A Simply Exponential Upper Bound on the Maximum Number of Stable Matchings

Anna R. Karlin
University of Washington
Seattle, WA, USA
karlin@cs.washington.edu

Shayan Oveis Gharan
University of Washington
Seattle, WA, USA
shayan@cs.washington.edu

Robbie Weber
University of Washington
Seattle, WA, USA
rtweber2@cs.washington.edu

ABSTRACT

Stable matching is a classical combinatorial problem that has been the subject of intense theoretical and empirical study since its introduction in 1962 in a seminal paper by Gale and Shapley. In this paper, we provide a new upper bound on $f(n)$, the maximum number of stable matchings that a stable matching instance with n men and n women can have. It has been a long-standing open problem to understand the asymptotic behavior of $f(n)$ as $n \rightarrow \infty$, first posed by Donald Knuth in the 1970s. Until now the best lower bound was approximately 2.28^n , and the best upper bound was $2^{n \log n - O(n)}$. In this paper, we show that for all n , $f(n) \leq c^n$ for some universal constant c . This matches the lower bound up to the base of the exponent. Our proof is based on a reduction to counting the number of downsets of a family of posets that we call “mixing”. The latter might be of independent interest.

CCS CONCEPTS

• **Mathematics of computing** → **Combinatorics**;

KEYWORDS

stable matching, stable marriage problem, counting, downset, poset

ACM Reference Format:

Anna R. Karlin, Shayan Oveis Gharan, and Robbie Weber. 2018. A Simply Exponential Upper Bound on the Maximum Number of Stable Matchings. In *Proceedings of 50th Annual ACM SIGACT Symposium on the Theory of Computing (STOC’18)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3188745.3188848>

1 INTRODUCTION

Stable matching is a classical combinatorial problem that has been the subject of intense theoretical and empirical study since its introduction in a seminal paper by Gale and Shapley in 1962 [8]. Variants of the algorithm introduced in [8] are widely used in practice, e.g. to match medical residents to hospitals. Stable matching is even the focus of the 2012 Nobel Prize in Economics [23].

A stable matching instance with n men and n women is defined by a set of preference lists, one per person. Person i ’s preference list gives a ranking over members of the opposite sex. The *Stable*

Matching Problem is to find a matching (i.e., a bijection) between the men and the women that is *stable*, that is, has no *blocking pairs*. A man m and a woman w form a blocking pair in a matching if they are not matched to each other, but both prefer the other to their partner in the matching. Gale and Shapley [8] showed that a stable matching always exists and gave an efficient algorithm to find one.¹ Since at least one stable matching always exists, a natural question is to determine the maximum number of stable matchings an instance of a given size can have. This problem was posed in the 1970s in a monograph by Knuth [13], and was the first of Gusfield and Irving’s twelve open problems in their 1989 textbook [10]. We denote the maximum number of stable matchings an instance with n men and n women can have by $f(n)$.

Progress on determining the asymptotics of $f(n)$ has been somewhat slow. The best lower bound is approximately 2.28^n , and the best upper bound prior to this paper was $2^{n \log n - O(n)}$. See the related work section for a detailed history.

In this paper, we present an improved upper bound.

THEOREM 1.1. *There is a universal constant c such that $f(n)$, the number of stable matchings in an instance with n men and n women, is at most c^n .*

To prove this theorem, we use a result of Irving and Leather [12] that shows that there is a bijection between the stable matchings of an instance $\mathcal{I}$ and the downsets² of a particular partially-ordered set (poset) associated with $\mathcal{I}$ known as the *rotation poset*. We show that the rotation poset associated with a stable matching instance has a particular property that we call *n-mixing*, and that any poset with this property has at most c^n downsets. All the steps in our proof are elementary.

The bound extends trivially to stable roommates instances. In the stable roommates problem, a set of n agents rank the other $n - 1$ agents in the set. The agents are paired off into roommate pairs, which are stable if no two agents would like to leave their partners and be matched to each other. A construction of Dean and Munshi [5], demonstrates that a stable roommates instance with n agents can be converted into a stable matching instance with n men and n women, such that the stable roommate assignments correspond to a subset of the stable matchings in the new instance. Using this construction, we can apply our upper bound to Stable Roommates.

THEOREM 1.2. *There is a universal constant c , such that the number of stable assignments in a stable roommate instance with n agents is at most c^n .*

¹Stable matching algorithms were actually developed and used as early as 1951 to match interns to hospitals [26].

