Stable roommates matchings, mirror posets, median graphs, and the local/global median phenomenon in stable matchings

Christine T. Cheng*
Department of Computer Science
University of Wisconsin–Milwaukee
ccheng@cs.uwm.edu

Anhua Lin
Department of Mathematical Sciences
Middle Tennessee State University
alin@mtsu.edu

September 16, 2009

Abstract

For stable marriage (SM) and solvable stable roommates (SR) instances, it is known that there are stable matchings that assign each participant to his or her (lower/upper) median stable partner. Moreover, for SM instances, a stable matching has this property if and only if it is a median of the distributive lattice formed by the instance's stable matchings.

In this paper, we show that the above local/global median phenomenon first observed in SM stable matchings also extends to SR stable matchings because SR stable matchings form a median graph. In the course of our investigations, we also prove that three seemingly different structures are pairwise duals of each other – median graphs give rise to mirror posets and vice versa, and mirror posets give rise to SR stable matchings and vice versa. Together, they imply that for every median graph G, there is an SR instance $I(G)$ whose graph of stable matchings is isomorphic to G. Our results are analogous to the pairwise duality results known for distributive lattices, posets, and SM stable matchings. Interestingly, they can also be inferred from the work of Feder in the early 1990's. Our constructions and proofs, however, are smoother generalizations of those used for SM instances.

1 Introduction

In the *stable roommates problem* (SR), there are $2n$ participants each of whom has a preference list that ranks all others in some linear order. A *matching* is a set consisting of n disjoint pairs of the participants. The matching is *unstable* if there is a pair of participants who prefer each other over their assigned partners in the matching; such a pair is said to *block* the matching. Intuitively, matchings with blocking pairs will likely unravel since there is always a temptation for the two participants who form a blocking pair to leave their partners in the matching and pair up. Hence, the goal of the problem is to find a *stable matching*, a matching with no blocking pairs.

A simpler version of SR, which is also sometimes referred to as the bipartite version of SR, is the *stable marriage problem* (SM). There are n men and n women each of whom ranks participants from the opposite sex only. This time, a stable matching is a set of n man-woman pairs such that no man and woman prefer each other over their assigned partners in the matching. An instance of SM can be transformed into an instance of SR by a simple trick: at the end of each man's preference list, append an arbitrary ordering of the other men; do the same for the women. It is straightforward to check that both instances have exactly the same stable matchings.

*Supported by NSF Award No. CCF-0830678.

Gale and Shapley [13] first proposed SM and SR in 1962. They showed that every SM instance has a stable matching and described a procedure that finds such a matching in $O(n^2)$ time. In contrast, they noted that there are SR instances with no stable matchings and left open the question of finding one if it exists. Two decades later, Irving [15] presented an $O(n^2)$ algorithm that can distinguish between the *solvable* SR instances – i.e., those with stable matchings – and the *unsolvable* ones: it outputs a stable matching for the solvable instances and reports none for the unsolvable instances. At least three books [20, 14, 25] and hundreds of papers have been written about SM, SR and their variants. Centralized stable matching algorithms also play an important role in society today as they are used to match medical residents to hospitals [24] and students to schools [1, 2].

By formulating SM and SR as linear programs, Teo and Sethuraman [27] discovered the surprising result that there are stable matchings which are locally fair – every participant is matched to their (lower or upper) median stable partner.

Theorem 1 *(Teo and Sethuraman) Let I be a solvable SR instance. Let $L = \{\mu_1, \mu_2, \ldots, \mu_l\}$ consist of distinct stable matchings of I. For each participant x, order his l partners in these matchings from his most preferred to his least preferred. Let $p_{i,L}(x)$ denote the ith partner of x in this sorted list.*

- *a. When l is odd, there is a stable matching of I that assigns each participant x to his median stable partner $p_{(l+1)/2,L}(x)$ in L.*
- *b. When l is even, there is a stable matching of I that assigns each participant x to $p_{l/2,L}(x)$ or $p_{l/2+1,L}(x)$, his lower or upper median stable partner in L.*
- *c. Additionally, when I is originally an SM instance, there is a stable matching that assigns each man m to $p_{i,L}(m)$, which simultaneously assigns each woman w to $p_{l-i+1,L}(w)$, for $i = 1, \ldots, l$.*

When I is an SM instance, the set of stable matchings of I, $M(I)$, forms a distributive lattice $\mathcal{M}(I) = (M(I), \preceq)$. The covering graph of $\mathcal{M}(I)$ (i.e., the undirected Hasse diagram of the distributive lattice) can then be thought of as the unordered structure that describes the connections between the stable matchings of I. Thus, the *distance* between two matchings μ and μ', $d(\mu, \mu')$, is the length of the shortest path connecting μ and μ' in this graph. A stable matching μ^* is a *median* of $\mathcal{M}(I)$ whenever $\sum_{\mu' \in M(I)} d(\mu^*, \mu') \leq \sum_{\mu' \in M(I)} d(\mu, \mu')$ for every $\mu \in M(I)$. That is, a median of $\mathcal{M}(I)$ is a stable matching whose total (or average) distance from all the stable matchings of I is the least. Medians of $\mathcal{M}(I)$ are arguably fair because they best represent all the stable matchings of I. Unlike the stable matchings in Theorem 1, however, their fairness is global in nature.

In Theorem 1(c), let $\alpha_{i,L}$ refer to the stable matching that matches each man m to $p_{i,L}(m)$ for $i = 1, \ldots, l$. Recently, Cheng [9] presented a characterization of these stable matchings that implied another surprising feature: when $L = M(I)$ and l is odd, $\alpha_{(l+1)/2,L}$ is the unique median of $\mathcal{M}(I)$. On the other hand, when l is even, the stable matchings μ such that $\alpha_{l/2,L} \preceq \mu \preceq \alpha_{l/2+1,L}$ are exactly the medians of $\mathcal{M}(I)$. Thus, quite remarkably, the two sets of fair stable matchings we have considered coincide: a stable matching is "locally" median (i.e., each person is matched to his median stable partner) if and only if it is "globally" median (i.e., the matching is a median of $\mathcal{M}(I)$). We shall call this the *local/global median phenomenon* of stable matchings.

In light of Theorem 1(a) and (b), a natural question to ask is whether the local/global median phenomenon extends to solvable SR instances. To answer, we need a structure for SR stable matchings that generalizes the distributive lattice of SM stable matchings. In [14], Gusfield and Irving described a way in which SR stable matchings can be viewed as having a meet semilattice structure. Their meet semilattice, however, is a bit "unnatural".[1] Furthermore, unlike the distributive lattice structure of SM stable matchings, very little is known about their semilattice. Thus, our objectives are three-fold: (1) find a more natural structure that governs SR stable matchings, (2) characterize the medians of the structure in a way that allows us to determine if the local/global median phenomenon holds for SR stable matchings, and (3) determine if the properties known about the distributive lattice of SM stable matchings are generalizable to the structure of the SR stable matchings.

To address our objectives, we take the traditional approach used by Gusfield and Irving in [14] to study the SR stable matchings; that is, we use the reduced rotation posets of the instance. Let I be a solvable SR instance. It is known that the reduced rotation poset of I, $\mathcal{R}'(I)$, encodes all the stable matchings of I (which can be exponentially large), and, yet, its size is always polynomial in the input size. We consider the general class of posets that includes $\mathcal{R}'(I)$ and call it the class of *mirror posets*. Here are our results:

- First, we prove that mirror posets give rise to *median graphs*. This implies that the set of stable matchings of I form a median graph $G(M(I))$. In this graph, two stable matchings are adjacent if and only if their encodings (with respect to the rotations in $\mathcal{R}'(I)$) differ in one rotation.

- We then show that the local/global median phenomenon also holds for the stable matchings of I. In particular, a stable matching of I matches each participant to his median stable partner if and only if the stable matching is also a median of $G(M(I))$. Applying results on medians in median graphs [4], we also make note of other nice properties that these median stable matchings possess.

- Next, we prove for any mirror poset $\mathcal{P}$, there is a solvable SR instance $I(\mathcal{P})$ so that the reduced rotation poset of $I(\mathcal{P})$ is isomorphic to $\mathcal{P}$.

- Finally, using a result of Barthélemy and Constantin [6], we prove a similar result for median graphs and mirror posets. That is, for any median graph G, there is a mirror poset $\mathcal{P}_G$ so that the median graph that arises from it, as noted in item 1, is isomorphic to G. Together with our third result, this implies that for every median graph G, there is a solvable SR instance I_G so that the graph of its stable matchings, $G(M(I_G))$, is isomorphic to G.

Our results strongly suggest that our proposed graph for the SR stable matchings is the appropriate generalization of the distributive lattice of SM stable matchings. When I is an SM instance, for example, $G(M(I))$ is the covering graph of $\mathcal{M}(I)$. The local/global median phenomenon for SR stable matchings holds when we use our graph as the global structure describing the SR stable matchings. The dualities between mirror posets and SR stable matchings, median graphs and mirror posets, and median graphs and SR stable matchings generalize the dualities between posets and SM stable matchings, distributive lattices and posets, and distributive lattices and SM stable matchings respectively [14]. Finally, because our graph is a median graph, it can be viewed

[1] Each stable matching μ is represented by P_μ, a set of ordered pairs (x, y) such that $\{x, y\}$ is in μ or x prefers y to his partner in μ. A stable matching μ_0 is then fixed. The elements of the semilattice consists of taking the symmetric difference between the P_{μ_0} and P_μ for each stable matching μ. A meet semilattice structure emerges because the elements are closed under intersection.

as a median semilattice [3], an ordered structure. This makes it possible to show that the meet semilattice proposed by Gusfield and Irving is isomorphic to our structure.

To a certain extent, our structural results on SR stable matchings are rediscoveries of some of Feder's results [10, 11, 12]. In the late 1980's/early 1990's, Subramanian [26] and Feder pioneered a different way of investigating SR by viewing SR stable matchings as the stable configurations of a network of non-expansive gates. This allowed Feder to prove the duality between 2-SAT (2-satisfiability) and SR instances, and the duality between median graphs and 2-SAT.

Theorem 2 *(Feder) For every SR instance with $2n$ people, its set of stable matchings can be described by a 2-SAT instance with $O(n^2)$ variables and clauses. The variables are of the form z_{xi} which denotes whether participant x is assigned to its first i choices. Conversely, every 2-SAT instance with $O(n)$ variables and $O(n^2)$ clauses characterizes the set of stable matchings of an SR instance with $O(n)$ participants.*

Theorem 3 *(Feder) For every median graph, there is a 2-SAT instance with no equivalent variables whose solutions are in one-to-one correspondence with the vertices of the graph. Conversely, for every 2-SAT instance with no equivalent variables, there is a median graph whose vertices are in one-to-one correspondence with the solutions of the 2-SAT instance.*

With a little bit of work (e.g., removing the equivalent variables in the 2-SAT representation of an SR instance), these two theorems can be combined to prove the duality between median graphs and solvable SR instances. Mirror posets and 2-SAT instances are also related. The (directed) implication graph of a solvable 2-SAT instance with no trivial and equivalent variables is a mirror poset. Conversely, mirror posets give rise to 2-SAT instances as shown in the proof of Theorem 4.3.4 in [14].

Interestingly, as far as we know, no one has taken advantage of the connections between median graphs and SR instances since Feder's work. In our opinion, this is due to the fact that the machinery he used is based on stable network configurations. For example, to determine his 2-SAT representation of the stable matchings of an SR instance, one would have to understand a more general algorithm that generates a 2-SAT representation of the stable configurations of an arbitrary network. Of course, there is nothing wrong with this, but it does make the transition from median graphs to SR instances and back harder to follow. The value of our paper then rests not only on the results but also the techniques we used. It serves as a bridge between the works of Feder [10, 11, 12] and Teo and Sethuraman [27] using the traditional approach of Gusfield and Irving [14]. In particular, our constructions are smooth generalizations of those used for SM, allowing the reader to have a better grasp of the results.

The rest of the paper is arranged as follows. In Section 2, we present a simple proof of Theorem 1 based on [18], describe Irving's algorithm for solving SR instances, and state how an SR instance's rotation poset encodes all its stable matchings. In Section 3, we show how mirror posets give rise to median graphs. In Section 4, we prove the local/global median phenomenon of stable matchings for solvable SR instances. We prove our duality result for mirror posets and SR stable matchings in Section 5, and the corresponding result for median graphs and mirror posets in Section 6. We conclude in Section 7.

2 Preliminaries

Klaus and Klijn [18] recently proved Theorem 1 without relying on linear programming. We present the proof here with some modifications.

Let I be an SR instance. From hereon, unless otherwise specified, we will always assume that such instances are solvable. Let $M(I)$ be its set of stable matchings. Given two stable matchings μ and μ' of I, a participant x *prefers* μ over μ' if x prefers his partner in μ over his partner in μ'. Here is a known property about stable matchings (Lemma 4.3.9 in [14]).

Lemma 1 *Let I be an SR instance where μ and μ' are two of its stable matchings. Suppose x and y are partners in μ but not in μ'. Then either x prefers μ over μ' and y prefers μ' over μ or vice versa; that is, x and y have opposite preferences over μ and μ'.*

The next lemma says that we can quantify the degree to which x and y disagree about their preference of μ over all other stable matchings of I.

Lemma 2 *Let $\pi = (\mu_1, \mu_2, \ldots, \mu_l) \in M(I)^l$. For each participant x, sort the multiset $\{p_{\mu_i}(x), 1 \leq i \leq l\}$ from x's most preferred to least preferred, and denote by $p_{i,\pi}(x)$ the ith person in this sorted list. For each x and for $i = 1, \ldots, l$, $y = p_{i,\pi}(x)$ if and only if $x = p_{l-i+1,\pi}(y)$.*

Proof Suppose x and y are partners in j of the stable matchings in π. Hence, there exists some i so that $y = p_{i,\pi}(x) = p_{i+1,\pi}(x) = \cdots = p_{i+j-1,\pi}(x)$. This means that there are $i-1$ stable matchings in π that x prefers more over those that match x to y, and $l-(i+j-1)$ stable matchings in π that x prefers less over those that match x to y. By Lemma 1, there are $i-1$ stable matchings in π that y prefers less over those that match y to x, and $l-(i+j-1)$ stable matchings in π that y prefers more over those that match y to x. Therefore, in y's sorted list of stable partners, $x = p_{l-i-j+2,\pi}(y) = p_{l-i-j+3,\pi}(y) = \cdots = p_{l-i+1,\pi}(y)$. Since we chose x, y, and i arbitrarily, the lemma follows. $\square$

The lemma leads to an easy proof of Theorem 1. Since the proof does not require the stable matchings in L to be distinct, we use π in place of L.

