switched. □

C.2 Extending to Nonsatisfactory Instances

The above algorithm for the construction of the lattice of hub-stable matchings is contingent on the instance being satisfactory; however, as noted in Corollary 4.23, any nonsatisfactory instance can be extended into a complete instance that preserves the behavior of ψ.

Theorem C.6. *For any $n' \times n''$ instance I, the lattice of hub-stable matchings can be constructed in $O(n^3)$ time, where $n = max(n', n'')$.*

Proof. If I is a satisfactory instance, then we can apply the above construction. Otherwise, let I' be any completion of I; since I' is a complete instance, we can determine $\psi_{I'}^{\infty}$ in $O(n^3)$ time. Thus, by Corollary 4.23, $\psi_I^{\infty} = \psi_{I'}^{\infty} \cap E(G(I))$ can be constructed in $O(n^3)$ time as well. Given ψ_I^{∞}, we can generate the lattice of hub-stable matchings on I in $O(n^2)$ time by finding the lattice of stable matchings on the instance generated from I by removing all edges not in ψ_I^{∞}. As a result, we can generate the lattice of hub-stable matchings on I in $O(n^3)$ time. $\square$

Appendix D

An Alternative Proof of Corollary 6.8

In this appendix, we look at an alternative proof of Corollary 6.8 which uses the properties of the hub ψ_I^∞ described in Theorem C.3. The results of this section are superceded by those of Section 6.2, and the techniques used are ultimately more complicated; however, it served as an important stepping stone in our discovery of the proof of Theorem 6.3.

We recall the following families of potential constraints:

Q_1**:** For all $(m,d) \in E(I)$, $wt(m,d) \geq 0$.

Q_2**:** For all $m \in V_m(I)$, $\sum_{d \in V_d(I)} wt(m,d) = 1$.

Q_3**:** For all $d \in V_d(I)$, $\sum_{m \in V_m(I)} wt(m,d) = 1$.

$Q_4(S)$**:** For all $(m,d) \in S$, $wt(m,d) + \sum_{d' <_m d} wt(m,d') + \sum_{m' <_d m} wt(m',d) \leq 1$.

$Q_5(S)$**:** For all $(m,d) \notin S$, $wt(m,d) = 0$.

By Theorem 6.2, we note that the polytope of fractional stable matchings over $I[\psi_I^\infty]$ - i.e. the polytope of fractional hub-stable matchings over I - can be constrained by $\{Q_1, Q_2, Q_3, Q_4(\psi_I^\infty), Q_5(\psi_I^\infty)\}$. However, in order to show that the constraints in $Q_5(\psi_I^\infty)$ are redundant, we will prove the following lemma:

Lemma D.1. *For any $wt \in P_h$ and edge $e \notin \psi_I^\infty$, $wt(e) = 0$.*

We prove this lemma via three sublemmas.

Lemma D.2. *Let P_h be defined as above, and M_0 be the mentor-optimal hub-stable matching for I. Then, for any $wt \in P_h$ and edge $e \in E(I)$ such that m_e strictly prefers d_e to $p_{M_0}(m_e)$, $wt(e) = 0$.*

Proof. By Theorem 2.25, we may set $V_m = \{m_1, m_2, \ldots, m_n\}$ and $V_d = \{d_1, d_2, \ldots, d_n\}$ such that $M_0 = (m_i, d_i) : i \in [n]$ is the mentor-optimal hub-stable matching and for all $i < j \leq n$, m_i prefers d_i to d_j. For the sake of contradiction, assume that there exists a $wt \in P_h$ and edge (m_k, d_j) such that m_k strictly prefers d_j to d_k, and $wt(m_k, d_j) > 0$; WLOG, we may assume that we select (m_k, d_j) such that k is minimized. (Note that by our choice of indexing, and the fact that m_k prefers d_j to d_k, $k > j$.)

We now consider the matching $M' \equiv M'_{\{k\}}$, as defined for Theorem C.3; we recall that $M' = M_{\{k\}} \cup \{(m_i, d_i) : k < i \leq n\}$, where $M_{\{k\}}$ is the mentor-optimal stable matching over $I[\{(m_i, d_{i'}) : i, i' \leq k\}]$. Theorem C.3 allows us to note the following:

1. M' is hub-stable.

2. $p_{M'}(m_k) = d_k$.

3. for all $i \in [k-1]$, $p_{M'}(m_i) = d_{i'}$ for some $i' \in [k-1]$, and each such $d_{i'}$ prefers m_i to m_k.

