Counting Popular Matchings in House Allocation Problems

Rupam Acharyya, Sourav Chakraborty, Nitesh Jha

Chennai Mathematical Institute
Chennai, India
{rupam,sourav,nj}@cmi.ac.in

Abstract. We study the problem of counting the number of *popular matchings* in a given instance. McDermid and Irving gave a poly-time algorithm for counting the number of popular matchings when the preference lists are strictly ordered. We first consider the case of ties in preference lists. Nasre proved that the problem of counting the number of popular matching is #P-hard when there are ties. We give an FPRAS for this problem.

We then consider the popular matching problem where preference lists are strictly ordered but each house has a capacity associated with it. We give a *switching graph characterization* of popular matchings in this case. Such characterizations were studied earlier for the case of strictly ordered preference lists (McDermid and Irving) and for preference lists with ties (Nasre). We use our characterization to prove that counting popular matchings in capacitated case is #P-hard.

1 Introduction

A *popular matching problem* instance I comprises a set $\mathcal{A}$ of *agents* and a set $\mathcal{H}$ of *houses*. Each agent a in $\mathcal{A}$ ranks (numbers) a subset of houses in $\mathcal{H}$ (lower rank specify higher preference). The ordered list of houses ranked by $a \in \mathcal{A}$ is called a's *preference list*. For an agent a, let E_a be the set of pairs (a, h) such that the house h appears on a's preference list. Define $E = \cup_{a \in \mathcal{A}} E_a$. The problem instance I is then represented by a bipartite graph $G = (\mathcal{A} \cup \mathcal{H}, E)$. A *matching* M of I is a matching of the bipartite graph G. We use $M(a)$ to denote the house assigned to agent a in M and $M(h)$ to denote the agent that is assigned house h in M. An agent *prefers* a matching M to a matching M' if (i) a is matched in M and unmatched in M', or (ii) a is matched in both M and M' but a prefers the house $M(a)$ to $M'(a)$. Let $\phi(M, M')$ denote the number of agents that prefer M to M'. We say M is *more popular than* M' if $\phi(M, M') > \phi(M', M)$, and denote it by $M \succ M'$. A matching M is called *popular* if there exists no matching M' such that $M' \succ M$.

The popular matching problem was introduced in [5] as a variation of the stable marriage problem [4]. The idea of popular matching has been studied extensively in various settings in recent times [1, 14, 12, 10, 8, 11, 13], mostly in

the context where only one side has preference of the other side but the other side has no preference at all. We will also focus on this setting. Much of the earlier work focuses on finding efficient algorithms to output a popular matching, if one exists.

The problem of counting the number of "solutions" to a combinatorial question falls into the complexity class #P. An area of interest that has recently gathered a certain amount of attention is the problem of counting stable matchings in graphs. The Gale-Shapely algorithm [4] gives a simple and efficient algorithm to output a stable matching, but counting them was proved to be #P-hard in [6]. Bhatnagar, Greenberg and Randall [2] showed that the random walks on the *stable marriage lattice* are slowly mixing, even in very restricted versions of the problem. [3] gives further evidence towards the conjecture that there may not exist an FPRAS at all for this problem.

Our motivation for this study is largely due to the similarity of structures between stable matchings and popular matchings (although no direct relationship is known). The interest is further fueled by the existence of a linear time algorithm to exactly count the number of popular matchings in the standard setting [12]. We look at generalizations of the standard version - preferences with ties and houses with capacities. In the case where preferences could have ties, it is already known that the counting version is #P-hard [13]. We give an FPRAS for this problem. In the case where houses have capacities, we prove that the counting version is #P-hard. While the FPRAS for the case of ties is achieved via a reduction to a well known algorithm, the #P-hardness for the capacitated case is involving, making it the more interesting setting of the problem.

We now formally describe the different variants of the popular matching problem (borrowing the notation from [14]) and also describe our results alongside.