Proof of Theorem 1. When l is odd, $(l+1)/2 = l-(l+1)/2+1$. By Lemma 2, this means that the $2n$ participants of I can be partitioned into n pairs $\{x, y\}$ such that $x = p_{(l+1)/2,\pi}(y)$ and $y = p_{(l+1)/2,\pi}(x)$ (i.e., x and y are each other's median stable partners in π). Denote this perfect matching of the participants of I as M. Now, consider any pair of participants $\{x', y'\} \notin M$. Suppose x' prefers y' over $p_{(l+1)/2,\pi}(x')$. This implies that there are at least $(l+1)/2$ stable matchings in π where x' prefers y' over his assigned partners. But since these matchings are stable, y' must prefer his partners in these matchings over x'. Hence, there are at least $(l+1)/2$ stable matchings in π where y' prefers his assigned partners over x'; i.e., y' prefers $p_{(l+1)/2,\pi}(y')$ over x'. Thus, $\{x', y'\}$ cannot be a blocking pair of M so M is a stable matching of I.

When l is even, let $\pi' = (\mu_1, \mu_2, \ldots, \mu_{l-1})$. According to the previous paragraph, the matching M' that assigns each participant to their median stable partner in π' is a stable matching of I. But for each participant x, x's median stable partner in π' is his lower or upper median stable partner in π. Hence, I has a stable matching that assigns each participant to his lower or upper median stable partner in π.

Finally, when I is a stable marriage instance, every stable matching consists of man-woman pairs. For $i = 1, \ldots, l$, let $M_i = \{(m, w) : w = p_{i,\pi}(m)\}$. Since $m = p_{l-i+1,\pi}(w)$ whenever $(m, w) \in M_i$, no two men can be matched to the same woman. Consequently, the women also have distinct partners in M_i so M_i is a perfect matching of the men and women. The argument for proving M_i is stable is similar to the one for M above. Suppose $(m', w') \notin M_i$. If m' prefers w'

over his partner in M_i, then there are at least $l-i+1$ stable matchings in π where m' prefers w' over his assigned partner. Since these are stable matchings, w' must prefer her assigned partners in each case to m'. Thus, w' prefers $p_{l-i+1,\pi}(w')$ over m' so M_i has no blocking pairs. □

One might wonder why we are not extending our proof of Theorem 1(c) to stable roommates instances and, consequently, show that there is a stable matching that matches each participant x to $p_{i,\pi}(x)$ or $p_{l-i+1,\pi}(x)$ for $i=1,\ldots,l$. Such a stable matching must be a subset of the set $\{\{x,y\} : y = p_{i,\pi}(x) \text{ or } x = p_{i,\pi}(y)\}$. But the set may not contain a perfect matching of the participants because of the "non-bipartiteness" of stable roommates instances.

2.1 Solving SR Instances

Let us now describe Irving's algorithm that computes a stable matching of an SR instance if it has one. A more thorough discussion can be found in [14]. Throughout the algorithm, a *table* is associated with the instance. Initially, it consists of the preference lists of the participants. Subsequently, the lists are shortened until one of two terminating conditions is satisfied: some list becomes empty, which means that the instance has no stable matching, or all the lists contain exactly one entry, which corresponds to a stable matching of the instance. A table is always *consistent* – i.e., x is present in y's list if and only if y is present in x's list. Thus, when a pair $\{x,y\}$ is deleted from a table, two operations are always involved: x is removed from y's list and y is removed from x's list. For a table T and a participant x, we shall use $f_T(x)$, $s_T(x)$, $l_T(x)$ to denote the first, second and last persons on x's list in T.

The algorithm consists of two phases. The first phase is very similar to the Gale-Shapley algorithm for solving a stable marriage instance. All participants are initially set to be *free*. Some of them then become *semi-engaged*. However, the *semi-engagement* relation (as opposed to the *engagement* relation) is not necessarily symmetric; for example, if x is semi-engaged to y, it may not be the case that y is semi-engaged to x. Here is main loop: While some free person x has a non-empty list, he proposes to the first person y on his list. If some person z is semi-engaged to y, set z to be free. Assign x to be semi-engaged to y. Delete all pairs $\{x',y\}$ from the table such that x' is a successor of x in y's list.

It is known that all possible executions of phase 1 lead to the same table. We shall denote it as T_0, the *phase-1 table*. It has the following properties: (i) $y = f_{T_0}(x)$ if and only if $x = l_{T_0}(y)$, and (ii) the pair $\{x,y\}$ is absent from T_0 (i.e., x is not in y's list and vice versa) if and only if x prefers $l_{T_0}(x)$ over y or y prefers $l_{T_0}(y)$ over x. Hence, if the while loop ended because some free person's list became empty, the algorithm returns that the instance has no stable matching. Otherwise, if every person's list contains exactly one participant, the table corresponds to a perfect matching because of property (i), which is stable because of property (ii). The algorithm outputs the stable matching. Finally, if some person's list contains two or more participants, the algorithm proceeds to the second phase.

When a table T with properties (i) and (ii) above has at least one list with two or more participants, it always has an *exposed rotation* $\rho = (x_0,y_0),(x_1,y_1),\ldots,(x_{r-1},y_{r-1})$ such that $y_i = f_T(x_i)$ and $y_{i+1} = s_T(x_i)$ for $i=0,\ldots,r-1$, where addition is done modulo r. The *X-set of ρ* is $\{x_0,x_1,\ldots,x_{r-1}\}$ while its *Y-set* is $\{y_0,y_1,\ldots,y_{r-1}\}$. Eliminating ρ from T is similar to making x_i semi-engaged to y_{i+1} for $i=0,\ldots,r-1$; that is, all pairs $\{z,y_{i+1}\}$ such that z is a successor of x_i in y_{i+1}'s list are deleted for $i=0,\ldots,r-1$. The resulting table is denoted as T/ρ. It is easy to check that if T has properties (i) and (ii) in the previous paragraph, then T/ρ maintains these properties.

Phase 2 of the algorithm now has an easy description. Set T to T_0. While none of the lists in T are empty and some list has more than one entry, find a rotation ρ exposed in T and eliminate it. Update T to T/ρ. Once again, when the while loop ends and some list in T is empty, the algorithm outputs that the instance has no stable matching. Otherwise, all lists have exactly one entry; the algorithm returns the corresponding stable matching. Irving's algorithm can be implemented in $O(n^2)$ time.

Suppose a stable matching μ is obtained by eliminating the sequence $\rho_1, \rho_2, \ldots, \rho_t$ of rotations from T_0. It turns out that the order of elimination is not important as long as rotations are eliminated from a table only when they are exposed. Hence, we write $\mu = T_0/\mathcal{Z}$, where $\mathcal{Z} = \{\rho_1, \rho_2, \ldots, \rho_t\}$; i.e., we specify μ by stating the set of rotations that were eliminated from T_0 to obtain μ. Furthermore, $\mu = T_0/\mathcal{Z} = T_0/\mathcal{Z}'$ if and only if $\mathcal{Z} = \mathcal{Z}'$. Thus, if the rotations in $\mathcal{Z}$ and $\mathcal{Z}'$ were eliminated from T_0 during two separate executions of phase 2 and $\mathcal{Z} \neq \mathcal{Z}'$, the resulting stable matchings are different.

2.2 Rotations of Solvable Stable Roommates Instances

Suppose I is an SR instance. Let $R(I)$ consist of all the exposed rotations that can be eliminated during an execution of phase 2 of Irving's algorithm. A rotation $\rho = (x_0, y_0), (x_1, y_1), \ldots, (x_{r-1}, y_{r-1})$ in $R(I)$ is *non-singular* if $\bar{\rho} = (y_1, x_0), (y_2, x_1), \ldots, (y_{r-1}, x_{r-2}), (y_0, x_{r-1})$ is also in $R(I)$; otherwise, it is called *singular*. If ρ is non-singular, we say that $\bar{\rho}$ is the *dual* of ρ. Notice that the dual of $\bar{\rho}$ is ρ by the definition. Additionally, the X-set and Y-set of ρ are, respectively, the Y-set and X-set of $\bar{\rho}$. Here is an important result.

Theorem 4 *Let μ be a stable matching of I. Suppose $\mu = T_0/\mathcal{Z}$. Then $\mathcal{Z}$ contains every singular rotation and exactly one of each dual pair of non-singular rotations of I.*

Let $\rho, \sigma \in R(I)$. We say that $\sigma \leq \rho$ if for every sequence of rotation eliminations that lead to a stable matching in which ρ appears, σ appears before ρ. That is, σ has to be eliminated before ρ can be exposed in a table. It is easy to verify that $\leq$ is a partial order. The pair $\mathcal{R}(I) = (R(I), \leq)$ is called the *rotation poset* of I. Here are some properties known about the precedence relations between the rotations.

Lemma 3 *Let ρ and σ be non-singular rotations and τ be a singular rotation of I. Then*
(i) $\rho \not\leq \bar{\rho}$
(ii) $\sigma \leq \rho$ if and only if $\bar{\rho} \leq \bar{\sigma}$
(iii) any predecessor of τ is a singular rotation.

Together, Theorem 4 and Lemma 3 implies that every stable matching μ of I can be obtained from the table T_0 by eliminating all the singular rotations first and then the non-singular rotations associated with μ next. Define T_0' to be the table obtained by eliminating from T_0 all the singular rotations of I. Let $R'(I)$ consist of all the non-singular rotations of I, and let $\mathcal{R}'(I) = (R'(I), \leq)$ be the corresponding subposet of $\mathcal{R}(I)$. A subset S of $\mathcal{R}'(I)$ is *complete* if it contains exactly one of each dual pair of rotations of I. It is *closed* if whenever $\rho \in S$, every rotation that precedes ρ in $\mathcal{R}'(I)$ is also in S.

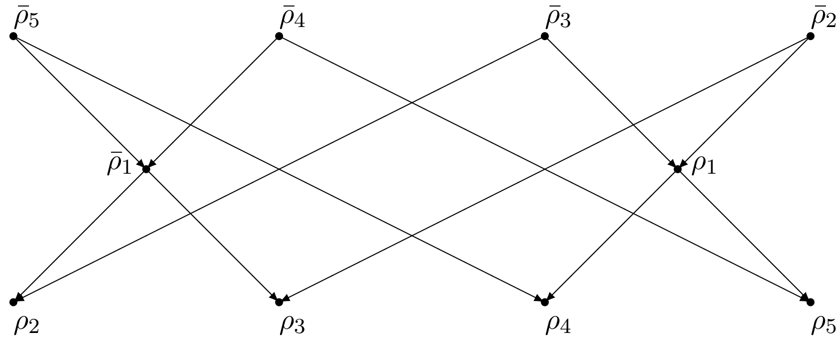

Figure 1: The Hasse diagram of the reduced rotation poset of the SR instance example.

Theorem 5 *Let I be an SR instance. There is a one-to-one correspondence between the stable matchings of I and the complete closed subsets of $\mathcal{R}'(I)$. In particular, if μ is the stable matching that corresponds to the complete closed subset S_μ of $\mathcal{R}'(I)$, then μ can be obtained from the phase-1 table by eliminating all the singular rotations of I and the rotations in S_μ.*

Constructing $\mathcal{R}'(I)$ takes $O(n^3 \log n)$ time. In an SR instance, a pair $\{x, y\}$ is *fixed* if x and y are partners in all the stable matchings of the instance. We state the next lemma, which is Lemma 4.4.1 in [14], because we will use it later.

Lemma 4 *In an SR instance,*

(i) $\{x, y\}$ is a fixed pair if and only if x's list in T_0' contains only y and y's contains only x;

(ii) otherwise, $\{x, y\}$ is a stable pair if and only if the pair (x, y) or the pair (y, x) is in a non-singular rotation of the instance.

2.3 An example

We now present an example to demonstrate the concepts we have just discussed. We will continually bring it up throughout the paper. Below is T_0, the phase-1 table of an SR instance with 20 participants. Observe that for each participant p, $f_{T_0}(p) = p'$ if and only if $l_{T_0}(p') = p$.

p_1:	p_8	p_{20}		p_{11}:	p_{17}	p_9	p_{10}	
p_2:	p_6	p_{19}		p_{12}:	p_9	p_{13}	p_7	
p_3:	p_{14}	p_{18}		p_{13}:	p_{10}	p_{12}	p_4	
p_4:	p_{13}	p_{17}		p_{14}:	p_{20}	p_{16}	p_{10}	p_3
p_5:	p_{19}	p_{15}	p_{16}	p_{15}:	p_7	p_5	p_8	
p_6:	p_{16}	p_8	p_2	p_{16}:	p_5	p_{14}	p_6	
p_7:	p_{12}	p_{15}		p_{17}:	p_4	p_{11}		
p_8:	p_{15}	p_6	p_1	p_{18}:	p_3	p_9		
p_9:	p_{18}	p_{11}	p_{12}	p_{19}:	p_2	p_5		
p_{10}:	p_{11}	p_{14}	p_{13}	p_{20}:	p_1	p_{14}		

The instance has no singular rotations but has five pairs of non-singular rotations which we describe below. Figure 1 shows its reduced rotation poset.