4. for all $i, i' \leq k$ such that m_i prefers $d_{i'}$ to $p_{M'}(m_i)$, $d_{i'}$ prefers $p_{M'}(d_{i'})$ to m_i (otherwise, (m_i, d'_i) would destabilize $M_{\{k\}}$ over $I[\{(m_i, d_{i'}) : i, i' \leq k\}]$).

Since $M' \subseteq \psi_I^\infty$, $Q_4(\{e\})$ is a constraint on P_h for every $e \in M'$, and so:

$$\sum_{d' \in V_d(I) : d' >_m d} wt(m, d') \geq \sum_{m' \in V_m(I) : m' <_d m} wt(m', d)$$

for all $(m, d) \in M'$. Adding these inequalities for all $(m, p_{M'}(m))$ such that $m \in \{m_i : i < k\}$ gives us that

$$\sum_{(m_i, d) : i < k, d >_{m_i} p_{M'}(m_i)} wt(m_i, d) \geq \sum_{(m, d_i) : i < k, m <_{d_i} p_{M'}(d_i)} wt(m, d_i).$$

(The right-hand side is summed over all d_i such that $d_i = p_{M'}(m_{i'})$ for some $i' < k$ - which, by item 3 above, is the set of all d_i such that $i < k$.)

By the assumption we made WLOG, every term in the former sum equals 0, so the latter sum equals 0 as well; however, because $wt : E(G(I)) \to [0, 1]$ only has non-negative outputs, $wt(m, d_i) = 0$ if $i < k$ and d_i strictly prefers $p_{M'}(d_i)$ to m. However, since m_k prefers d_j to d_k, then $j < k$ (by Theorem 2.25) and d_j strictly prefers $p_{M'}(d_j)$ to m_k - implying that $wt(m_k, d_j) = 0$. This contradicts the assumption that $wt(m_k, d_j) > 0$, so by contradiction, we see that if $wt \in P_I$ and m_k prefers d_j to d_k, then $wt(m_k, d_j) = 0$. □

Corollary D.3. *Let P_h be defined as above, and M_0 be the mentor-optimal hub-stable matching for I. Then, for any $wt \in P_h$ and edge $e \in E(I)$ such that d_e strictly prefers $p_{M_0}(d_e)$ to m_e, $wt(e) = 0$.*

Proof. Since $M_0 \subseteq \psi_I^\infty$, $\sum_{d \in V_d(I): d >_{m_i} d_i} wt(m_i, d) \geq \sum_{m \in V_m(I): m <_{d_i} m_i} wt(m, d_i)$ for all $i \in [n]$. Adding these inequalities for all $i \in [n]$, we see that:

$$\sum_{(m_i, d_j): d_j >_{m_i} d_i} wt(m_i, d_j) \geq \sum_{(m_i, d_j): m_i <_{d_j} m_j} wt(m_i, d_j).$$

By lemma D.2, every term in the former sum equals 0, so the latter sum equals 0 as well; however, because $wt : E(G(I)) \to [0, 1]$ only has non-negative outputs, $wt(m_i, d_j) = 0$ if d_j strictly prefers m_j to m_i. □

Corollary D.4. *Let P_I be defined as above, and M_1 be the student-optimal hub-stable matching for I. Then, for any $wt \in P_I$ and edge $e \in E(I)$ such that m_e prefers $p_{M_0}(m_e)$ to d_e, $wt(e) = 0$.*

Together, Lemma D.2, Corollary D.3, and Corollary D.4 tell us that Lemma D.1 holds iff it holds for every instance I where the mentor-optimal hub-stable matching matches each mentor with their top choice and each student with their bottom choice, and the student-optimal hub-stable matching matches each student with their top choice and each mentor with their bottom choice.

Lemma D.5. *Suppose that I is a satisfactory instance where the mentor-optimal hub-stable matching matches each mentor with their top choice and each student with their bottom choice, and the student-optimal hub-stable matching matches each student with their top choice and each mentor with their bottom choice. We define P_h as above. Then, for every $wt \in P_h$ and $e \in E(I) - \psi_I^\infty$, $wt(e) = 0$.*

Proof. We prove this result by induction on q, the number of edges in $G(I)$. For the base case, when $q = n$, $G(I)$ must be a perfect matching for I to be satisfactory, and the above holds.