House Allocation problem (HA) These are the instances $G = (\mathcal{A} \cup \mathcal{H}, E)$ where the preference list of each agent $a \in \mathcal{A}$ is a linear order. Let $n = |\mathcal{A}| + |\mathcal{H}|$ and $m = |E|$. In [1], Abraham et al. give a complete characterization of popular matchings in an HA instance, using which they give an $O(m + n)$ time algorithm to check if the instance admits a popular matching and to obtain the largest such matching, if one exists. The question of counting popular matchings was first addressed in [12], where McDermid et al. give a new characterization by introducing a powerful structure called the *switching graph* of an instance. The switching graph encodes all the popular matchings via *switching paths* and *switching cycles*. Using this structure, they give a linear time algorithm to count the number of popular matchings.

House Allocation problem with Ties (HAT) An instance $G = (\mathcal{A} \cup \mathcal{H}, E)$ of HAT can have applicants whose preference list contains ties. For example, the preference list of an agent could be $[h_3, (h_1, h_4), h_2]$, meaning, house h_3 gets rank 1, houses h_1 and h_4 get a tied rank 2 and house h_2 gets the rank 3. A characterization for popular matchings in HAT was given in [1]. The characterization is used to give an $O(\sqrt{n}m)$ time algorithm to solve the maximum cardinality

popular matching problem. We outline their characterization briefly in Section 2 where we consider the problem of counting popular matchings in HAT. In [13], Nasre gives a proof of #P-hardness of this problem. We give an FPRAS for this problem by reducing it to the problem of counting perfect matchings in a bipartite graph.

Capacitated House Allocation Problem (CHA) A popular matching instance in CHA has a *capacity* c_i associated with each house $h_i \in \mathcal{H}$, allowing at most c_i agents to be matched to house h_i. The preference list of each agent is strictly ordered. A characterization for popular matchings in CHA was given in [14], along with an algorithm to find the largest popular matching (if one exists) in time $O(\sqrt{C}n_1 + m)$, where $n_1 = |\mathcal{A}|$, $m = |E|$ and C is the total capacity of the houses. In Section 3, we consider the problem counting popular matchings in CHA. We give a switching graph characterization of popular matchings in CHA. This is similar to the switching graph characterization for HA in [12]. Our construction is also motivated from [13], which gives a switching graph characterization of HAT. We use our characterization to prove that it is #P-Complete to compute the number of popular matchings in CHA.

Remark: A natural reduction exists from a CHA instance $G = (\mathcal{A} \cup \mathcal{H}, E)$ to an HAT instance. The reduction is as follows. Treat each house $h_i \in \mathcal{H}$ with capacity c as c different houses $h_i^1, \ldots, h_i^c$ of unit capacity, which are always tied together and appear together wherever h_i appears in any agent's preference list. Let the HAT instance thus obtained be G'. It is clear that every popular matching of G is a popular matching of G'. Hence, for example, an algorithm which finds a maximum cardinality popular matching for HAT can be used to find a maximum cardinality popular matching for the CHA instance G. In the context of counting, it is important to note that one popular matching of G may translate to many popular matchings in G'. It is not clear if there is a useful map between these two sets that may help in obtaining either hardness or algorithmic results for counting problems.

2 Counting in House Allocation Problem with Ties

In this section we consider the problem of counting the number of popular matchings in House Allocation problem with Ties (HAT). We first describe the characterization given in [1] here using similar notations. Let $G = (\mathcal{A} \cup \mathcal{H}, E)$ be an HAT instance. For any agent $a \in \mathcal{A}$, let $f(a)$ denote the set of first choices of a. For any house $h \in \mathcal{H}$, define $f(h) := \{a \in \mathcal{A}, f(a) = h\}$. A house h for which $f(h) \neq \phi$ is called an f-house. To simplify the definitions, we add a unique last-resort house $l(a)$ with lowest priority for each agent $a \in A$. This forces every popular matching to be an applicant complete matching.

Definition 1. (Section 3.1 in [1]). *The **first choice graph** of G is defined to be $G_1 = (\mathcal{A} \cup \mathcal{H}, E_1)$, where E_1 is the set of all rank one edges.*

Lemma 1. (Lemma 3.1 in [1]). *If M is a popular matching of G, then $M \cap E_1$ is a maximum matching of G_1.*

Let M_1 be any maximum matching of G_1. The matching M_1 can be used to identify the houses h that are always matched to an agent in the set $f(h)$. In this direction, we observe that M_1 defines a partition of the vertices $\mathcal{A} \cup \mathcal{H}$ into three disjoint sets - *even*, *odd* and *unreachable*: a vertex is *even* (resp. *odd*) if there is an even (resp. odd) length alternating path from an unmatched vertex (with respect to M_1) to v; a vertex v is *unreachable* if there is no alternating path from an unmatched vertex to v. Denote the sets *even*, *odd* and *unreachable* by $\mathcal{E}$, $\mathcal{O}$ and $\mathcal{U}$ respectively. The following is a well-known theorem in matching theory [9].