$\rho_1 = (p_{10}, p_{14}), (p_{12}, p_{13}), (p_{15}, p_7), (p_{16}, p_5)$ $\bar{\rho}_1 = (p_{13}, p_{10}), (p_7, p_{12}), (p_5, p_{15}), (p_{14}, p_{16})$
$\rho_2 = (p_9, p_{11}), (p_{13}, p_{12}), (p_{17}, p_4)$ $\bar{\rho}_2 = (p_{12}, p_9), (p_4, p_{13}), (p_{11}, p_{17})$
$\rho_3 = (p_{11}, p_9), (p_{14}, p_{10}), (p_{18}, p_3)$ $\bar{\rho}_3 = (p_{10}, p_{11}), (p_3, p_{14}), (p_9, p_{18})$
$\rho_4 = (p_{15}, p_5), (p_6, p_8), (p_{19}, p_2)$ $\bar{\rho}_4 = (p_8, p_{15}), (p_2, p_6), (p_5, p_{19})$
$\rho_5 = (p_{16}, p_{14}), (p_8, p_6), (p_{20}, p_1)$ $\bar{\rho}_5 = (p_6, p_{16}), (p_1, p_8), (p_{14}, p_{20})$

Here are its six stable matchings together with their corresponding complete closed subsets.

$$\begin{aligned}
\mu_1 &= T_0/\{\bar{\rho}_5, \bar{\rho}_4, \bar{\rho}_3, \bar{\rho}_2, \bar{\rho}_1\} \\
&= \{\{p_3, p_{18}\}, \{p_4, p_{17}\}, \{p_1, p_{20}\}, \{p_2, p_{19}\}, \{p_9, p_{11}\}, \\
&\quad \{p_{10}, p_{14}\}, \{p_{12}, p_{13}\}, \{p_5, p_{16}\}, \{p_6, p_8\}, \{p_7, p_{15}\}\} \\
\mu_2 &= T_0/\{\bar{\rho}_5, \bar{\rho}_4, \bar{\rho}_3, \bar{\rho}_2, \rho_1\} \\
&= \{\{p_3, p_{18}\}, \{p_4, p_{17}\}, \{p_1, p_{20}\}, \{p_2, p_{19}\}, \{p_9, p_{11}\}, \\
&\quad \{p_{10}, p_{13}\}, \{p_{12}, p_7\}, \{p_5, p_{15}\}, \{p_6, p_8\}, \{p_{14}, p_{16}\}\} \\
\mu_3 &= T_0/\{\bar{\rho}_5, \bar{\rho}_4, \bar{\rho}_3, \bar{\rho}_1, \rho_2\} \\
&= \{\{p_3, p_{18}\}, \{p_4, p_{13}\}, \{p_1, p_{20}\}, \{p_2, p_{19}\}, \{p_9, p_{12}\}, \\
&\quad \{p_{10}, p_{14}\}, \{p_{11}, p_{17}\}, \{p_5, p_{16}\}, \{p_6, p_8\}, \{p_7, p_{15}\}\} \\
\mu_4 &= T_0/\{\bar{\rho}_5, \bar{\rho}_4, \bar{\rho}_2, \bar{\rho}_1, \rho_3\} \\
&= \{\{p_3, p_{14}\}, \{p_4, p_{17}\}, \{p_1, p_{20}\}, \{p_2, p_{19}\}, \{p_9, p_{18}\}, \\
&\quad \{p_{10}, p_{11}\}, \{p_{12}, p_{13}\}, \{p_5, p_{16}\}, \{p_6, p_8\}, \{p_7, p_{15}\}\} \\
\mu_5 &= T_0/\{\bar{\rho}_5, \bar{\rho}_3, \bar{\rho}_2, \rho_1, \rho_4\} \\
&= \{\{p_3, p_{18}\}, \{p_4, p_{17}\}, \{p_1, p_{20}\}, \{p_2, p_6\}, \{p_9, p_{11}\}, \\
&\quad \{p_{10}, p_{13}\}, \{p_{12}, p_7\}, \{p_5, p_{19}\}, \{p_8, p_{15}\}, \{p_{14}, p_{16}\}\} \\
\mu_6 &= T_0/\{\bar{\rho}_4, \bar{\rho}_3, \bar{\rho}_2, \rho_1, \rho_5\} \\
&= \{\{p_3, p_{18}\}, \{p_4, p_{17}\}, \{p_1, p_8\}, \{p_2, p_{19}\}, \{p_9, p_{11}\}, \\
&\quad \{p_{10}, p_{13}\}, \{p_{12}, p_7\}, \{p_5, p_{15}\}, \{p_6, p_{16}\}, \{p_{14}, p_{20}\}\}
\end{aligned}$$

Below, the participants' stable partners in the six stable matchings were collected and sorted according to their preferences.

	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}
1	p_8	p_6	p_{14}	p_{13}	p_{19}	p_{16}	p_{12}	p_{15}	p_{18}	p_{11}
2	p_{20}	p_{19}	p_{18}	p_{17}	p_{15}	p_8	p_{12}	p_6	p_{11}	p_{14}
3	p_{20}	p_{19}	p_{18}	p_{17}	p_{15}	p_8	p_{12}	p_6	p_{11}	p_{14}
4	p_{20}	p_{19}	p_{18}	p_{17}	p_{16}	p_8	p_{15}	p_6	p_{11}	p_{13}
5	p_{20}	p_{19}	p_{18}	p_{17}	p_{16}	p_8	p_{15}	p_6	p_{11}	p_{13}
6	p_{20}	p_{19}	p_{18}	p_{17}	p_{16}	p_2	p_{15}	p_1	p_{12}	p_{13}

	p_{11}	p_{12}	p_{13}	p_{14}	p_{15}	p_{16}	p_{17}	p_{18}	p_{19}	p_{20}
1	p_{17}	p_9	p_{10}	p_{20}	p_7	p_5	p_4	p_3	p_2	p_1
2	p_9	p_{13}	p_{10}	p_{16}	p_7	p_5	p_4	p_3	p_2	p_1
3	p_9	p_{13}	p_{10}	p_{16}	p_7	p_5	p_4	p_3	p_2	p_1
4	p_9	p_7	p_{12}	p_{10}	p_5	p_{14}	p_4	p_3	p_2	p_1
5	p_9	p_7	p_{12}	p_{10}	p_5	p_{14}	p_4	p_3	p_2	p_1
6	p_{10}	p_7	p_4	p_3	p_8	p_6	p_{11}	p_9	p_5	p_{14}

When $\pi = (\mu_1, \mu_2, \mu_3, \mu_4, \mu_5, \mu_6)$, it is easy to verify that μ_1 and μ_2 are the only stable matchings that match participants to their lower or upper median stable partners in π.

3 From mirror posets to median graphs

Let us now define a class of posets based on Lemma 3, and a class of graphs called *median graphs*. In this section, we will prove that mirror posets give rise to median graphs.

Definition 1 *A poset $\mathcal{P} = (P, \leq)$ is a mirror poset if P can be partitioned into pairs of dual elements, where the dual of $\rho \in P$ is denoted as $\bar{\rho}$, such that*
(i) $\rho \not\leq \bar{\rho}$ for each $\rho \in P$ and
(ii) $\sigma < \rho$ if and only if $\bar{\rho} < \bar{\sigma}$ for any $\rho, \sigma \in P$.

Let $G = (V, E)$ be a graph, and let $d(u, v)$ denote the length of the shortest path from u to v in G. Suppose $\pi = (v_1, v_2, \ldots, v_l) \in V^l$. Define the *distance of vertex u from π* as

$$D(u, \pi) = \sum_{i=1}^{l} d(u, v_i).$$

A vertex $u^* \in V$ is a *median of* π if $D(u^*, \pi) \leq D(u, \pi)$ for each $u \in V$. When π consists exactly of the vertices of G, we shall also say that a median of π is a *median of* G. The *median set of* π consists of all the medians of π.

Definition 2 *A graph is a median graph if every family of three vertices always has a unique median. Equivalently, a graph is a median graph if for any three vertices u, v, w, there exists a unique vertex that lies in a shortest path from u to v, u to w, and v to w.*

An immediate consequence of the definition is that every median graph is connected and bipartite. Median graphs were first studied by Avann [3], and introduced independently by Nebeský [23] and Mulder and Schrijver [22]. Surveys by Klavžar and Mulder [19], Bandelt and Chepoi [5], and Knuth [21] summarizes the extensive research work that has been done in this area. Important examples of median graphs include trees and hypercubes.

Let $\mathcal{P} = (P, \leq)$ be a mirror poset. A subset $S \subseteq P$ is *partially complete* if it contains at most one element from each dual pair in P; it is *complete* if it contains exactly one element from each dual pair in P. It is *closed* if whenever $\rho \in S$ and $\sigma \leq \rho$ then $\sigma \in S$. We shall use $\mathcal{S}_\mathcal{P}$ to denote the set that consists of all the complete closed subsets of P in $\mathcal{P}$. For any subset $T \subseteq P$, an element $\rho \in T$ is a *minimal* element of T if none of its predecessors in $\mathcal{P}$ are in T; it is a *maximal* element of T if none of its successors in $\mathcal{P}$ are in T.

We construct a graph on $\mathcal{S}_\mathcal{P}$ as follows. Let $G(\mathcal{S}_\mathcal{P})$ be the graph whose vertex set is $\mathcal{S}_\mathcal{P}$, and whose edge set consists of pairs (S, S') if and only if they differ in only one dual element (i.e., there is an element $\rho \in P$ so that $\rho \in S$ and $\bar{\rho} \in S'$, and $S - \{\rho\} = S' - \{\bar{\rho}\}$.) We state several technical results that will be used to prove the main result of this section: $G(\mathcal{S}_\mathcal{P})$ is a median graph. While it suffices to show that every family of three vertices in $G(\mathcal{S}_\mathcal{P})$ has a unique median, we actually do more work and identify the medians of an arbitrary sequence of vertices in $G(\mathcal{S}_\mathcal{P})$ because the result will turn out to be useful in the next section.

Lemma 5 *Let $\mathcal{P} = (P, \leq)$ be a mirror poset. Suppose S is a partially complete and closed subset of P. Then there exists $T \in \mathcal{S}_{\mathcal{P}}$ so that $S \subseteq T$. In other words, S can be extended to a complete closed subset of $\mathcal{P}$.*

Proof Suppose S is a partially complete and closed subset of P. This means that there exists dual pairs $\{\rho_i, \bar{\rho_i}\}$, $i = 1, \ldots, k$ so that neither ρ_i nor $\bar{\rho_i}$ is part of S. Let ρ be a minimal element of $\{\rho_1, \bar{\rho_1}, \ldots, \rho_k, \bar{\rho_k}\}$. Suppose $\sigma \leq \rho$. Either $\sigma \in S$ or $\sigma \notin S$ because $\bar{\sigma} \in S$. Since $\mathcal{P}$ is a mirror poset, $\bar{\rho} \leq \bar{\sigma}$. Thus, if $\sigma \notin S$, $\bar{\rho}$ must be in S since S is closed. But this is a contradiction since neither ρ nor $\bar{\rho}$ are in S. Hence, $\sigma \in S$. This means that all the predecessors of ρ are in S so $S \cup \{\rho\}$ is closed. The set $S \cup \{\rho\}$ is also partially complete since S was partially complete and does not contain $\bar{\rho}$. When $S \cup \{\rho\} \notin \mathcal{S}_{\mathcal{P}}$, we again choose a minimal element from the set containing all dual pairs not in $S \cup \{\rho\}$ and add it to this set. By the same reasoning, the resulting set must be partially complete and closed. We repeat this until every dual pair has an element in the set. $\Box$

Lemma 6 *The graph $G(\mathcal{S}_{\mathcal{P}})$ is connected. Moreover, for $S, S' \in \mathcal{S}_{\mathcal{P}}$, $d(S, S') = |S' - S|$.*

Proof Let us first prove the following claim.

Claim: Let $S, S' \in \mathcal{S}_{\mathcal{P}}$. Suppose ρ is a minimal element of $S' - S$. Then $S - \{\bar{\rho}\} \cup \{\rho\} \in \mathcal{S}_{\mathcal{P}}$.
Proof of claim: First, let us show that $\bar{\rho}$ is a maximal element of S so that $S - \{\bar{\rho}\}$ remains closed. Suppose this is not the case so there exists $\sigma \in S$ such that $\bar{\rho} \leq \sigma$. This also means that $\bar{\sigma} \leq \rho$. Now, either $\sigma \in S'$ or it is not. If it is, then S' is not closed since $\bar{\rho} \notin S'$ – a contradiction. If $\sigma \notin S'$ then $\bar{\sigma} \in S'$. This means that both $\bar{\sigma}$ and ρ are in $S' - S$ so ρ is no longer a minimal element of $S' - S$ which is, again, a contradiction. Hence, σ does not exist.

Second, let us show that all the predecessors of ρ also belong to $S - \{\bar{\rho}\}$ so that $S - \{\bar{\rho}\} \cup \{\rho\}$ is still closed. Let $\tau \leq \rho$ so $\bar{\rho} \leq \bar{\tau}$. (Note that $\tau \neq \bar{\rho}$ by the first property of mirror posets.) Since we just showed that $\bar{\rho}$ is a maximal element of S, $\bar{\tau} \notin S$. But S is a complete set of P so $\tau \in S$. Thus, $S - \{\bar{\rho}\} \cup \{\rho\}$ is a closed subset of P. Finally, $S - \{\bar{\rho}\} \cup \{\rho\}$ is complete since every dual pair still has exactly one element in it.

Consider $S, S' \in \mathcal{S}_{\mathcal{P}}$. Topologically order the elements in $S' - S$ and let the result be $\rho_1, \rho_2, \ldots, \rho_k$. Let $T_0 = S$ and $T_i = T_{i-1} \cup \{\rho_i\} - \{\bar{\rho_i}\}$ for $i = 1, \ldots, k$ so $T_k = S'$. Since $T_0 \in \mathcal{S}_{\mathcal{P}}$ and ρ_1 a minimal element of $S' - T_0$, the above claim states that $T_1 \in \mathcal{S}_{\mathcal{P}}$. Furthermore, ρ_2 is a minimal element of $S' - T_1$; again, by the above claim $T_2 \in \mathcal{S}_{\mathcal{P}}$. Applying this argument repeatedly, we have that each $T_i \in \mathcal{S}_{\mathcal{P}}$. In addition, it is easy to check that T_i and T_{i+1} are adjacent in $G(\mathcal{S}_{\mathcal{P}})$ for $i = 0, \ldots, k-1$. Therefore, $T_0, T_1, \ldots, T_k$ is a path from S to S' in $G(\mathcal{S}_{\mathcal{P}})$. This path shows that $d(S, S') \leq |S' - S|$. But in fact any path from S to S' must be have length at least $|S' - S|$ since the elements in $S' - S$ must be introduced into the sets encountered along the path one at a time. Hence, $d(S, S') = |S' - S|$. $\Box$

Lemma 7 *Let $\pi = (S_1, S_2, \ldots, S_l) \in \mathcal{S}^l_{\mathcal{P}}$. For each $\rho \in P$, define $n_{\rho,\pi}$ to be the number of times ρ appears in the sets $S_1, S_2, \ldots, S_l$. Let $S_{maj}(\pi) = \{\rho : n_{\rho,\pi} > l/2\}$; i.e., $S_{maj}(\pi)$ consists of all rotations ρ that appear in majority of the closed subsets in π. Then, $S_{maj}(\pi)$ is closed and partially complete. Additionally, when l is odd, $S_{maj}(\pi)$ is complete.*

Proof When $\rho \in S_i$, each $\sigma \leq \rho$ also belongs to S_i since S_i is closed; hence, $n_{\sigma,\pi} \geq n_{\rho,\pi}$. Thus, $S_{maj}(\pi)$ is closed because when $n_{\rho,\pi} > l/2$, it is also the case that $n_{\sigma,\pi} > l/2$.