Now, assume that the above is true for every instance I' such that $G(I')$ has less than q edges for some $q > n$; we will show that it is true for I such that $G(I)$ has q edges. As noted in our construction of ψ^∞, if $G(I)$ has $> n$ edges, there exists a sequence of mentors $m_1, m_2, \ldots, m_k$ and students $d_1, d_2, \ldots, d_k$ such that for each $i \in [k]$, m_i's top choice is d_i and their second choice is d_{i+1} (with the index taken mod k). (WLOG, we may assume that $V_m(I) = \{m_i : i \in [n]\}$, $V_d(I) = \{d_i : i \in [n]\}$, and each m_i's top preference is d_i.) We can confirm that the matching $M_2 \equiv \{(m_1, d_2), (m_2, d_3), \ldots, (m_{k-1}, d_k), (m_k, d_1), (m_{k+1}, d_{k+1}), \ldots, (m_n, d_n)\}$ is stable; as a result, for every edge (m, d) in $\{(m_1, d_2), (m_2, d_3), \ldots, (m_{k-1}, d_k), (m_k, d_1)\}$, P_I is constrained by $\sum_{d' \in V_d(I): d' >_m d} wt(m, d') \geq \sum_{m' \in V_m(I): m' <_d m} wt(m', d)$. Adding these inequalities together, we see that

$$\sum_{(m_i, d): i \in [k], d >_{m_i} p_{M_2}(m_i)} wt(m_i, d) \geq \sum_{(m, d_i): i \in [k], m <_{d_i} p_{M_2}(d_i)} wt(m, d_i).$$

However, the edges in the former sum are $\{(m_1, d_1), \ldots, (m_k, d_k)\}$, since M_2 matches every m_i with their second choice. These edges also appear in the latter sum (since each d_i has m_i at the bottom of their preference list), so by the non-negative condition on wt, this inequality is tight and $wt(m, d_i) = 0$ if $i \in [k]$, $m \neq m_i$, and d_i prefers m_{i-1} (mod k) to m. (Note that, by the construction of ψ_I^∞, no such edge appears in that set.)

In addition, for every $i \in [k-1]$, by the condition ascribed by (m_i, d_{i+1}), $wt(m_i, d_i) = wt(m_{i+1}, d_{i+1})$ - as such, there exists a constant C_{wt} such that $wt(m_i, d_i) = C_{wt}$ for all $i \in [k]$. As a result, we notice that $wt' \equiv wt - C_{wt} \cdot \{(m_i, d_i) : i \in [n]\} + C_{wt} \cdot M_2$ is also in P_I. Furthermore, wt' is a weight function on the faithful truncation $I' = I_{(M_2, \emptyset)}$. By Corollary 4.36, $wt' \in P_{I'}$, so by the inductive assumption, for every $e \in E(I')$ not in $\psi_{I'}^\infty$, $wt'(e) = 0$. This means that $wt(e) = 0$ for each such edge as well.

Since for every $e \notin \psi_I^\infty$, d_e either prefers m_e to $p_{M_2}(d_e)$ (in which case $e \in E(I') - \psi_{I'}^\infty$ by Corollary 4.36) or prefers $p_{M_2}(d_e)$ to m_e, we have shown that $wt(e) = 0$ for every such edge, and we are done. ☐

We now can prove Lemma D.1.

Proof. Let M_0 and M_1 be the mentor-optimal and student-optimal hub-stable matchings of I respectively; by Lemma D.2 and Corollary D.4, $wt(e) = 0$ unless $e \in S$, where S is the set of all edges such that m_e ranks d_e between $p_{M_0}(m_e)$ and $p_{M_1}(m_e)$, and d_e ranks m_e between $p_{M_1}(d_e)$ and $p_{M_0}(d_e)$. This implies that $P_I = P_{I[S]}$. By Lemma D.5, $wt \in P_{I[S]}$ implies that $wt(e) = 0$ for every $e \notin \psi_{I'}^{\infty}$. However, by Corollary 4.37, $\psi_{I[S]}^{\infty} \subseteq \psi_I^{\infty}$, and so $wt(e) = 0$ for every $e \notin \psi_I^{\infty}$. ☐

We are now able to prove Corollary 6.8.

Proof. In our description of P_I, we note by Lemma D.1, the conditions that $wt(m, d) = 0$ for all $(m, d) \notin \psi_I^{\infty}$ is implicitly enforced. However, by Theorem 6.2, this set of conditions is exactly the set of conditions on the convex hull of stable matchings of $I[\psi_I^{\infty}]$. By the definition of ψ_I^{∞}, the stable matchings of $I[\psi_I^{\infty}]$ are the hub-stable matchings of I, so we are done. ☐

References

[BDW02] Harry Buhrman and Ronald De Wolf, *Complexity measures and decision tree complexity: a survey*, Theoretical Computer Science **288** (2002), no. 1, 21–43.