Lemma 2 (Gallai-Edmonds Decomposition). *Let G_1 and M_1 define the partition $\mathcal{E}$, $\mathcal{O}$ and $\mathcal{U}$ as above. Then,*

(a) The sets $\mathcal{E}$, $\mathcal{O}$ and $\mathcal{U}$ are pairwise disjoint, and every maximum matching in G_1 partitions the vertices of G_1 into the same partition of even, odd and unreachable vertices.
(b) In any maximum matching of G_1, every vertex in $\mathcal{U}$ is matched with another vertex in $\mathcal{U}$, and every vertex in $\mathcal{O}$ is matched with some vertex in $\mathcal{E}$. No maximum matching contains an edge between a vertex in $\mathcal{O}$ and a vertex in $\mathcal{O} \cup \mathcal{U}$. The size of a maximum matching is $|\mathcal{O}| + |\mathcal{U}|/2$.
(c) G_1 contains no edge connecting a vertex in $\mathcal{E}$ with a vertex in $\mathcal{U}$.

We show the decomposition of G_1 in Figure 1, where we look at the bipartitions of $\mathcal{U}$, $\mathcal{O}$, and $\mathcal{E}$ into their left and right parts, denoted by subscripts l and r respectively. Since G_1 only contained edges resulting from first-choices, every house in $\mathcal{U}_r$ and $\mathcal{O}_r$ is an f-house. From Lemma 2, each such house $h \in \mathcal{U}_r \cup \mathcal{O}_r$ is matched with an agent in $f(h)$ in every maximum matching of G_1, and correspondingly in every popular matching of G (Lemma 1).

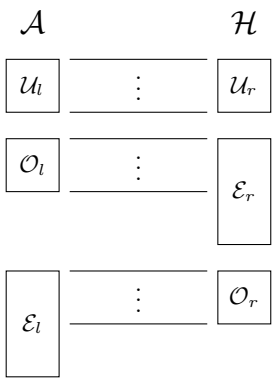

Fig. 1. Gallai-Edmonds decomposition of the first-choice graph of G

For each agent a, define $s(a)$ to be a's most preferred house(s) in $\mathcal{E}_r$. Note that $s(a)$ always exists after the inclusion of last-resort houses $l(a)$. The following is proved in [1].

Lemma 3. (Lemma 3.5 in [1]). *A matching M is popular in G if and only if*

1. *$M \cap E_1$ is a maximum matching of G_1, and*
2. *for each applicant a, $M(a) \in f(a) \cup s(a)$.*

The following hardness result is from [13].

Lemma 4. (Theorem 3 in [13]). *Counting the number of popular matchings in HAT is #P-hard.*

We now give an FPRAS for counting the number of popular matchings in the case of ties. As before, let $G = (\mathcal{A} \cup \mathcal{H}, E)$ be our HAT instance. We assume that that G admits at least one popular matching (this can be tested using the characterization). We reduce our problem to the problem of counting perfect matchings in a bipartite graph. We start with the first-choice graph G_1 of G, and perform a Gallai-Edmonds decomposition of G_1 using any maximum matching of G_1. In order to get a perfect matching instance, we extend the structure obtained from Gallai-Edmonds decomposition described in Figure 1. Let $\mathcal{F}$ be the set of f-houses and $\mathcal{S}$ be the set of s-houses. We make use of the following observations in the decomposition.