For any dual pair $\{\rho, \bar{\rho}\}$, $n_{\rho,\pi} + n_{\bar{\rho},\pi} = l$. Since at most one of these values can be greater than $l/2$, at most one of ρ and $\bar{\rho}$ can belong to $S_{maj}(\pi)$. And when l is odd, $S_{maj}(\pi)$ is complete because exactly one of $n_{\rho,\pi}$ and $n_{\bar{\rho},\pi}$ is greater than $l/2$. $\Box$

Theorem 6 *Let* $\pi = (S_1, S_2, \ldots, S_l) \in \mathcal{S}_{\mathcal{P}}^l$.

(i) When l *is odd,* $S_{maj}(\pi)$ *is the unique median of* π *in* $G(\mathcal{S}_{\mathcal{P}})$.

(ii) When l *is even, either* $S_{maj}(\pi)$ *is the unique median of* π *in* $G(\mathcal{S}_{\mathcal{P}})$ *or the extensions of* $S_{maj}(\pi)$ *are exactly the medians of* π *in* $G(\mathcal{S}_{\mathcal{P}})$.

Proof Since $d(S, S_i) = |S_i - S|$, $\rho \in S$ contributes a value of 1 to $d(S, S_i)$ if $\rho \notin S_i$ and 0 otherwise. Thus, $d(S, S_i)$ is the sum of the contributions of each $\rho \in S$ to $d(S, S_i)$. Therefore,

$$D(S, \pi) = \sum_{i=1}^{l} d(S, S_i) = \sum_{\rho \in S} n_{\bar{\rho},\pi}$$

because the number of times that $\rho \in S$ does not appear in $S_1, \ldots, S_l$ is equal to the number of times that $\bar{\rho}$ appears in those sets. This also implies that

$$\begin{aligned} D(S, \pi) - D(S', \pi) &= \sum_{\rho \in S} n_{\bar{\rho},\pi} - \sum_{\rho \in S'} n_{\bar{\rho},\pi} \\ &= \sum_{\rho \in S - S'} n_{\bar{\rho},\pi} - \sum_{\rho \in S' - S} n_{\bar{\rho},\pi} \\ &= \sum_{\rho \in S - S'} (n_{\bar{\rho},\pi} - n_{\rho,\pi}) \end{aligned}$$

since S and S' are complete subsets so $\rho \in S - S'$ if and only if $\bar{\rho} \in S' - S$.

When l is odd, Lemma 7 states that $S_{maj}(\pi) \in \mathcal{S}_{\mathcal{P}}$. Furthermore, for each $\rho \in S_{maj}(\pi)$ $n_{\bar{\rho},\pi} < n_{\rho,\pi}$ so $n_{\bar{\rho},\pi} - n_{\rho,\pi} < 0$. It follows that for each $S' \neq S_{maj}(\pi)$, $D(S_{maj}(\pi), \pi) - D(S', \pi) < 0$ so $S_{maj}(\pi)$ is the unique median of π in $G(\mathcal{S}_{\mathcal{P}})$.

So suppose l is even. If there is no dual pair $\{\rho, \bar{\rho}\}$ with $n_{\rho,\pi} = n_{\bar{\rho},\pi} = l/2$, $S_{maj}(\pi)$ is complete and closed. The above discussion again implies that $S_{maj}(\pi)$ is the unique median of π. But when this is not the case, $S_{maj}(\pi)$ is partially complete and closed. By Lemma 5, it can be extended into a complete closed subset of $\mathcal{P}$. Let $T_1, T_2 \in \mathcal{S}_{\mathcal{P}}$ so that $S_{maj}(\pi) \subset T_1, T_2$; that is, T_1 and T_2 are extensions of $S_{maj}(\pi)$. We note that when $\rho \in T_1 - T_2$, $n_{\bar{\rho},\pi} = n_{\rho,\pi}$ so $D(T_1, \pi) - D(T_2, \pi) = 0$. In other words, all extensions of $S_{maj}(\pi)$ have the same distances from π. Now suppose S' is not an extension of $S_{maj}(\pi)$ but T is. Then for every $\rho \in T - S'$, either $n_{\bar{\rho},\pi} = n_{\rho,\pi}$ or $n_{\bar{\rho},\pi} < n_{\rho,\pi}$; moreover, there must be at least one element $\rho \in T - S'$ with $n_{\bar{\rho},\pi} < n_{\rho,\pi}$. Thus, it must be the case that $D(T, \pi) - D(S', \pi) < 0$. This shows that the extensions of $S_{maj}(\pi)$ are exactly the medians of π in $G(\mathcal{S}_{\mathcal{P}})$. $\square$

An immediate consequence of the above theorem is that every family of three vertices of $G(\mathcal{S}_{\mathcal{P}})$ has a unique median.

Corollary 1 *The graph* $G(\mathcal{S}_{\mathcal{P}})$ *is a median graph.*

4 The graph of stable roommates matchings and its medians

When I is an SR instance, $\mathcal{R}'(I)$ is a mirror poset whose complete closed subsets are in one-to-one correspondence with the stable matchings of I. Our discussion in the previous section suggests a natural graph for the stable matchings of I. For each $\mu \in M(I)$, let S_μ denote the complete closed subset of $\mathcal{R}'(I)$ that corresponds to μ.

Definition 3 *Let $G(M(I))$ denote the graph whose vertices are the stable matchings of I, and two stable matchings μ and μ' are adjacent if and only if S_μ and $S_{\mu'}$ differ by one rotation.*

Thus, $G(M(I))$ is isomorphic to $G(\mathcal{S}_{\mathcal{R}'(I)})$. According to Corollary 1, it is a median graph. Furthermore, Theorem 6 describes the complete closed subsets of $\mathcal{R}'(I)$ that correspond to the medians of a sequence of stable matchings in $G(M(I))$. After the next lemma, we present the connections between these medians and the stable matchings discovered by Teo and Sethuraman.

Lemma 8 *Let I be an SR instance. Suppose that participant x appears in the X-sets of the non-singular rotations $\rho_1, \ldots, \rho_r$. In particular, $(x, y_i) \in \rho_i$ for $i = 1, \ldots, r$, and x prefers y_i over y_{i+1} for $i = 1, \ldots, r-1$. Then*

(i) $\rho_1 \leq \rho_2 \leq \ldots \leq \rho_r$,

(ii) $(y_{i+1}, x) \in \bar{\rho}_i$ *for* $i = 1, \ldots, r-1$ *and*

(iii) the stable partners of x are exactly $y_1, y_2, \ldots, y_r, y_{r+1}$ ordered from her most preferred to her least preferred, where $(y_{r+1}, x) \in \bar{\rho}_r$.

Proof Let $1 \leq i \leq r-1$. Since x prefers y_i over y_{i+1}, y_i appears before y_{i+1} in x's list. In order for ρ_{i+1} to be exposed in a table, y_{i+1} has to be the first person in x's list. Thus, y_i has to be removed from x's list. But ρ_i is the only non-singular rotation that contains (x, y_i) so it has to be eliminated first. It follows that $\rho_1, \rho_2, \ldots, \rho_r$ forms a chain in $\mathcal{R}'(I)$.

Suppose $(z, x) \in \bar{\rho}_i$. This means that once ρ_i has been eliminated, the front of x's list is z. If $z \neq y_{i+1}$, then z has to be removed in order for ρ_{i+1} to be exposed. Thus, (x, z) must be part of some rotation that lies between ρ_i and ρ_{i+1}. But no such rotation exists according to part (i), so $z = y_{i+1}$.

According to Lemma 4, every stable partner of x must appear with x in some non-singular rotation of I. Since x is in the X-set of σ if and only if x is in the Y-set of $\bar{\sigma}$, we simply need to consider the participants that appear with x in $\rho_1, \ldots, \rho_r$ and $\bar{\rho}_1, \ldots, \bar{\rho}_r$. Applying part (ii), the stable partners of x consists exactly of $y_1, y_2, \ldots, y_r$ and y_{r+1} where $(y_{r+1}, x) \in \bar{\rho}_r$. Clearly, x prefers y_r over y_{r+1}. $\square$

Theorem 7 *Let I be an SR instance and $\pi = (\mu_1, \mu_2, \ldots, \mu_l) \in \mathcal{M}(I)^l$. For each participant x and for $i = 1, \ldots, l$, define $p_{i,\pi}(x)$ as before.*

a. When l is odd, the stable matching that assigns each participant x to $p_{(l+1)/2,\pi}(x)$ is the unique median of π in $G(M(I))$.

b. When l is even, the stable matchings that assign each participant x to $p_{l/2,\pi}(x)$ or $p_{l/2+1,\pi}(x)$ are medians of π in $G(M(I))$. Additionally, when every stable matching of I is part of π, all medians of π assign each participant x to $p_{l/2,\pi}(x)$ or $p_{l/2+1,\pi}(x)$.

Proof In this proof, let $(S_{\mu_1}, S_{\mu_2}, \ldots, S_{\mu_l})$ be an alternate representation for $\pi = (\mu_1, \mu_2, \ldots, \mu_l)$. We then define $n_{\rho,\pi}$ and $S_{maj}(\pi)$ as before; i.e., $n_{\rho,\pi}$ is the number of times rotation ρ appears in $S_{\mu_1}, S_{\mu_2}, \ldots, S_{\mu_l}$, and $S_{maj}(\pi)$ consists of those rotations that appear in majority of $S_{\mu_1}, S_{\mu_2}, \ldots, S_{\mu_l}$.

For each participant x, let $R'(x)$ denote the set containing all the non-singular rotations that have x as part of their X-sets. Let μ^* be a stable matching that matches each participant to her (lower or upper) median stable partner in π. To prove part (i), we will show that $S_{\mu^*} \cap R'(x) = S_{maj}(\pi) \cap R'(x)$ for every participant x. Hence, $S_{\mu^*} = S_{maj}(\pi)$. According to Theorem 6, μ^* is the unique median of π. To prove the first half of part (ii), we will show that $S_{maj}(\pi) \cap R'(x) \subseteq S_{\mu^*} \cap R'(x)$ for every participant x. This implies that $S_{maj}(\pi) \subseteq S_{\mu^*}$; i.e., either $S_{\mu^*} = S_{maj}(\pi)$ or S_{μ^*} is an extension of $S_{maj}(\pi)$. Again, according to Theorem 6, μ^* is a median of π.

Let x be some participant. If $R'(x) = \emptyset$, then $S_{maj}(\pi) \cap R'(x) = S_\mu \cap R'(x) = \emptyset$. So suppose $R'(x) \neq \emptyset$. From Lemma 8, the rotations in $R'(x)$ can be arranged as $\rho_1, \ldots, \rho_r$ which forms a chain in $\mathcal{R}'(I)$. Furthermore, the stable partners of x are $y_1, y_2, \ldots, y_r, y_{r+1}$ where $(x, y_i) \in \rho_i$ for $i = 1, \ldots, r$, $(y_{r+1}, x) \in \bar{\rho}_r$, and the partners are arranged from x's most preferred to least preferred stable partner. It is also straightforward to verify that if x is matched to y_{i+1} in some stable matching μ, then $S_\mu \cap R'(x) = \{\rho_1, \ldots, \rho_i\}$.

Assume l is odd. Suppose x is matched to y_{i^*+1} in μ^*. Then $S_{\mu^*} \cap R'(x) = \{\rho_1, \ldots, \rho_{i^*}\}$. By our definition of μ^*, we also have $p_{(l+1)/2,\pi}(x) = y_{i^*+1}$. This means that at least $(l+1)/2$ of the stable matchings in π match x to $y_{i^*+1}, y_{i^*+2}, \ldots$, or y_{r+1}. In all of these stable matchings, ρ_{i^*} must be in their corresponding closed subsets. Hence, $n_{\rho_{i^*},\pi} \geq (l+1)/2$. Applying a similar reasoning, we also conclude that $n_{\rho_{i^*+1},\pi} \leq (l-1)/2$. And since $n_{\rho,\pi} \leq n_{\sigma,\pi}$ whenever $\sigma < \rho$ in $\mathcal{R}(I)$, it follows that $S_{maj}(\pi) \cap R'(x) = \{\rho_1, \ldots, \rho_{i^*}\}$. In other words, $S_{maj} \cap R'(x) = S_{\mu^*}(\pi) \cap R'(x)$ for an arbitrary participant x.

Suppose l is even. Again, suppose x is matched to y_{i^*+1} in μ^* so $S_{\mu^*} \cap R'(x) = \{\rho_1, \ldots, \rho_{i^*}\}$. Using the same argument in the previous paragraph, when $p_{l/2,\pi}(x) = y_{j+1}$, $S_{maj}(\pi) \cap R'(x) = \{\rho_1, \ldots, \rho_j\}$. But by our assumption, $p_{l/2,\pi}(x)$ or $p_{l/2+1,\pi}(x)$ is y_{i^*+1}. Hence, $j \leq i^*$. Therefore, $S_{maj}(\pi) \cap R'(x) \subseteq S_{\mu^*} \cap R'(x)$.