²See section 2 for definitions of all the relevant terminology.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.
STOC’18, June 25–29, 2018, Los Angeles, CA, USA

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5559-9/18/06...\$15.00
<https://doi.org/10.1145/3188745.3188848>

1.1 Related Work

Lower Bounds. It is trivial to provide instances with $2^{n/2}$ stable matchings by combining disjoint instances of size 2. Irving and Leather constructed a family of instances [12] which has since been shown by Knuth³ to contain at least $\Omega(2.28^n)$ matchings. Irving and Leather's family only has instances for n which is a power of 2. Benjamin, Converse, and Krieger also provided a lower bound on $f(n)$ by creating a family of instances with $\Omega(2^n \sqrt{n})$ matchings [2]. While this is fewer matchings than the instances in [12], Benjamin et al.'s family has instances for every even n , not just powers of 2. In 2002, Thurber extended Irving and Leather's lower bound to all values of n . For n powers of 2, Thurber's construction exactly coincides with Irving and Leather's. For all other n , the construction produces a lower bound of $2.28^n / c^{\log n}$ for some constant c [28]. To date, this lower bound of $\Omega(2.28^n)$ is the best known. We refer the reader to Manlove's textbook for a more thorough description of the history of these lower bounds [18].

Upper Bounds. Trivially, there are at most $n!$ stable matchings (as there are at most $n!$ bijections between the men and women). The first progress on upper bounds that we are aware of was made by Stathoupolos in his 2011 Master's thesis [27], where he proves that the number of stable matchings is at most $O(n! / c^n)$ for some constant c . A more recent paper of Drgas-Burchardt and Świtalski shows a weaker upper bound of approximately $\frac{3}{4}n!$ [6]. All previous upper bounds have the form $2^{n \log n - O(n)}$.

Restricted Preferences. The number of possible stable matchings has also been studied under various models restricting or randomizing the allowable preference lists. If all preference lists are equally likely and selected independently for each agent, Pittel shows that the expected number of stable matchings is $O(n \log n)$ [19]. Applying Markov's Inequality shows that the number of stable matchings is polynomial in n with probability $1 - o(1)$. Therefore, the lower bound instances described above are a vanishingly small fraction of all instances. Work of Hoffman, Levy, and Mossel (described in Levy's PhD thesis [16]) shows that under a Mallows model [17], where preference lists are selected with probability proportional to the number of inversions in the list, the number of stable matchings is C^n with high probability (where the constant C depends on the exact parameters of the model).

The number of attainable partners⁴ a person can have has also been the subject of much research. Knuth, Motwani, and Pittel show that the number of attainable partners is $O(\log n)$ with high probability if the lists are uniformly random [15]. Immorlica and Mahdian show that if agents on one side of the instance have random preference lists of length k (and consider all other agents unacceptable) the expected number of agents with more than one attainable partner depends only on k (and not on n) [11]. Ashlagi, Kanoria, and Leshno show that if the number of men and women is unbalanced, with uniformly random lists, the fraction of agents with more than one attainable partner goes to 0 as the size of the market grows [1]. Intuitively, the advantage to being on the proposing side "disappears" when the market becomes unbalanced (and sufficiently

large). Follow-up work of Pittel finds the expected number of stable matchings in this setting and related results [20].

Various related questions have been examined for alternative notions of stability and in the stable roommates setting. See [21, 22] and references therein.

Counting. A natural computational problem is to count the number of stable matchings in a given instance as efficiently as possible. Irving and Leather show that finding the exact number of matchings is #P-complete [12], so finding an approximate count is a more realistic goal. Bhatnagar, Greenberg, and Randall consider instances where preference lists come from restricted models [3]; for example, those in which the preference lists reflect linear combinations of k "attributes" of the other set, or where every agent appears in a "range" of k consecutive positions in every preference list. In both of these cases, they show that a natural Markov Chain Monte Carlo approach does not produce a good approximation (as the chain does not mix efficiently). As part of their proof, they show the number of stable matchings can still be as large as c^n for some constant $c < 2.28$, even in these restricted cases.

A formal hardness result was later shown by Dyer, Goldberg, Greenhill, and Jerrum [7]. They show approximately counting the number of stable matchings is equivalent to approximately counting for a class of problems, canonically represented by #BIS.⁵ This hardness result was strengthened by Chebolu, Goldberg, and Martin [4] to hold even if the instances come from some of the restricted classes of [3].

The heart of all of these results is the rotation poset (originally developed in [12]), which we use and describe in section 4.