— Every agent in $\mathcal{U}_l$ and $\mathcal{O}_l$ gets one of their first-choice houses in every popular matching.
— $\mathcal{E}_r$ can be further partitioned into the following sets:
  – $\mathcal{E}_r^f := \{h \in \mathcal{F} \cap \overline{\mathcal{S}}, h \in \mathcal{E}_r\}$,
  – $\mathcal{E}_r^s := \{h \in \overline{\mathcal{F}} \cap \mathcal{S}, h \in \mathcal{E}_r\}$,
  – $\mathcal{E}_r^{f/s} := \{h \in \mathcal{F} \cap \mathcal{S}, h \in \mathcal{E}_r\}$, and
  – $\mathcal{E}_r^\star := \{h \notin \mathcal{F} \cup \mathcal{S}, h \in \mathcal{E}_r\}$.
— O_l can only match with houses in $\mathcal{E}_r^f \cup \mathcal{E}_r^{f/s}$ in every popular matching.

These observations are described in Figure 2(a).

Next, we observe that every agent in $\mathcal{E}_l$ that is already not matched to a house in $\mathcal{O}_r$, must match to a house in $\mathcal{E}_r^s \cup \mathcal{E}_r^{f/s}$. We facilitate this by adding all edges $(a, s(a))$ for each agent in $\mathcal{E}_l$. Finally, we add a set of dummy agent vertices $\mathcal{D}$ on the left side to balance the bipartition. The size of $\mathcal{D}$ is $|\mathcal{A}| - (|\mathcal{H}| - |\mathcal{E}_r^\star|)$. This difference is non-negative as long as the preference-lists of agents are complete. We make the bipartition $(\mathcal{D}, \mathcal{E}_r^f \cup \mathcal{E}_r^{f/s} \cup \mathcal{E}_r^s)$ a complete bipartite graph by adding the appropriate edges. This allows us to move from one popular matching to another by switching between first and second-choices and, among second choices of agents. Finally, we remove set $\mathcal{E}_r^\star$ from the right side. The new structure is described in Figure 2(b). Denote the new graph by G'.

Lemma 5. *The number of popular matchings in G is $|D|!$ times the number of perfect matchings in G'.*

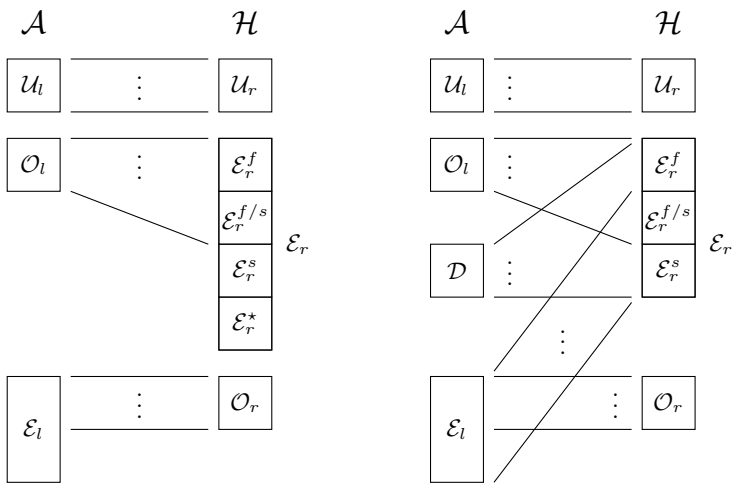

Fig. 2(a) **Fig. 2(b)**

Fig. 2. Reduction to a perfect-matching instance by extending the Gallai-Edmonds decomposition of G_1.

Proof. Consider a perfect matching M of G'. Let the matching M' be obtained by removing from M all the edges coming out of the set $\mathcal{D}$. Observe that $M' \cap E_1$ is a maximum matching of G_1. This is because the sets $\mathcal{U}_l$, O_l and $\mathcal{O}_r$ are always matched in M' (or else M would not be a perfect matching of G') and that the size of a maximum matching in G_1 is $(|\mathcal{U}_l| + |\mathcal{O}_l| + |\mathcal{O}_r|)$ by Lemma 2. Also, each agent in $\mathcal{A}$ is matched to either a house in $\mathcal{F}$ or in $\mathcal{S}$ by the construction of graph G'. Using Lemma 3, we conclude that M is a popular matching of G. Finally, observe that every popular matching in M in G can be augmented to a perfect matching of G' by adding exactly $|\mathcal{D}|$ edges. This follows again from Lemma 2 and Lemma 3. □

We now make use of the following result of Jerrum et al. from [7].