Finally, let us prove the second half of part (ii). Let every stable matching of I occur in π, and let μ_{med} be a median of π. According to Theorem 6, $S_{maj}(\pi) \subseteq S_{\mu_{med}}$. Suppose x is matched to y_{i+1} in μ_{med}. This means that $\rho_1, \ldots, \rho_i \in S_{\mu_{med}}$ but $\rho_{i+1}, \ldots, \rho_{r+1} \notin S_{\mu_{med}}$. In particular, $\rho_{i+1}, \ldots, \rho_{r+1} \notin S_{maj}$. If $p_{l/2,\pi}(x) = y_{j+1}$ for some $j > i$, then as in the previous paragraphs $n_{\rho_j,\pi} \geq l/2 + 1$; i.e., $\rho_j \in S_{maj}$. But $j \geq i+1$ so this is a contradiction. It follows that $p_{l/2,\pi}(x) = y_{j+1}$ where $j \leq i$. If $p_{l/2,\pi}(x) = y_{i+1}$ or $p_{l/2,\pi}(x) = y_{j+1}$, $j < i$ but $p_{l/2+1,\pi}(x) = y_{i+1}$, we are done since x is matched to either her lower or upper median stable partner. The only case we have to consider is $p_{l/2,\pi}(x) = y_{j+1}$, $j < i$ but $p_{l/2+1,\pi}(x) = y_k$, $k > i+1$. This implies that x is never matched to y_{i+1} in the stable matchings in π, which is a contradiction since every stable matching of I occurs in π. Since x is an arbitrary participant of I, every participant in μ_{med} must be matched to her lower or upper median stable partner. $\square$

Corollary 2 *Let I be an SR instance. A stable matching μ matches each participant x to her (lower or upper) median stable partner in I if and only if μ is a median of the graph $G(M(I))$.*

In other words, the local/global median phenomenon that was observed in SM stable matchings generalizes to SR stable matchings. It does so because the graph underlying the set of stable roommates matchings is a median graph.

Example continued. In the example in Section 2.3, $S_{\mu_1} = \{\bar{\rho}_5, \bar{\rho}_4, \bar{\rho}_3, \bar{\rho}_2, \bar{\rho}_1\}$, $S_{\mu_2} = \{\bar{\rho}_5, \bar{\rho}_4, \bar{\rho}_3, \bar{\rho}_2, \rho_1\}$, $S_{\mu_3} = \{\bar{\rho}_5, \bar{\rho}_4, \bar{\rho}_3, \rho_2, \bar{\rho}_1\}$, $S_{\mu_4} = \{\bar{\rho}_5, \bar{\rho}_4, \rho_3, \bar{\rho}_2, \bar{\rho}_1\}$, $S_{\mu_5} = \{\bar{\rho}_5, \rho_4, \bar{\rho}_3, \bar{\rho}_2, \rho_1\}$, $S_{\mu_6} = \{\rho_5, \bar{\rho}_4, \bar{\rho}_3, \bar{\rho}_2, \rho_1\}$. Thus, when $\pi = (\mu_1, \mu_2, \mu_3, \mu_4, \mu_5, \mu_6)$, $S_{maj}(\pi) = \{\bar{\rho}_5, \bar{\rho}_4, \bar{\rho}_3, \bar{\rho}_2\}$ – which is a closed but not complete subset of $\mathcal{R}'(I)$. Its two extensions are S_{μ_1} and S_{μ_2}. Thus, μ_1 and μ_2 are the medians of $G(M(I))$. This can also be verified by drawing $G(M(I))$, shown in Figure 2 where u_i corresponds to μ_i, $i = 1, \ldots, 6$. But μ_1 and μ_2 are also the only two stable matchings in π that match participants to their median stable partners as noted in Section 2.3. Hence, μ_1 and μ_2 are median stable matchings in a local *and* global sense.

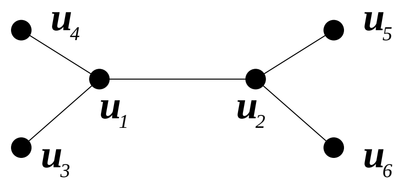

Figure 2: The graph of stable matchings for the instance given in Section 2.3. Node u_i corresponds to the stable matching μ_i for $i = 1, \ldots, 6$.

Below, we note a few more properties of the median set of $\pi = (\mu_1, \ldots, \mu_l) \in M(I)^l$ based on the work of Bandelt and Barthélemy [4] on the medians of median graphs.

- In a graph $G = (V, E)$, the interval $I(u, v)$ between any two vertices u and v is the set $\{t : d(u, t) + d(t, v) = d(u, v)\}$. The median set of π in $G(M(I))$ is always some interval $I(\alpha(\pi), \beta(\pi))$ in G. In particular, the set induces a connected subgraph of $G(M(I))$. The values of $\alpha(\pi)$ and $\beta(\pi)$ can be determined from Proposition 6 of [4].

- For each $\mu \in M(I)$, let $N(\mu)$ consist of all stable matchings adjacent to μ in $G(M(I))$. A stable matching μ^* is said to be a *local median* of π if $D(\mu^*, \pi) \leq D(\mu, \pi)$ for each $\mu \in N(\mu^*)$. It turns out that because $G(M(I))$ is a median graph, μ^* is a local median of π if and only if it is a median of π.

- A stable matching μ^* is a *Condorcet vertex* of π if the number of stable matchings in π closer to μ^* is greater than or equal to the corresponding number for every other stable matching μ. The *Condorcet set* of π contains all the Condorcet vertices of π. The median set of π is exactly the Condorcet set of π whenever $G(M(I))$ is a cube-free median graph (i.e., a graph that does not contain the cube, which is formed by connecting the corresponding vertices of two 4-cycles, as a subgraph).

5 Mirror posets and stable roommate matchings

When I is an SR instance, $\mathcal{R}'(I)$ is a mirror poset. In this section, we prove the converse – that every mirror poset gives rise to an SR instance. Our construction is a generalization of the one used by Irving and Leather [16] to create a small SM instance from an arbitrary poset. The proof, however, is more involved because only half of the rotations are revealed from the construction; careful analysis is needed to argue the existence of the other half of the rotations.

For any two vertices α and β of a directed graph, let $\langle \alpha, \beta \rangle$ denote the directed arc from α to β and $(\alpha, \beta) = (\beta, \alpha)$ the undirected edge between the two elements. We say that β is a *neighbor* of α when there is an edge $\langle \alpha, \beta \rangle$ or $\langle \beta, \alpha \rangle$ in the directed graph.

Let $\mathcal{P} = (P, \leq)$ be a mirror poset with $2n$ elements, and let S be one of its complete closed subsets. Without loss of generality, assume that when the elements of S are topologically ordered, the result is $\bar{\sigma}_n, \bar{\sigma}_{n-1}, \ldots, \bar{\sigma}_1$. Thus, the other elements of P are $\sigma_1, \ldots, \sigma_n$, and $\bar{\sigma}_n, \ldots, \bar{\sigma}_1, \sigma_1, \ldots, \sigma_n$ is a topological ordering of the elements of P. Let $H(\mathcal{P})$ denote the Hasse diagram of $\mathcal{P}$.

`procedure construct-instance` $(\mathcal{P}, (\bar{\sigma}_n, \bar{\sigma}_{n-1}, \ldots, \bar{\sigma}_1))$

begin
  form $\mathcal{P}'$ from $\mathcal{P}$ by adjoining an element $\bar{\sigma}_{n+1}$ that precedes all others and an element σ_{n+1}
    that is a successor of all others;
  for each edge $\langle \alpha, \beta \rangle$ in the Hasse diagram $H(\mathcal{P}')$ of $\mathcal{P}'$
    create a participant $p(\alpha, \beta) = p(\beta, \alpha)$;
  for each pair of edges $\langle \alpha, \beta \rangle$ and $\langle \bar{\beta}, \bar{\alpha} \rangle$
    place $p(\alpha, \beta)$ and $p(\bar{\beta}, \bar{\alpha})$ on each other's (initially empty) preference lists;
  {This completes iteration 0.}
  for $i := 1$ to n
  begin
    $N(i) := (\alpha_0^i, \ldots, \alpha_{r-1}^i)$, an arbitrary ordering of the neighbors of σ_i;
    $E(i) := (x_0^i, \ldots, x_{r-1}^i)$ where $x_j^i := p(\sigma_i, \alpha_j^i)$;
    $L(i) := (y_0^i, \ldots, y_{r-1}^i)$, the ordered set of people such that y_j^i is (currently) last on x_j^i's list;
    for $j := 0$ to $r - 1$ do
      place y_{j+1}^i at the end of x_j^i's list;
      place x_j^i at the beginning of y_{j+1}^i's list;
      {addition in the subscript is taken mod r}
    {This completes iteration i.}
  end
  for each participant
    if his list is incomplete
      add missing participants at the end of his list in an arbitrary order;
  {This completes iteration $n + 1$.}
  return the created SR instance as $I(\mathcal{P})$;
end

Example continued. Suppose we used the mirror poset shown in Figure 1 as the input to the algorithm with σ_i corresponding to ρ_i and $\bar{\sigma}_i$ corresponding to $\bar{\rho}_i$ for $i = 1, \ldots, 5$. The result after iteration n is a table identical to the phase-1 table T_0 with the following correspondence:

$p(\bar{\sigma}_6, \bar{\sigma}_5)$:	p_1	$p(\bar{\sigma}_5, \sigma_4)$:	p_6	$p(\bar{\sigma}_2, \sigma_3)$:	p_{11}	$p(\sigma_1, \sigma_5)$:	p_{16}
$p(\bar{\sigma}_6, \bar{\sigma}_4)$:	p_2	$p(\bar{\sigma}_4, \bar{\sigma}_1)$:	p_7	$p(\bar{\sigma}_2, \sigma_1)$:	p_{12}	$p(\sigma_2, \sigma_6)$:	p_{17}
$p(\bar{\sigma}_6, \bar{\sigma}_3)$:	p_3	$p(\bar{\sigma}_4, \sigma_5)$:	p_8	$p(\bar{\sigma}_1, \sigma_2)$:	p_{13}	$p(\sigma_3, \sigma_6)$:	p_{18}
$p(\bar{\sigma}_6, \bar{\sigma}_2)$:	p_4	$p(\bar{\sigma}_3, \sigma_2)$:	p_9	$p(\bar{\sigma}_1, \sigma_3)$:	p_{14}	$p(\sigma_4, \sigma_6)$:	p_{19}
$p(\bar{\sigma}_5, \bar{\sigma}_1)$:	p_5	$p(\bar{\sigma}_3, \sigma_1)$:	p_{10}	$p(\sigma_1, \sigma_4)$:	p_{15}	$p(\sigma_5, \sigma_6)$:	p_{20}

Lemma 9 *In* `construct-instance`*, let A_i be the preference table at the end of iteration i, for $i = 0, \ldots, n+1$. Let $\rho_i = (x_0^i, y_0^i), (x_1^i, y_1^i), \ldots, (x_{r-1}^i, y_{r-1}^i)$ so that $\bar{\rho}_i = (y_0^i, x_{r-1}^i), (y_1^i, x_0^i), \ldots, (y_{r-1}^i, x_{r-2}^i)$ for $i = 1, \ldots, n$. Suppose we use Irving's algorithm to find a stable matching of $I(\mathcal{P})$. Then*

a. A_n is the phase-1 table,

b. $\bar{\rho}_i$ is an exposed rotation of A_i and $A_{i-1} = A_i / \bar{\rho}_i$, for $i = 1, \ldots, n$,

c. $\{\bar{\rho}_1, \bar{\rho}_2, \ldots, \bar{\rho}_n\}$ contains all the singular rotations and exactly one of each dual pair of non-singular rotations of $I(\mathcal{P})$.

Proof Notice that `construct-instance` always maintained the property that $y = l_{A_i}(x)$ if and only if $x = f_{A_i}(y)$, for $i = 0, \ldots, n$. Since A_{n+1} was created from A_n by simply completing the preference lists of the participants, when y is considered during phase 1 of Irving's algorithm, all participants that were appended to the end of the list of $x = f_{A_n}(y)$ will be removed. Thus, A_n is the phase-1 table. Furthermore, it is easy to verify that by the way A_i was created from A_{i-1}, $\bar{\rho}_i$ is an exposed rotation in A_i and $A_{i-1} = A_i/\bar{\rho}_i$. Thus $A_0 = A_n/\{\bar{\rho}_n, \bar{\rho}_{n-1}, \ldots, \bar{\rho}_1\}$. Since A_0 has no empty lists, it corresponds to a stable matching of $I(\mathcal{P})$. According to Theorem 4, every singular rotation of $I(\mathcal{P})$, and exactly one of each dual pair of non-singular rotations of $I(\mathcal{P})$ had to be eliminated from A_n to obtain A_0. The third part of the theorem follows. $\square$

Even though we have not yet established that the ordered lists ρ_i, $i = 1, \ldots, n$, are rotations, we will still say that the X-set and Y-set of each ρ_i are $\{x^i_0, x^i_1, \ldots, x^i_{r-1}\}$ and $\{y^i_0, y^i_1, \ldots, y^i_{r-1}\}$ respectively. For $\phi = \rho_i$ or $\bar{\rho}_i$, $i = 1, \ldots, n$, let $X(\phi)$ and $Y(\phi)$ denote the X-set and Y-set of ϕ respectively. For each participant z of $I(\mathcal{P})$, the next lemmas describe when $z \in X(\phi)$ and $z \in Y(\phi)$, enabling us to understand later the relationship between the σ_i's and the ρ_i's, and the $\bar{\sigma}_i$'s and the $\bar{\rho}_i$'s. The lemmas make use of the fact that in $H(\mathcal{P}')$ if there are edges from $\bar{\sigma}_j$ to $\bar{\sigma}_i$ and from σ_i to σ_j then $j > i$, and if there are edges from $\bar{\sigma}_i$ to σ_j and from $\bar{\sigma}_j$ to σ_i then $i \neq j$.

Lemma 10 *(The X-sets.) Let $1 \leq i < j \leq n+1$.*

a. Suppose $\langle \sigma_i, \sigma_j \rangle$ is an edge in $H(\mathcal{P}')$. Then $p(\sigma_i, \sigma_j) \in X(\rho_k)$ if and only if $k = i$ or $k = j \leq n$.

b. Suppose $\langle \bar{\sigma}_j, \bar{\sigma}_i \rangle$ is an edge in $H(\mathcal{P}')$. Then $p(\bar{\sigma}_j, \bar{\sigma}_i) \notin X(\rho_k)$ for any k.

c. Suppose $\langle \bar{\sigma}_j, \sigma_i \rangle$ is an edge in $H(\mathcal{P}')$. Then $p(\bar{\sigma}_j, \sigma_i) \in X(\rho_k)$ if and only if $k = i$.

d. Suppose $\langle \bar{\sigma}_i, \sigma_j \rangle$ is an edge in $H(\mathcal{P}')$. Then $p(\bar{\sigma}_i, \sigma_j) \in X(\rho_k)$ if and only if $k = j \leq n$.