Stable Matching in General. We refer the reader to books by Roth and Sotomayor [25], Gusfield and Irving [10], Manlove [18], and Knuth [14] for more about the topic of stable matching. For many examples of stable matching in the real world, see [24].

2 PRELIMINARIES AND MAIN TECHNICAL THEOREM

In this section, we review standard terminology regarding partially ordered sets, describe the key property of a poset we will use, and state our main technical theorem (Theorem 2.5).

Definition 2.1 (Poset). A partially ordered set (or poset) $(V, <)$, is defined by a set V and a binary relation, $<$, on V satisfying:

- **Antisymmetry:** For all distinct $u, v \in V$ if $u < v$ then $v \not< u$, and
- **Transitivity:** for all $u, v, w \in V$ if $u < v$ and $v < w$ then $u < w$.

Two elements $u, v \in V$ are *comparable* if $u < v$ or $v < u$. They are *incomparable* otherwise.

If $u < v$, we say u is *dominated* by v and v *dominates* u .

Definition 2.2 (Chain). A set S of elements is called a *chain* if each pair of elements in S is comparable. In other words, for $\ell > 0$, a chain of length ℓ is a sequence of elements $v_1 < v_2 < v_3 < \dots < v_\ell$.

³Personal communication, as described in [10].

⁴Woman w is an attainable partner of man m if there is a stable matching in which they are matched to each other.

⁵More specifically, they show an FPRAS for the number of stable matchings exists if and only if one exists for #BIS, approximately counting the number of independent sets in a bipartite graph. Goldberg and Jerrum conjecture that no such FPRAS exists [9]. See [4] for formal definitions.

Figure 1: A $2n$ -mixing poset with respect to chains defined by the red and blue paths. A path in the graph from v to u indicates that $u < v$. That is, the poset is the transitive closure of the arrows shown. For any set U of k elements, there are at least $2\sqrt{k}$ chains that contain one of these elements.

A set of elements is called an *antichain* if they are pairwise incomparable.

Definition 2.3 (Downset). A *downset* of a partial order is an antichain and all elements dominated by some element of that antichain.

Observe that a downset is closed under $<$. That is, for any downset S , if $v \in S$ and $u < v$ then $u \in S$.

The following is the key property of the posets associated with stable matching instances that we will use in the proof.

Definition 2.4 (n -mixing). A poset $(V, <)$ is n -mixing if there exist n chains $C_1, \dots, C_n$ (not necessarily disjoint) such that

- i) Every element of V belongs to at least one of the chains, i.e., $\bigcup_{i=1}^n C_i = V$,
- ii) For any $U \subseteq V$, at least $2\sqrt{|U|}$ of the C_i contain an element of U .

Observe that if a poset is formed by n disjoint chains, each of length ℓ , it has about ℓ^n downsets, and ℓ could be arbitrarily bigger than n . But such a poset is not mixing, since taking U to be the set of elements on one of the chains violates property ii) of mixing. For an example of a mixing poset, see Figure 1. We can now state our main technical theorem.

THEOREM 2.5. *There is a universal constant c , such that if a poset is n -mixing, then it has at most c^n downsets.*

Note that the n -mixing property immediately implies that the poset has at most n^2 elements; just let $U = V$ in the above definition. A poset with n^2 elements covered by n chains can have at most $(n+1)^n$ downsets (this is achieved for n equal length chains). So, the main contribution of the above theorem is to improve this trivial upper bound to c^n for some constant c .

Theorem 2.5 is the main technical contribution of this work. The proof is contained in section 3. To complete the proof of Theorem 1.1 we use the following theorem relating mixing posets to stable matchings.

THEOREM 2.6. *For every stable matching instance I with a total of n men and women, there exists an n -mixing poset $(\mathcal{R}, <)$, called the rotation poset, such that the number of downsets of the poset is equal to the number of stable matchings of I .*

Note that Theorem 2.5 and Theorem 2.6 immediately imply Theorem 1.1. We prove Theorem 2.6 in section 4, by combining existing observations about the rotation poset.

3 PROOF OF MAIN TECHNICAL THEOREM

In this section we prove Theorem 2.5. The proof proceeds by finding an element v of the poset which dominates and is dominated by many elements. We then count downsets by considering those downsets that contain v and those that do not. Since v dominates and is dominated by many elements, the size of each remaining instance is significantly smaller, yielding the bound.