Lemma 6. (Theorem 1.1 in [7]). *There exists an FPRAS for the problem of counting number of perfect matchings a bipartite graph.*

From Lemma 5 and Lemma 6, we have the following.

Theorem 1. *There exists an FPRAS for counting the number of popular matchings in the House Allocation problem with Ties.*

3 Counting in Capacitated House Allocation Problem

In this section, we consider the structure of popular matchings in Capacitated House Allocation problem (CHA). A CHA instance I consists of agents $\mathcal{A}$ and

houses $\mathcal{H}$. Let $|\mathcal{A}| = n$ and $|\mathcal{H}| = m$. Let $c : \mathcal{H} \to \mathbb{Z}_{>0}$ be the capacity function for houses. Each agent orders a subset of the houses in a strict order creating its *preference list*. The preference list of $a_i \in \mathcal{A}$ defines a set of edges E_i from a_i to houses in $\mathcal{H}$. Define $E = \cup_{i \in [n]} E_i$. The problem instance I can then be represented by a bipartite graph $G = (\mathcal{A} \cup \mathcal{H}, E)$.

For the instance I, a *matching* M is a subset of E such that each agent appears in at most one edge in M and each house h appears in at most $c(h)$ edges in M. The definitions of *more popular than* relationship between two matchings and *popular matching* is same as described earlier in Section 1.

We now outline a characterization of popular matchings in CHA from [14]. As before, denote by $f(a)$ the first choice of an agent $a \in \mathcal{A}$. A house which is the first choice of at least one agent is called an f-house. For each house $h \in \mathcal{H}$, define $f(h) = \{a \in \mathcal{A}, f(a) = h\}$. For each agent $a \in \mathcal{A}$, we add a unique last-resort house $l(a)$ with least priority and capacity 1.

Lemma 7. (Lemma 1 in [14]) *If M is a popular matching then for each f-house h, $|M(h) \cap f(h)| = \min\{c(h), |f(h)|\}$.*

For each agent $a \in \mathcal{A}$, define $s(a)$ to be the highest ranked house h on a's preference list such that one of the following is true:

- h is not an f-house, or,
- h is an f-house but $h \neq f(a)$ and $|f(h)| < c(h)$.

Notice that $s(a)$ always exists after the inclusion of last-resort houses $l(a)$. The following lemma gives the characterization of popular matchings in G.

Lemma 8. (Theorem 1 in [14]) *A matching M is popular if and only if*

1. *for every f-house $h \in \mathcal{H}$,*
   - *if $|f(h)| \leq c(h)$, then every agent in $f(h)$ is matched to the house h,*
   - *else, house h is matched to exactly $c(h)$ agents, all belonging to $f(h)$,*
2. *M is an agent complete matching such that for each agent $a \in \mathcal{A}$, $M(a) \in \{f(a), s(a)\}$.*

3.1 Switching Graph Characterization of CHA

We now give a *switching graph* characterization of popular matchings for instances from this class. Our results are motivated from similar characterizations for HA in [12] and for HAT in [13]. A switching graph for an instance allows us to move from one popular matching to another by making well defined walks on the switching graph.

Consider a popular matching M of an instance G of CHA. The switching graph of G with respect to M is a directed weighted graph $G_M = (\mathcal{H}, E_M)$, with the edge set E_M defined as follows. For every agent $a \in \mathcal{A}$,

- add a directed edge from M(a) to $\{f(a), s(a)\} \setminus M(a)$,

– if $M(a) = f(a)$, assign a weight of -1 on this edge, otherwise assign a weight of $+1$.

Associated with the switching graph G_M, we have an *unsaturation degree* function $u_M : \mathcal{H} \to \mathbb{Z}_{\geq 0}$, defined $u_M(h) = c(h) - |M(h)|$. A vertex h is called *saturated* if its unsaturation degree is 0, i.e. $u_M(h) = 0$. If $u_M(h) > 0$, h is called *unsaturated*. We make use of the following terminology in the foregoing discussion. We now describe some useful properties of the switching graph G_M.