Proof The lemma follows directly from the observation that the X-set of ρ_k consists of participants $p(\sigma_k, \alpha)$ where α is a neighbor of σ_k, and $1 \leq k \leq n$. $\square$

Lemma 11 *(The Y-sets.) Let $1 \leq i < j \leq n+1$.*

a. Suppose $\langle \sigma_i, \sigma_j \rangle$ is an edge in $H(\mathcal{P}')$. Then $p(\sigma_i, \sigma_j) \notin Y(\rho_k)$ for any k.

b. Suppose $\langle \bar{\sigma}_j, \bar{\sigma}_i \rangle$ is an edge in $H(\mathcal{P}')$. Then $p(\bar{\sigma}_j, \bar{\sigma}_i)$ is in the Y-set of one or more of the ρ_k's. Moreover, these ordered lists can be arranged as $\rho_{k_0}, \rho_{k_1}, \ldots, \rho_{k_{\bar{i}\bar{j}}}$ so that $\rho_{k_0} = \rho_i$, and $\sigma_{k_0}, \sigma_{k_1}, \ldots, \sigma_{k_{\bar{i}\bar{j}}}$ is a directed path in $H(\mathcal{P})$.

c. Suppose $\langle \bar{\sigma}_j, \sigma_i \rangle$ is an edge in $H(\mathcal{P}')$. If $j = n+1$, $p(\bar{\sigma}_j, \sigma_i) \notin Y(\rho_k)$ for any k. If $j \leq n$, $p(\bar{\sigma}_j, \sigma_i)$ is in the Y-set of one or more of the ρ_k's. Moreover, these ordered lists can be arranged as $\rho_{k_0}, \rho_{k_1}, \ldots, \rho_{k_{i\bar{j}}}$ so that $\rho_{k_0} = \rho_j$, and $\sigma_{k_0}, \sigma_{k_1}, \ldots, \sigma_{k_{i\bar{j}}}$ is a directed path in $H(\mathcal{P})$.

d. Suppose $\langle \bar{\sigma}_i, \sigma_j \rangle$ is an edge in $H(\mathcal{P}')$. Then $p(\bar{\sigma}_i, \sigma_j)$ is in the Y-set of one or more of the ρ_k's. Moreover, these ordered lists can be arranged as $\rho_{k_0}, \rho_{k_1}, \ldots, \rho_{k_{\bar{i}j}}$ so that $\rho_{k_0} = \rho_i$, and $\sigma_{k_0}, \sigma_{k_1}, \ldots, \sigma_{k_{\bar{i}j}}$ is a directed path in $H(\mathcal{P})$.

Proof First, we note that $p(\sigma_i, \sigma_j)$ is never last in the list of any $p(\sigma_k, \psi)$, ψ a neighbor of σ_k, in tables $A_0, \ldots, A_n$. Hence, $p(\sigma_i, \sigma_j) \notin Y(\rho_k)$ for any k. By the same reasoning, when $j = n+1$ $p(\bar{\sigma}_j, \sigma_i) \notin Y(\rho_k)$ for any k. On the other hand, the smallest index k for which $p(\bar{\sigma}_j, \bar{\sigma}_i) \in Y(\rho_k)$ is when $k = i$. Similarly, the smallest index k for which $p(\bar{\sigma}_j, \sigma_i) \in Y(\rho_k)$, $j \leq n$, is when $k = j$, and for which $p(\bar{\sigma}_i, \sigma_j) \in Y(\rho_k)$, is when $k = i$.

A participant z belongs to $Y(\rho_k)$ if and only if z is the last person in some $p(\sigma_k, \psi)$'s list, ψ a neighbor of σ_k, at the beginning of the kth iteration. After the kth iteration, z is again the last person in some $p(\sigma_k, \alpha)$'s list, α a neighbor of σ_k. If $\alpha = \bar{\sigma}_t$ for some t, $\alpha = \sigma_m$ for some $m < k$, or $\alpha = \sigma_{n+1}$, then z will stay as the last person on $p(\sigma_k, \alpha)$'s list until the end of the nth iteration because $p(\sigma_k, \alpha)$ will not be part of $E(r)$, $r > k$. Hence, $z \notin Y(\rho_r)$, $r > k$. On the other hand, if $\alpha = \sigma_m$, $k < m \leq n$, z will stay at the end of $p(\sigma_k, \alpha)$'s list from the kth to the $(m-1)$st iteration. At the mth iteration, $p(\sigma_k, \sigma_m)$ will be part of $E(m)$ since σ_k is a neighbor of σ_m, and $z \in Y(\rho_m)$. And since $k < m$, it must be the case that $\langle \sigma_k, \sigma_m \rangle$ is an edge of $H(\mathcal{P})$.

We have shown that if $z \in Y(\rho_k)$ and m is the next largest index after k such that $z \in Y(\rho_m)$, $\langle \sigma_k, \sigma_m \rangle$ is an edge of $H(\mathcal{P})$. Parts (b), (c) and (d) follow. □

Lemma 12 *Let $1 \leq i < j \leq n$. If $\langle \bar{\sigma}_j, \bar{\sigma}_i \rangle$ is an edge of $H(\mathcal{P})$, then $\bar{\rho}_j$ precedes $\bar{\rho}_i$. Similarly, if $\langle \bar{\sigma}_j, \sigma_i \rangle$ is an edge of $H(\mathcal{P})$ and ρ_i is a rotation, then $\bar{\rho}_j$ also precedes ρ_i.*

Proof If $\langle \bar{\sigma}_j, \bar{\sigma}_i \rangle$ is an edge of $H(\mathcal{P})$, $\langle \sigma_i, \sigma_j \rangle$ is also an edge of $H(\mathcal{P})$. According to Lemma 10, $p(\sigma_i, \sigma_j) \in X(\rho_i) \cap X(\rho_j) = Y(\bar{\rho}_i) \cap Y(\bar{\rho}_j)$. Thus, during the ith iteration, there is some participant y that was placed at the end of $p(\sigma_i, \sigma_j)$'s list so that $(y, p(\sigma_i, \sigma_j))$ belongs to $\bar{\rho}_i$. Later, during the jth iteration, there is another participant y' that was placed at the end of $p(\sigma_i, \sigma_j)$'s list so that $(y', p(\sigma_i, \sigma_j))$ belongs to $\bar{\rho}_j$. Now, in order for $\bar{\rho}_i$ to be exposed, y has to be at the end of $p(\sigma_i, \sigma_j)$'s list. In particular, y' has to be removed. Since the pair $(p(\sigma_i, \sigma_j), y')$ is part of only one rotation of $I(\mathcal{P})$, $\bar{\rho}_j$ has to be eliminated first.

Suppose $\langle \bar{\sigma}_j, \sigma_i \rangle$ is an edge of $H(\mathcal{P})$. According to Lemmas 10 and 11, $p(\bar{\sigma}_j, \sigma_i) \in X(\rho_i) \cap Y(\rho_j) = X(\rho_i) \cap X(\bar{\rho}_j)$. Let $(p(\bar{\sigma}_j, \sigma_i), y)$ be in ρ_i; i.e., y is at the end of $p(\bar{\sigma}_j, \sigma_i)$'s list at the beginning of iteration i. During the jth iteration, there is some participant x that was placed in front $p(\bar{\sigma}_j, \sigma_i)$'s list so that $(p(\bar{\sigma}_j, \sigma_i), x)$ belongs to $\bar{\rho}_j$. If ρ_i is a rotation, all the participants in front of y in $p(\bar{\sigma}_j, \sigma_i)$'s list have to be removed in order for it to be exposed. Since x is in front of y and $(p(\bar{\sigma}_j, \sigma_i), x)$ is part of only one rotation, $\bar{\rho}_j$ has to be eliminated first. □

We are now ready to prove that each ρ_i, $i = 1, \ldots, n$ is a rotation of $I(\mathcal{P})$. In other words, $I(\mathcal{P})$ has no singular rotations so $R(I(\mathcal{P})) = R'(I(\mathcal{P})) = \{\bar{\rho}_1, \ldots, \bar{\rho}_n, \rho_1, \ldots, \rho_n\}$. Furthermore, we will show that there is an isomorphism from $\mathcal{R}'(I(\mathcal{P}))$ to $\mathcal{P}$. We will prove these two important results simultaneously using induction. We were unable to split them apart because they depend on each other.

Theorem 8 *Define $R_k = \{\rho_i, \bar{\rho}_i, 1 \leq i \leq k\}$ and $P_k = \{\sigma_i, \bar{\sigma}_i, 1 \leq i \leq k\}$ for $k = 1, \ldots, n$. Let $f_k : R_k \to P_k$ be the mapping where $f_k(\rho_i) = \sigma_i$ and $f_k(\bar{\rho}_i) = \bar{\sigma}_i$, $i = 1, \ldots, k$. For $k = 1, \ldots, n$,*

a. ρ_k is a rotation and

b. f_k is an isomorphism from $\mathcal{R}_k$ to $\mathcal{P}_k$ where $\mathcal{R}_k$ is the subposet of $\mathcal{R}'(I(\mathcal{P}))$ induced by R_k and $\mathcal{P}_k$ is the subposet of $\mathcal{P}$ induced by P_k.

Proof When $k = 1$, it is easy to verify that both $\bar{\rho}_1$ and ρ_1 are exposed in the table A_1 and are not comparable – which is the case with $\bar{\sigma}_1$ and σ_1. So suppose the theorem is true for indices $1, \ldots, k-1$. Let us now prove it for index k. Our proof will consist of three steps. Let $\mathcal{R}_k - \{\rho_k\}$ be the subposet of $\mathcal{R}_k$ when ρ_k is removed, and let $\mathcal{P}_k - \{\sigma_k\}$ be the subposest of $\mathcal{P}_k$ when σ_k is removed. First, we will show that $\mathcal{R}_k - \{\rho_k\}$ is isomorphic to $\mathcal{P}_k - \{\sigma_k\}$. Next, we will argue that ρ_k is a rotation of $I(\mathcal{P})$. Finally, we will prove that $\mathcal{R}_k$ is isomorphic to $\mathcal{P}_k$.

Step 1. Since by our induction hypothesis $\mathcal{R}_{k-1}$ and $\mathcal{P}_{k-1}$ are isomorphic, we simply need to prove that $\bar{\rho}_k$ is an immediate predecessor of ϕ, $\phi \in R_{k-1}$, if and only if $\bar{\sigma}_k$ is an immediate predecessor of $f_{k-1}(\phi)$ (i.e., $\langle \bar{\sigma}_k, f_{k-1}(\phi)\rangle$ is an edge in $H(\mathcal{P})$). We consider two cases.

Case 1: $\phi = \bar{\rho}_j$ for some $j < k$.

If $\bar{\rho}_k$ is an immediate predecessor of $\bar{\rho}_j$, then $X(\bar{\rho}_k) \cap X(\bar{\rho}_j) = Y(\rho_k) \cap Y(\rho_j) \neq \emptyset$. That is, there is a participant that lies in $Y(\rho_k)$ and $Y(\rho_j)$. By Lemma 11, this means that there is a directed path from σ_j to σ_k in $H(\mathcal{P})$. Since $\mathcal{P}$ is a mirror poset, there is also a directed path from $\bar{\sigma}_k$ to $\bar{\sigma}_j$. Suppose $\bar{\sigma}_k$ is not an immediate predecessor of $\bar{\sigma}_j$ so there exists some $\bar{\sigma}_t$, $j < t < k$, so that $\langle \bar{\sigma}_k, \bar{\sigma}_t\rangle$ is an edge of $H(\mathcal{P})$, and $\bar{\sigma}_t < \bar{\sigma}_j$. By Lemma 12, $\bar{\rho}_k$ precedes $\bar{\rho}_t$. By the induction hypothesis, $\bar{\rho}_t$ precedes $\bar{\rho}_j$. These two statements imply that $\bar{\rho}_k$ is not an immediate predecessor of $\bar{\rho}_j$, a contradiction. Thus, $\langle \bar{\sigma}_k, \bar{\sigma}_j\rangle$ is an edge in $H(\mathcal{P})$.

On the other hand, suppose $\langle \bar{\sigma}_k, \bar{\sigma}_j\rangle$ is an edge of $H(\mathcal{P})$. From Lemma 12, $\bar{\rho}_k$ precedes $\bar{\rho}_j$. If $\bar{\rho}_k$ is not an immediate predecessor of $\bar{\rho}_j$, then there exists a $\bar{\rho}_t$, $j < t < k$, so that $\bar{\rho}_k$ immediately precedes $\bar{\rho}_t$ which precedes $\bar{\rho}_j$. From the previous paragraph, $\langle \bar{\sigma}_k, \bar{\sigma}_t\rangle$ is an edge of $H(\mathcal{P})$. From the induction hypothesis, $\bar{\sigma}_t$ precedes $\bar{\sigma}_j$. Again, these two statements imply that $\bar{\sigma}_k$ is not an immediate predecessor of $\bar{\sigma}_j$, a contradiction. So $\bar{\rho}_k$ is an immediate predecessor of $\bar{\rho}_j$.

Case 2: $\phi = \rho_j$ for some $j < k$.

Assume $\bar{\rho}_k$ is an immediate predecessor of ρ_j. Then $X(\bar{\rho}_k) \cap X(\rho_j) = Y(\rho_k) \cap X(\rho_j) \neq \emptyset$. According to the definition of the X-set of ρ_j and to Lemma 11, every participant in $Y(\rho_k) \cap X(\rho_j)$ has to be of the form $p(\sigma_j, \bar{\sigma}_t)$, where $\bar{\sigma}_t$ is a neighbor of σ_j. Additionally, there is a directed path from σ_t to σ_k in $H(\mathcal{P})$ so $t \leq k$. If $p(\sigma_j, \bar{\sigma}_k) \in Y(\rho_k) \cap X(\rho_j)$, we are done as this immediately implies that $\langle \bar{\sigma}_k, \sigma_j\rangle$ is an edge of $H(\mathcal{P})$. Otherwise, since $H(\mathcal{P})$ is a mirror poset, there is a directed path from $\bar{\sigma}_k$ to $\bar{\sigma}_t$. From Case 1, this implies that $\bar{\rho}_k$ precedes $\bar{\rho}_t$. Since $\langle \bar{\sigma}_t, \sigma_j\rangle$ is an edge of $H(\mathcal{P})$, by Lemma 12 $\bar{\rho}_t$ precedes ρ_j. Thus, $\bar{\rho}_k$ is not an immediate predecessor of ρ_j, a contradiction. Hence, it has to be the case that $\bar{\sigma}_k$ is an immediate predecessor of σ_j.