[Bir37] Garrett Birkhoff, *Ring of sets*, Duke Mathematical Journal **3** (1937), no. 3, 443–454.

[Bir46] ———, *Tres observationes sobre el algebra lineal [three observations on linear algebra]*, Univ. Nac. Tucamán, Revista A **5** (1946), 147–151.

[Bla84] Charles Blair, *Every finite distributive lattice is a set of stable matchings*, Journal of Combinatorial Theory, Series A **37** (1984), 353–356.

[CHS20] John Chiarelli, Pooya Hatami, and Michael Saks, *An asymptotically tight bound on the number of relevant variables in a bounded degree boolean function*, Combinatorica **40** (2020), 237–244.

[Dan63] George Dantzig, *Linear programming and extensions*, Princeton University Press, 1963.

[Dig16] Mircea Digulescu, *Strategic play in stable marriage problem*, ArXiv preprint DOI:10.13140/RG.2.2.20331.75041 (2016).

[Dig19] ———, *Farsighted collution in stable marriage problem*, ResearchGate preprint DOI: 10.13140/RG.2.2.27615.10400 (2019).

[Ehl07] Lars Ehlers, *Von neumann-morgenstern stable sets in matching problems*, Journal of Economic Theory **134** (2007), 537–547.

[FI18] Yuval Filmus and Ferdinand Ihringer, *Boolean constant degree functions on the slice are juntas*, arXiv preprint arXiv:1801.06338 (2018).

[GI89] Dan Gusfield and Robert W. Irving, *The stable marriage problem*, The MIT Press, 1989.

[GIL87] Dan Gusfield, Robert W. Irving, and Paul Leather, *An efficient algorithm for the "optimal" stable marriage*, Journal of the ACM **34** (1987), no. 3, 532–543.

[GILS87] Dan Gusfield, Robert W. Irving, Paul Leather, and Michael Saks, *Every finite distributive lattice is a set of stable matchings for a small stable marriage instance*, Journal of Combinatorial Theory, Series A **44** (1987), 304–309.

[GS62] David Gale and Lloyd Shapley, *College admissions and the stability of marriage*, The American Mathematical Monthly **69** (1962), no. 1, 9–15.

[GS85] David Gale and Marilda Sotomayor, *Some remarks on the stable matching problem*, Discrete Applied Mathematics **11** (1985), 223–232.

[Gus87] Dan Gusfield, *Three fast algorithms for four problems in stable marriage*, 1987, pp. 111–128.

[HKP11] Pooya Hatami, Raghav Kulkarni, and Denis Pankratov, *Variations on the sensitivity conjecture*, Theory of Computing Library, Graduate Surveys **4** (2011), 1–27.

[IL86] Robert Irving and Paul Leather, *The complexity of counting stable marriages*, SIAM Journal on Computing **34** (1986), 655–667.

[Knu76] Donald E. Knuth, *Mariages stables*, Les Presses de l'Université de Montréal, 1976.

[MP88] Marvin Minsky and Seymour S. Papert, *Perceptrons*.

[NS94] Noam Nisan and Mario Szegedy, *On the degree of boolean functions as real polynomials*, Computational complexity **4** (1994), no. 4, 301–313.

[O'D14] Ryan O'Donnell, *Analysis of boolean functions*, Cambridge University Press, 2014.

[RS] Vladimir Retakh and Michael Saks, *On the rational relationships among pseudo-roots of a non-commutative polynomial*, Pending publication.

[Sig14] Mark Siggers, *On the representations of finite distributive lattices*, ArXiv preprint (2014).

[ST17] Igor Shinkar and Avishay Tal, 2017, Private communication.

[Tas12] Tamir Tassa, *Finding all maximally-matchable edges in a bipartite graph*, Theoretical Computer Science **423** (2012), 50–58.

[VV89] John Vande Vate, *Linear programming brings marital bliss*, Operations Research Letters **8** (1989), no. 3, 147–153.

[Wak08] Jun Wako, *A note on existence and uniqueness of vnm stable sets in marriage games*, ICALP 2008 Proceedings of Workshop: Matching Under Preferences-Algorithms and Complexity (2008), 157–168.

[Wak10] ———, *A polynomial-time algorithm to find von neumann-morgenstern stable matchings in marriage games*, Algorithmica **58** (2010), 188–220.

[Wel19] Jake Wellens, *A tighter bound on the number of relevant variables in a bounded degree boolean function*, arXiv preprint 1903.08214 (2019).