▷ *Property 1:* Each vertex h can have out-degree at most $c(h)$.
Proof. Each edge is from a matched house to an unmatched house and since the house h has a maximum capacity $c(h)$, it can only get matched to at most $c(h)$ agents. □

▷ *Property 2:* Let M and M' be two different popular matchings in G and let G_M and $G_{M'}$ denote the switching graphs respectively. For any vertex house h, the number of -1 outgoing edges from h is invariant across G_M and $G_{M'}$. The number of $+1$ incoming edges on h is also invariant across G_M and $G_{M'}$.
Proof. From Lemma 7, in any popular matching, each f-house h is matched to exactly $min\{|f(h)|, c(h)\}$ agents and this is also the number of outgoing edges with weight -1. A similar argument can be made for $+1$ weighted incoming edges. □

▷ *Property 3:* No $+1$ weighted edge can end at an unsaturated vertex.
Proof. If a $+1$ weighted edge is incident on a vertex h, this means that the house h is an f-house for some agent a that is still not matched to it in M. But if h is unsaturated then it still has some unused capacity. The matching M' obtained by just promoting a to h is popular than M, which is a contradiction. □

▷ *Property 4:* There can be no incoming -1 weighted edge on a saturated vertex if all its outgoing edges have weight -1.
Proof. A -1 weighted edge on a vertex h implies that the house h is an s-house for some agent a. But if h is saturated with all outgoing edges having a weight of -1, then all the capacity of h has been used up by agents who had h as their first choice. But by definition, h can not be an s-house for any other agent. □

▷ *Property 5:* For a given vertex h, if there exists at least one $+1$ weighted incoming edge, then all outgoing edges are of weight -1 and there can be no -1 weighted incoming edge on h.
Proof. Let agent a correspond to any $+1$ weighted incoming edge. Suppose h has an outgoing $+1$ edge ending at a vertex h' and agent a' corresponds to this edge. We can promote agents a and a' to their first choices and demote any agent which is assigned house h'. This leads to a matching more popular than M. Hence all outgoing edges from h must be of weight -1. Further,

Property 3 and Property 4 together imply that there can be no incoming edge on h of weight -1. □

Switching Moves We now describe the operation on the switching graph which takes us from one popular matching to another. We make use of the following terminology with reference to the switching graph G_M. Note that the term "path" ("cycle") implies a "directed path" ("directed cycle"). A "+1 edge"("−1 edge") means an "edge with weight +1" ("edge with weight −1").

- A path is called an *alternating path* if it starts with a +1 edge, ends at a −1 edge and alternates between +1 and −1 edges.
- A *switching path* is an alternating path that ends at an unsaturated vertex.
- A *switching cycle* is an even length cycle of alternating −1 and +1 weighted edges.
- A *switching set* is a union of edge-disjoint switching cycles and switching paths, such that at most k switching paths end a vertex of unsaturation degree k.
- A *switching move* is an operation on G_M by a switching set S in which, for every edge e in S, we reversed the direction of e and flip the weight of e ($+1 \leftrightarrow -1$).

Observe that every *valid* switching graph inherently implies a matching (in the context of CHA) of G.

Let $G_M = (\mathcal{H}, E_M)$ and $G_{M'} = (\mathcal{H}, E_{M'})$ be the switching graphs associated with popular matchings M and M' of the CHA instance $G = (\mathcal{A} \cup \mathcal{H}, E)$. Observe that the underlying undirected graph of G_M and $G_{M'}$ are same. We have the following.

Theorem 2. *Let S be the set of edges in G_M that get reversed in $G_{M'}$. Then, S is a switching set for G_M.*

We prove this algorithmically in stages.

Lemma 9. *Every directed cycle in S is a switching cycle of G_M.*

Proof. Let C be any cycle in S. From Property 5 of switching graphs, we know that no vertex in C can have an incoming edge and an outgoing edge of same weight +1. Similarly, since S is the set of edges in G_M which have opposite directions and opposite weights in $G_{M'}$, we observe that S can not contain any vertex with incoming and outgoing edges both having weight −1 (again from Property 5). This forces the weights of cycle C to alternate between +1 and −1. Moreover, this alternation forces the cycle to be of even length.

At this stage we apply the following algorithm to the set S.