Suppose $\langle \bar{\sigma}_k, \sigma_j\rangle$ is an edge of $H(\mathcal{P})$. By Lemma 12, $\bar{\rho}_k$ precedes ρ_j. If $\bar{\rho}_k$ is not an immediate predecessor of ρ_j, then, as in Case 1, there is a rotation η so that $\bar{\rho}_k$ immediately precedes η which precedes ρ_j. If $\eta = \rho_t$, $t < k$, $\langle \bar{\sigma}_k, \sigma_t\rangle$ is an edge of $H(\mathcal{P})$ from the previous paragraph. If $\eta = \bar{\rho}_t$, $t < k$, $\langle \bar{\sigma}_k, \bar{\sigma}_t\rangle$ is an edge of $H(\mathcal{P})$ from Case 1. Furthermore, from the induction hypothesis, η is a predecessor of σ_j whether $\eta = \rho_t$ or $\bar{\rho}_t$. Since this contradicts our assumption that $\bar{\sigma}_k$ is an immediate predecessor of σ_j, η must not exist and $\bar{\rho}_k$ is an immediate precedessor ρ_j.

Step 2. From Lemma 9, $A_k = A_n/\{\bar{\rho}_n, \bar{\rho}_{n-1}, \ldots, \bar{\rho}_{k+1}\}$. Suppose Φ_k consists of all rotations $\phi \in R_k - \{\rho_k\}$ such that $f_k(\phi)$ is a predecessor of σ_k. Let a topological ordering of the elements in Φ_k be $\phi_1, \phi_2, \ldots, \phi_{i_k}$. For each $\phi_j \in \Phi_k$, notice that a predecessor of ϕ_j either belongs to $\{\bar{\rho}_n, \bar{\rho}_{n-1}, \ldots, \bar{\rho}_{k+1}\}$ or to $\{\phi_1, \ldots, \phi_{j-1}\}$ because there is an isomorphism between $\mathcal{R}_k - \{\rho_k\}$ and $\mathcal{P}_k - \{\sigma_k\}$. Furthermore, $\bar{\rho}_k \notin \Phi_k$ since it would mean that $\bar{\sigma}_k$ precedes σ_k. Thus, we can eliminate the elements of Φ_k from A_k starting from ϕ_1 all the way to ϕ_{i_k}. Let A'_k be the resulting table; i.e., $A'_k = A_k/\Phi_k$. Let us now show that ρ_k is exposed in A'_k.

Consider any two consecutive pairs in ρ_k. Assume they are $(p(\sigma_k, \alpha), y)$ and $(p(\sigma_k, \alpha'), y')$ where both α and α' are neighbors of σ_k. Thus, the last two persons in $p(\sigma_k, \alpha)$'s list in A_k are y followed by y', and $(y', p(\sigma_k, \alpha))$ belongs to $\bar{\rho}_k$. Here are the different possibilities for α and $p(\sigma_k, \alpha)$'s list in A'_k.

Case 1: $\alpha = \sigma_t$ or $\bar{\sigma}_t$, $t > k$. For both possibilities, $p(\sigma_k, \alpha)$ had only one person in his list prior to iteration k of `construct-instance`: $p(\bar{\sigma}_k, \bar{\alpha})$. During iteration k, another person was placed at the end of his list. Hence, $p(\sigma_k, \alpha)$'s list in A_k consists of $y = p(\bar{\sigma}_k, \bar{\alpha})$ followed by y', the person added during iteration k. Since neither $\bar{\rho}_k$ nor ρ_k was eliminated from A'_k, the list of $p(\sigma_k, \alpha)$ in A'_k still consists only of y followed by y'.

Case 2: $\alpha = \sigma_t$, $t < k$. Building on Case 1, it is easy to verify that the list of $p(\sigma_k, \alpha)$ in A_k consists of $p(\bar{\sigma}_k, \bar{\alpha})$ followed by y and then by y', where y was added during iteration t and y' during iteration k. Furthermore, $(p(\sigma_k, \alpha), p(\bar{\sigma}_k, \bar{\alpha}))$ belongs to ρ_t. By assumption, $\alpha = \sigma_t$ is a neighbor of σ_k. Since $t < k$, σ_t must also precede σ_k so $\rho_t \in \Phi_k$. Hence, ρ_t has been eliminated prior to the creation of A'_k, and the list of $p(\sigma_k, \alpha)$ in A'_k just consists of y followed by y'.

Case 3: $\alpha = \bar{\sigma}_t$, $t < k$. Prior to iteration t, $p(\sigma_k, \alpha)$ had only $p(\bar{\sigma}_k, \bar{\alpha})$ is his list. Between iterations t and $k-1$, several other persons may have been added to the front of the list because $p(\sigma_k, \alpha)$ belonged to some of the Y-sets of rotations in $\{\rho_t, \rho_{t+1}, \ldots, \rho_{k-1}\}$. Then during iteration k, another person was added to the end of the list. Hence, $y = p(\bar{\sigma}_k, \bar{\alpha})$, and y' is this last person added to the list. Let $p(\sigma_k, \alpha)$'s list in A_k consist of $z_{s+1}, z_s, \ldots, z_2, z_1 = y, y'$. Let $(z_i, p(\sigma_k, \alpha))$ belong to ρ_{j_i} so $(p(\sigma_k, \alpha), z_{i+1})$ belongs to $\bar{\rho}_{j_i}$ for $i = 1, \ldots, s$. According to Lemma 11, $\sigma_{j_1}, \sigma_{j_2}, \ldots, \sigma_{j_s}$ forms a directed path in $H(\mathcal{P})$, and $\sigma_{j_1} = \sigma_t$. Therefore, $\bar{\sigma}_{j_s}, \ldots, \bar{\sigma}_{j_2}, \bar{\sigma}_{j_1} = \bar{\sigma}_t$ is also a directed path in $H(\mathcal{P})$. Since $\alpha = \bar{\sigma}_t$ is a neighbor of σ_k, $\bar{\sigma}_t$ must precede σ_k. It follows that $\bar{\rho}_{j_s}, \ldots, \bar{\rho}_{j_2}, \bar{\rho}_{j_1} = \bar{\rho}_t$ all belong to Φ_k, and all have been eliminated prior to the construction of A'_k. Once again, the list of $p(\sigma_k, \alpha)$ in A'_k just consists of y followed by y'.

We have shown that for all possibilities of α, $p(\sigma_k, \alpha)$'s list in A'_k consists only of y and y'. Since α was chosen arbitrarily, ρ_k is exposed in A'_k.

Step 3. To prove that $\mathcal{R}_k$ is isomorphic to $\mathcal{P}_k$, we now have to show that for every $\phi \in R_k - \{\rho_k\}$, ϕ is an immediate predecessor of ρ_k if and only if $f_r(\phi)$ is an immediate predecessor of σ_k.

Since ρ_k is a rotation of $I(\mathcal{P})$ by step 2, according to Lemma 3 $\mathcal{R}_k$ is a mirror poset. Thus, ϕ is an immediate predecessor of ρ_k if and only if $\bar{\rho}_k$ is an immediate predecessor of $\bar{\phi}$. By step 1, the latter is true if and only if $\bar{\sigma}_k$ is an immediate predecessor of $f_r(\bar{\phi})$. But since $\mathcal{P}_k$ is a mirror poset, the previous statement is true if and only if $f_r(\phi)$ is an immediate predecessor of σ_k.

By induction, we have now shown that the theorem is true. □

Theorem 9 *Let $\mathcal{P}$ be a mirror poset with $2n$ elements. There is an SR instance $I(\mathcal{P})$ with $O(n^2)$ participants such that $\mathcal{R}'(I(\mathcal{P}))$ is isomorphic to $\mathcal{P}$. Additionally, when the dual of each element in $\mathcal{P}$ is given, $I(\mathcal{P})$ can be constructed in $O(n^2)$ time.*

Proof Since $\mathcal{R}'(I(\mathcal{P})) = \mathcal{R}_n$ and $\mathcal{P} = \mathcal{P}_n$ in Theorem 8, it follows that $\mathcal{R}'(I(\mathcal{P}))$ is isomorphic to $\mathcal{P}$. The number of participants in $I(\mathcal{P})$ equals the number of edges in $H(\mathcal{P})$ and so is $O(n^2)$.

Prior to `construct-instance`, a complete closed subset S of $\mathcal{P}$ is needed. This can be done as follows. First, do a topological ordering of the elements of P. Suppose the result is $\tau_1, \tau_2, \ldots, \tau_{2n}$. Initialize S to the empty set. Then greedily add elements into S; that is, for $i = 1$ to $2n$ if the dual of τ_i is not in S, add τ_i to S. Clearly, at the end of $2n$ iterations S is a complete subset of $\mathcal{P}$. To verify that it is also closed, suppose τ_i is a predecessor of $\tau_j \in S$. This means that τ_i occurs before τ_j in the topolgocial ordering. If $\tau_i \notin S$, $\bar{\tau}_i \in S$ and occurs before τ_i in the topological ordering. And since $\mathcal{P}$ is a mirror poset, $\bar{\tau}_j$ is a predecessor of $\bar{\tau}_i$ and so must occur before it in the topological ordering. Hence, the dual of τ_j occurs before τ_j in the topological ordering and should have been added to S, not τ_j – a contradiction. It follows that $\tau_i \in S$. The topological ordering of the elements of $\mathcal{P}$ takes $O(n^2)$ time. When the dual of each element in $\mathcal{P}$ is given, the greedy method described above takes $O(n)$ time. Hence, finding a complete closed subset of $\mathcal{P}$ takes $O(n^2)$ time.

Finally, iteration 0 of `construct instance` takes $O(n^2)$ time. For $i = 1, \ldots, n$, iteration i takes $O(|N(i)|)$ time, and iteration $n+1$ takes $O(n^2)$ time. Since $\sum_{i=1}^{n} |N(i)| = O(n^2)$, it follows that constructing $I(\mathcal{P})$ when the dual of each element in $\mathcal{P}$ is known takes $O(n^2)$ time. $\square$

6 Putting it all together

In Section 3, we showed that mirror posets give rise to median graphs. We will now prove that median graphs also give rise to mirror posets. Combining this result with Theorem 9, we will then establish the duality between median graphs and SR stable matchings.

Recall that the interval between u and v in a graph $G = (V, E)$ is the set $I(u, v) = \{t \in V : d(u, v) = d(u, t) + d(t, v)\}$. For each $u \in V$, define the *canonical order* $\leq_u$ as follows: for any $a, b \in V$, $a \leq_u b$ if $a \in I(u, b)$. The poset $(V, \leq_u)$ can then be thought as u's view of the graph G. There is a structure that is intimately related with median graphs.

Definition 4 *A median semilattice $\mathcal{Q} = (Q, \leq)$ is a meet semilattice (i.e., the greatest lower bound of any two elements always exists) so that (i) for any $\rho \in Q$, $\{\sigma \in Q : \sigma \leq \rho\}$ is a distributive lattice, and (ii) any three elements have a unique least upper bound whenever every pair does.*

In a poset $\mathcal{P} = (P, \leq)$, ρ *covers* σ if for any τ such that $\sigma \leq \tau \leq \rho$, $\tau = \sigma$ or ρ. In the *covering graph* of $\mathcal{P}$, P is the set of vertices and two elements σ and ρ are adjacent if and only if σ covers ρ or vice versa. Equivalently, the covering graph of $\mathcal{P}$ is the undirected Hasse diagram of $\mathcal{P}$.

Theorem 10 *(Avann [3]) The covering graph of any median semilattice is a median graph. Conversely, every median graph gives rise to a median semilattice with respect to any canonical order $\leq_u$, where u is a vertex of the graph.*

Thus, when I is an SR instance, instead of the graph $G(M(I))$, we can use $(M(I), \leq_\mu)$, $\mu \in M(I)$, as the ordered structure underlying the stable matchings of I. Interestingly, if I is also an SM instance and $\mu = \mu_M$, the man-optimal stable matching of I, $(M(I), \leq_{\mu_M})$ becomes a distributive lattice because the woman-optimal stable matching of I is the unique stable matching that is farthest from μ_M in $G(M(I))$. In fact, $(M(I), \leq_{\mu_M})$ is *the* distributive lattice that we referred to in the introduction of the paper. A similar observation was made in [14] except that the elements of the semilattice were modifications of the matchings in $M(I)$.

We also note the following property of median graphs and semilattices.

Lemma 13 *Suppose $G = (V, E)$ is a median graph and $u \in V$. The covering graph of the median semilattice $(V, \leq_u)$ is G.*

Proof The lemma stems from the observation that if $d(u, v) = d(u, w)$ then $(v, w) \notin E$. To see this, let P_1 and P_2 be shortest paths from u to v and u to w respectively. Let z be the last node in P_1 that is also in P_2. Notice that the $u - z$ subpaths in P_1 and in P_2 must have the same length because otherwise P_1 and P_2 are not shortest paths. Hence, the $z - v$ subpath in P_1 has the same length as the $z - w$ subpath in P_2. If (v, w) exists, then this edge together with the $z - v$ and $z - w$ subpaths form an odd cycle in G. But G is bipartite because it is a median graph. Thus, (v, w) does not exist.

The above observation implies that if $(v, w) \in E$, either v is closer to u than w or vice versa. If it is the former, w covers v since a shortest path from u to v together with (v, w) is a shortest path from u to w. If it is the latter, v covers w for the same reason. Hence, (v, w) is in the covering graph of $(V, \leq_u)$. $\square$

An element of a poset is *join-irreducible* if it covers exactly one element of the poset. The next lemma shows how mirror posets can be derived from median graphs.