`Reduction`(S):
1. `while` (*there exists a switching cycle C in S*):

let $S := S \setminus C$

2. `while` (S *is non-empty*):

(a) *find a longest path* P *in* S *which alternates between weights* $+1$ *and* -1

(b) *let* $S := S \setminus P$

At the end of every iteration of the `while` loop in Step 1, Lemma 9 still holds true. We now prove a very crucial invariant of the while loop in Step 2.

Lemma 10. *In every iteration of the* `while` *loop in Step 2 of the algorithm* `Reduction`*, the longest path in step* $2(a)$ *is a switching path for* G_M*.*

Proof. Let us denote the stages of the run of algorithm `Reduction` by t. Initially, at $t = 0$, before any of the while loops run, S is exactly the difference of edges in E_M and E'_M. Let the while loop in Step 1 runs t_1 times and the while loop in Step 2 runs t_2 times.

Let the current stage be $t = t_1 + i$. Let P be the maximal path in step $2(a)$ at this stage. We show that P starts with an edge of weight $+1$. For contradiction, let (h_i, h_j) be an edge of weight -1 and that this is the first edge of path P. Let a_{ij} be the agent associated with the edge (h_i, h_j).

The Property 5 of switching sets precludes any incoming edge of weight -1 on the vertex h_i. Hence, no switching path could have ended at h_i at any stage $t < t_1 + i$. Similarly, no switching cycle with an incoming edge -1 was incident on h_i at an earlier stage.

Let us assume that there were r cycles that were incident at h_i at $t = 0$. At stage $t = t_1 + i$, let the number of outgoing -1 edges be m. Hence at $t = 0$, h_i had r incoming $+1$ edges and $r + m$ outgoing -1 edges. But this would also imply that at $t = 0$, h_i had $r + m$ incoming $+1$ edges in $G_{M'}$. This contradicts Property 2, requiring the number of incoming $+1$ edges to be constant in the switching graphs corresponding to different popular matchings.

A similar argument can be made for the fact that the path P can only end at an edge with weight -1 and that P ends at an unsaturated vertex.

The following theorem establishes the characterization for popular matchings in CHA.

Theorem 3. *If* G_M *is the switching graph of the CHA instance* G *with respect to a popular matching* M*, then*
(i) every switching move on G_M *generates another popular matching, and*
(ii) every popular matching of G *can be generated by a switching move on* M*.*

Proof.

(i) We verify that the new matching generated by applying a switching move on G_M satisfies the characterization in Lemma 8. Call the new switching graph $G_{M'}$ and the associated matching M'. First, observe that M' is indeed an agent complete matching since $G_{M'}$ still has a directed edge for each agent in $\mathcal{A}$. Next, each agent a is still matched to $f(a)$ or $s(a)$ as the switching move either reverses an edge of G_M or leaves it as it is. Finally, for each

house h, $f(h) \subseteq M'(h)$ if $|f(h)| < c(h)$ and $|M'(h)| = c(h)$ with $M'(h) \subseteq f(h)$ otherwise. This is true because $|M'(h)| = |M(h)|$, by the definition of switching moves.

(ii) This is implied by Theorem 2.

3.2 Hardness of Counting

In this section we prove the #P-hardness of counting popular matchings in CHA. We reduce the problem of counting the number of matchings in a bipartite graph to our problem.

Let $G = (A \cup B, E)$ be a bipartite matching instance in which we want to count the number of matchings. From G we create a CHA instance I such that the number of popular matchings of I is same as the number of matchings of G.

Observe that a description of a switching graph gives the following information about its instance:

- the set of agents A,
- for each agent $a \in \mathcal{A}$, it gives $f(a)$ and $s(a)$, and
- for each s-house or f-house h, the unsaturation degree gives the capacity $c(h)$.

Using this information, we can create the description of the instance I so that it meets our requirement. For simplicity, we assume G to be connected (as isolated vertices do not affect the count). We orient all the edges of G from A to B and call the directed graph $G' = (A \cup B, E')$. Using G', we construct a graph S, which will be the switching graph.

Let $|A| = n_1$, $|B| = n_2$ and $|E'| = m$. S is constructed by augmenting G'. We keep all the vertices and edges of G' in S and assign each edge a weight of -1. Further, for each vertex $u \in A$, add a copy u' and add a directed edge from u' to u, and assign a weight of $+1$ to the edge. Call the new set of vertices A'. The sets A' and B contain s-houses and the set A contains f-houses. We label every vertex in A' and A as *saturated* and for each vertex v in B, we label v as *unsaturated* with *unsaturation degree* 1. Hence, the switching graph S has $2n_1 + n_2$ vertices and $n_1 + m$ edges.

The CHA instance I corresponding to the switching graph S has $2n_1 + n_2$ houses and $n_1 + m$ agents. Each agent has a preference list of length 2 that is naturally defined by the weight of edges in S.

Let the popular matching represented by S be M_ϕ. This corresponds to the empty matching of G. Every non-empty matching of G can be obtained by a switching move on S. We make this more explicit in the following theorem.

Theorem 4. *The number of matchings in G is same as the number of popular matchings in I.*

Proof. We prove this by showing that each matching in G corresponds to a unique set of edge disjoint switching paths in the switching graph S of I.

Consider a matching M of G and let $(u, v) \in M$. We look at the length 2 directed path in S that is obtained by extending (u, v) in the reverse direction: $u' \to u \to v$ with $u' \in A'$. It's easy to see that this is a switching path for I. Moreover, the set of switching paths obtained from any matching of G forms a valid switching set (as every pair of such paths arising from a matching are always edge disjoint).

For the converse, observe that S can only have switching paths of length 2 and it has no switching cycles. An edge disjoint set of such paths corresponds to a matching of G. By the definition of S, it's easy to see every matching in M can be obtained by a switching set of S.

Conclusions and Acknowledgement: We obtained an FPRAS for the #P-hard problem of counting popular matchings where instances could have ties. We presented a switching graph characterization for Capacitated House Allocation problem. Though our motivation for studying this structure was to prove a hardness result for counting popular matchings in CHA, the characterization may itself be of importance to many other problems of interest. This also completes the picture of House Allocation problems to a wider extent as such characterizations were only known for HA and HAT instances. We believe that this structure could be used to give an FPRAS for the case of CHA. This remains an open question.

We thank Meghana Nasre for fruitful discussions. We also thank anonymous reviewers for their input.

References

1. David J. Abraham, Robert W. Irving, Telikepalli Kavitha, and Kurt Mehlhorn. Popular matchings. *SIAM J. Comput.*, 37(4):1030–1045, 2007.
2. Nayantara Bhatnagar, Sam Greenberg, and Dana Randall. Sampling stable marriages: why spouse-swapping won't work. In *SODA*, pages 1223–1232, 2008.
3. Prasad Chebolu, Leslie Ann Goldberg, and Russell A. Martin. The complexity of approximately counting stable matchings. In *APPROX-RANDOM*, pages 81–94, 2010.
4. David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
5. Peter Gärdenfors. Match making: assignments based on bilateral preferences. *Behavioral Science*, 20(3):166–173, 1975.
6. Robert W. Irving and Paul Leather. The complexity of counting stable marriages. *SIAM J. Comput.*, 15(3):655–667, 1986.
7. Mark Jerrum, Alistair Sinclair, and Eric Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries. In *STOC*, pages 712–721, 2001.
8. Telikepalli Kavitha, Julián Mestre, and Meghana Nasre. Popular mixed matchings. *Theor. Comput. Sci.*, 412(24):2679–2690, 2011.

9. L. Lovász and M. D. Plummer. *Matching theory*, volume 121 of *North-Holland Mathematics Studies*. North-Holland Publishing Co., Amsterdam, 1986. Annals of Discrete Mathematics, 29.
10. Mohammad Mahdian. Random popular matchings. In *ACM Conference on Electronic Commerce*, pages 238–242, 2006.
11. Richard Matthew McCutchen. The least-unpopularity-factor and least-unpopularity-margin criteria for matching problems with one-sided preferences. In *LATIN*, pages 593–604, 2008.
12. Eric McDermid and Robert W. Irving. Popular matchings: structure and algorithms. *J. Comb. Optim.*, 22(3):339–358, 2011.
13. Meghana Nasre. Popular matchings: Structure and cheating strategies. In *STACS*, pages 412–423, 2013.
14. Colin T. S. Sng and David Manlove. Popular matchings in the weighted capacitated house allocation problem. *J. Discrete Algorithms*, 8(2):102–116, 2010.