Lemma 14 *Let $G = (V, E)$ be a median graph and $u \in V$. In the median semilattice $(V, \leq_u)$, let J be the set of join-irreducible elements. Let $\bar{J} = \{\bar{v} | v \in J\}$. For each $v \in J$, let v and $\bar{v}$ be duals of each other. Define the relation $\leq'$ on $J \cup \bar{J}$ as follows:*

(i) for any pair $v, w \in J$, if $v \leq_u w$ let $v \leq' w$ and $\bar{w} \leq' \bar{v}$, and

(ii) for any pair $v, w \in J$, if v and w do not have an upper bound in $(V, \leq_u)$ let $\bar{v} \leq' w$ and $\bar{w} \leq' v$.

The relation $\leq'$ is a partial order and $(J \cup \bar{J}, \leq')$ is a mirror poset.

Proof For each $v \in J$, since the dual of $\bar{v}$ is v, we define $\bar{\bar{v}} = v$. Note that (1) $\leq_u$ is also a partial order, and (2) in the ordered relation $\leq'$, elements from J are never related to elements in $\bar{J}$; only elements from $\bar{J}$ are related to elements in J. Since $\leq_u$ is reflexive, $\leq'$ is also reflexive. Facts (1) and (2) also imply that if $a \leq' b$ and $b \leq' a$ for any two elements $a, b \in J \cup \bar{J}$, $a = b$; that is, $\leq'$ is anti-symmetric. Finally, suppose $a \leq' b$ and $b \leq' c$. If $b \in J$, $c \in J$ by fact (2). If, additionally, $a \in J$, then $a \leq' c$ since $\leq_u$ is also transitive. If $a \in \bar{J}$, $a \leq' b$ means that a's corresponding element $\bar{a} \in J$ and b have no upper bound in $(V, \leq_u)$. Since $b \leq_u c$, $\bar{a}$ and c cannot have an upper bound in $(V, \leq_u)$ either. Thus, $a \leq' c$. The case when $b \in \bar{J}$ is proved similarly. Hence, $\leq'$ is transitive. Since $\leq'$ is reflexive, anti-symmetric, and transitive, $\leq'$ is a partial order.

We also note that $v \not\leq' \bar{v}$ because of Fact 2, and $\bar{v} \not\leq' v$ because the negation of this will imply that v and v have no upper bound – an obvious contradiction. By construction, $(J \cup \bar{J}, \leq')$ also has the property that whenever $a, b \in J \cup \bar{J}$ and $a \leq' b$, it follows that $\bar{b} \leq' \bar{a}$. Hence, $(J \cup \bar{J}, \leq')$ is a mirror poset. □

Example continued. Suppose we start with the median graph shown in Figure 2. Set $u = u_i$, for some $1 \leq i \leq 6$, so that $(V, \leq_u)$ is a median semilattice. All other nodes in V are join-irreducible elements of $(V, \leq_u)$. We leave it up to the reader to verify that $(J \cup \bar{J}, \leq')$ is a mirror poset that is isomorphic to the one shown in Figure 1.

According to Corollary 1, the graph induced by the complete closed subsets of $(J \cup \bar{J}, \leq')$, $G(\mathcal{S}_{(J \cup \bar{J}, \leq')})$, is also a median graph. What we would like to prove next is that this graph is in fact isomorphic to G. Instead of arguing it from scratch, however, we make use of a duality result by Barthélemy and Constantin [6].

Let $\mathcal{Q}$ be a median semilattice. Again, denote its set of join-irreducible elements by J. Let $(J, \leq, E_{\mathcal{Q}})$ be the triple where $(J, \leq)$ is the subposet of $\mathcal{Q}$ induced by J, and $(J, E_{\mathcal{Q}})$ is the graph where two elements of J are adjacent if and only if they have no upper bound in $\mathcal{Q}$. Let $\mathcal{T}_{(J, \leq, E_{\mathcal{Q}})}$ consist of all subsets of J that are closed under $(J, \leq)$ and form an independent set in $(J, E_{\mathcal{Q}})$ (i.e., no two elements of the set are adjacent in $(J, E_{\mathcal{Q}})$).

Theorem 11 *(Barthélemy and Constantin [6]) Let $\mathcal{Q} = (Q, \leq)$ be a median semilattice and J its set of join-irreducible elements. Then $(\mathcal{T}_{(J, \leq, E_{\mathcal{Q}})}, \subseteq)$ is a median semilattice that is isomorphic to $\mathcal{Q}$. Additionally, in the covering graph of $(\mathcal{T}_{(J, \leq, E_{\mathcal{Q}})}, \subseteq)$, two subsets are adjacent if and only if they differ by one element.*

Birkhoff's representation theorem for distributive lattices [7] states that given a distributive lattice $\mathcal{D}$, the closed subsets of the poset induced by its join-irreducible elements form a distributive lattice that is isomorphic to $\mathcal{D}$. Barthélemy and Constantin showed that the join-irreducible elements of a median semilattice $\mathcal{Q}$ play a similar but more complicated role.

Theorem 12 *Let $G = (V, E)$ be a median graph and $u \in V$. In the median semilattice $(V, \leq_u)$, let J be the set of join-irreducible elements. Define the mirror poset $(J \cup \bar{J}, \leq')$ as in Lemma 14. The graph $G(\mathcal{S}_{(J\cup\bar{J},\leq')})$ is isomorphic to G.*

Proof Let $\mathcal{Q} = (V, \leq_u)$. Since $\mathcal{Q}$ is a median semilattice, the set $\mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})}$ is well defined. Let's begin by proving the following claim.

Claim: There is a one-to-one correspondence between the subsets in $\mathcal{T}_{(J,\leq,E_{\mathcal{Q}})}$ and the subsets in $S_{(J\cup\bar{J},\leq')}$.
Proof of claim: Let f be a function from $\mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})}$ to $\mathcal{S}_{(J\cup\bar{J},\leq')}$ such that for each $T \in \mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})}$, $f(T) = S$ where $S = T \cup \{\bar{t} \in \bar{J} : t \notin T\}$. First, we need to verify that S indeed belongs to $\mathcal{S}_{(J\cup\bar{J},\leq')}$. Notice that S is a complete subset of $(J \cup \bar{J}, \leq')$ since every element of J or its dual is in S. Let us prove that S is closed as well. Let $v \in J$. Suppose $v \in T$ and $w \leq' v$. If $w \in J$, then $w \in T$ because T is a closed subset of $(J, \leq_u)$. If $w \in \bar{J}$, then $\bar{w} \in J$ and v have no upper bound in $\mathcal{Q}$. This means that $\bar{w}$ and v share an edge in $(J, E_{\mathcal{Q}})$. Thus, $\bar{w} \notin T$ so $w \in S$. On the other hand, suppose $v \notin T$ so that $\bar{v} \in S$ and $x \leq' \bar{v}$. Clearly, $x \in \bar{J}$. Since $(J \cup \bar{J}, \leq')$ is a mirror poset, $v \leq' \bar{x}$. But $v \notin T$ and T is a closed subset of $(J, \leq_u)$ so $\bar{x} \notin T$ as well. Consequently, $x \in S$.

Now, for each $S \in \mathcal{S}_{(J\cup\bar{J},\leq')}$, the set $S \cap J$ is a closed subset of $(J, \leq_u)$ because S is a closed subset of $(J \cup \bar{J}, \leq')$. Additionally, if $v, w \in J$ do not have an upper bound in $\mathcal{Q}$, $\bar{v} \leq' w$ and $\bar{w} \leq' v$. Since S is a complete subset of $(J \cup \bar{J}, \leq')$, both v and w cannot belong to S. Thus, $S \cap J$ is also an independent subset of $(J, E_{\mathcal{Q}})$ so $S \cap J \in \mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})}$. This implies that f is onto since $S \cap J$ is a pre-image of S under f. It is also 1-1 because whenever $T \neq T'$, $f(T) \neq f(T')$. It follows that f is a bijection from $\mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})}$ to $\mathcal{S}_{(J\cup\bar{J},\leq')}$.

Let $G(\mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})})$ denote the graph induced by the independent closed subsets of $(J, \leq_u, E_{\mathcal{Q}})$ where two subsets are adjacent if and only if they differ by one element. Notice that $G(\mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})})$ is the covering graph of the median semilattice $(\mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})}, \subseteq)$. Furthermore, in the proof of the claim above two subsets T and T' differ by one element if and only if $f(T)$ and $f(T')$ differ in one dual element so $G(\mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})})$ is isomorphic to $G(\mathcal{S}_{(J\cup\bar{J},\leq')})$.

According to Theorem 11, $\mathcal{Q}$ is isomorphic to $(\mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})}, \subseteq)$ so their respective covering graphs must also be isomorphic to each other. Since they are G and $G(\mathcal{T}_{(J,\leq_u,E_{\mathcal{Q}})})$ respectively, it follows that G and $G(\mathcal{S}_{(J\cup\bar{J},\leq')})$ are also isomorphic. □

Theorem 13 *Let $G = (V, E)$ be a median graph. There is an SR instance $I(G)$ so that the graph of its stable matchings is isomorphic to G.*

Proof Pick an arbitrary vertex $u \in V$. Let $\mathcal{P}_G = (J \cup \bar{J}, \leq')$ be the mirror poset defined in Lemma 14. Construct the SR instance $I(\mathcal{P}_G)$ using `construct-instance` in Section 5. According to Theorem 9, its reduced rotation poset $\mathcal{R}'(I(\mathcal{P}_G))$ is isomorphic to $\mathcal{P}_G$. Now both of these posets are mirror posets so the graphs that arise from their complete closed subsets are also isomorphic. For $\mathcal{R}'(I(\mathcal{P}_G))$, the said graph is isomorphic to the graph of stable matchings of $I(\mathcal{P}_G)$ by definition. For $\mathcal{P}_G$, the said graph is isomorphic to G according to Theorem 12. Thus, if we let $I(G) = I(\mathcal{P}_G)$, then $G(M(I(G)))$ is isomorphic to G. □

7 Conclusion

We have shown that the local/global median phenomenon that was first observed in SM instances also occurs for solvable SR instances by proving that the underlying structure that governs SR stable matchings is a median graph. Earlier results on medians in median graphs also imply that these median stable matchings have many other nice properties. Unfortunately, finding a median stable matching of an SM instance is #P-hard [8, 9]. Interesting research directions include characterizing SM and SR instances where the search problem is easy or developing algorithms that find stable matchings that approximate the median stable matchings. Some work has been done for SM instances [9, 17]; a lot more seems possible since our characterization suggests two interpretations of the median stable matchings.

We also showed that three structures – SR stable matchings, mirror posets, and median graphs are pairwise duals of each other. These results can also be inferred from Feder's work [10, 11, 12]. However, our constructions and proofs are smooth generalizations of the ones used for SM, which makes them easier to follow.

References

[1] A. Abdulkadiroglu, P. Pathak, and A. Roth. The New York City high school match. *American Economic Review, Papers and Proceedings*, 95:364–367, 2005.

[2] A. Abdulkadiroglu, P. Pathak, A. Roth, and T. Sönmez. The Boston public school match. *American Economic Review, Papers and Proceedings*, 95:368–371, 2005.

[3] S. Avann. Metric ternary distributive semi-lattices. *Proceedings of the American Mathematical Society*, 12:407–414, 1961.

[4] H. Bandelt and J. Barthélemy. Medians in median graphs. *Discrete Applied Mathematics*, 8:131–142, 1984.

[5] H. Bandelt and V. Chepoi. Metric graph theory and geometry: a survey. *Contemporary Mathematics*, pages 49–86, 2008.

[6] J. Barthélemy and J. Constantin. Median graphs, parallelism and posets. *Discrete Mathematics*, 111:49–63, 1993.

[7] G. Birkhoff. Rings of sets. *Duke Mathematical Journal*, 3:443–454, 1937.

[8] C. Cheng. The generalized median stable matchings: finding them is not that easy. In *Proceedings of the 8th Latin Theoretical Informatics Conference*, pages 568–579, 2008.

[9] C. Cheng. Understanding the generalized median stable matchings. Accepted to Algorithmica, 2009.

[10] T. Feder. A new fixed-point approach for stable networks and stable marriages. In *Proceedings of the 21st ACM Symposium on Theory of Computing*, pages 513–522, 1989.

[11] T. Feder. A new fixed-point approach for stable networks and stable marriages. *Journal of Computer and System Sciences*, 45:233–284, 1992.

[12] T. Feder. *Stable networks and product graphs*, volume 116 of *Memoirs of the American Mathematical Society.* AMS Bookstore, 1995.

[13] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69, 1962.

[14] D. Gusfield and R. Irving. *The Stable Marriage Problem: Structure and Algorithms.* The MIT Press, 1989.

[15] R. Irving. An efficient algorithm for the stable roommates problem. *Journal of Algorithms*, 6:577–595, 1985.

[16] R. Irving and P. Leather. The complexity of counting stable marriages. *SIAM Journal on Computing*, 15:655–667, 1986.

[17] S. Kijima and T. Nemoto. Randomized approximation for the generalize median stable matchings. A preprint, RIMS-1648, 2008.

[18] B. Klaus and F. Klijn. Smith and Rawls share a room: stability and medians. Meteor RM/07-026, 2008.

[19] S. Klavžar and H. Mulder. Median graphs: characterizations, location theory and related structures. *Journal of Combinatorial Mathematics and Combinatorial Computing*, 30:103–127, 1999.

[20] D. Knuth. *Mariages Stables.* Les Presses de l'Université de Montréal, 1976.

[21] D. Knuth. *The Art of Computer Programming*, volume IV. Addison-Wesley, 2008.

[22] H. Martyn and A. Schrijver. Median graphs and helly hypergraphs. *Discrete Mathematics*, 25:41–50, 1979.

[23] L. Nebeský. Median graphs. *Commentationes Mathematicae Universitatis Carolinae*, 12:317–325, 1971.

[24] A. Roth and E. Peranson. The redesign of the matching market of American physicians: Some engineering aspects of economic design. *American Economic Review*, 89:748–780, 1999.

[25] A. Roth and M. Sotomayor. *Two-sided matching: a study in game-theoretic modeling and analysis.* Cambridge University Press, 1990.

[26] A. Subramanian. A new approach to stable matching problems. *SIAM Journal on Computing*, 23:671–701, 1994.

[27] C.-P. Teo and J. Sethuraman. The geometry of fractional stable matchings and its applications. *Mathematics of Operations Research*, 23:874–891, 1998.