Formally, we say an element is α -critical if it dominates α elements and is dominated by α elements. The key lemma is that there is always an $\Omega((|V|/n)^{3/2})$ -critical element.

Lemma 3.1. *Let $(V, <)$ be an n -mixing poset with respect to chains $C_1, \dots, C_n$, and define $d = \frac{|V|}{n}$. For some universal constants $d_0 > 1$ and $c_0 > 0$, there is an element $v \in V$ such that v is $(c_0 d^{3/2})$ -critical as long as $d \geq d_0$.*

We prove Lemma 3.1 in subsection 3.1 via a counting argument.

In the rest of this section we prove Theorem 2.5 using Lemma 3.1. We bound the number of downsets by induction on d .

Our base case is when $d = d_0$. In this case, the number of downsets is maximized when the chains are all the same length, so we have an upper bound of $(d_0 + 1)^n$.

For larger d , first we identify a $(c_0 d^{3/2})$ -critical element v . The number of downsets containing v is the number of downsets in the poset remaining after we delete v and everything it dominates. Similarly, the number of downsets not containing v is the number of downsets in the poset remaining after we delete v and everything dominating v . In both cases, the resulting poset is still n -mixing, so we can induct. We call such a step (choosing a critical element) an *iteration*.

It remains to bound the number of downsets that this process enumerates. By the n -mixing property, there are at most n^2 elements in the initial poset. We partition the iterations into phases, where in phase i we reduce the size of the poset from $\frac{n^2}{2^i}$ elements to $\frac{n^2}{2^{i+1}}$ elements. By definition, in phase i , $d \geq \frac{n}{2^{i+1}}$. So, we can bound the number of iterations required in phase i (call it k_i) by:

$$c_0 \left(\frac{n}{2^{i+1}} \right)^{3/2} k_i > \frac{n^2}{2^{i+1}}.$$

Rearranging, we see it suffices to choose $k_i = 2^{(i+1)/2} \sqrt{n}/c_0$. We continue until $d = d_0$. Summing across all phases, the number of choices to make is at most

$$\frac{\sqrt{n}}{c_0} \sum_{i=0}^{\log n} 2^{(i+1)/2} < 5 \frac{n}{c_0}.$$

So, the algorithm enumerates at most $(d_0 + 1)^n$ downsets in the base case and it makes at most $5n/c_0$ choices during the inductive process. Thus, the number of downsets is at most $(d_0 + 1)^n 2^{5n/c_0} = c^n$ as required.

3.1 Proof of Main Technical Lemma

Finally, we prove [Lemma 3.1](#), i.e. we show that any n -mixing poset $(V, <)$ contains a $c_0 d^{3/2}$ -critical element as long as $d \geq d_0$. (Recall that $d = |V|/n$.)

We make use of the standard graph representation of the partial order: In this graph there is a node for each element of the partial order, and a directed edge from v to u if $u < v$. Of course, this directed graph is acyclic. Henceforth, we refer only to this DAG rather than to the poset and partition the nodes of the DAG into levels as follows: Level 1 nodes are those with no outgoing edges, i.e., sinks of the DAG, and level i nodes are those whose longest path to a sink (i.e., a level 1 node) has exactly i nodes. Note that each level is an antichain.

Next, we create n disjoint subchains $S_1, \dots, S_n$. The subchain S_i will be a subset of the nodes in C_i . We perform the assignment of nodes to subchains by processing up the DAG level by level. Initialize every subchain S_i to be empty. For each node u , consider the set of indices $I(u) = \{j : u \in C_j\}$. Assign u to the subchain for an index in $I(u)$ which currently has the fewest nodes among those chains, i.e. $\arg \min_{j \in I(u)} |S_j|$, breaking ties arbitrarily. If u is the k^{th} node assigned to a subchain S_i , then we say the *height* of u is k . By construction, the following properties hold:

- (a) The S_i 's are chains, since $S_i \subseteq C_i$.
- (b) The S_i 's are disjoint.
- (c) If u has height h , then u dominates at least $h - 1$ nodes in each of the subchains S_j such that $j \in I(u)$ (i.e., $\{j : u \in C_j\}$).

Claim 3.2. *Let $\bar{D}$ be the set of nodes of height at least $\lfloor d/2 \rfloor$. Each node $u \in \bar{D}$ dominates at least $c_0 \cdot d^{3/2}$ nodes, for $d \geq d_0$, where c_0 and d_0 are universal constants.*

PROOF. Suppose that $u \in S_i \cap \bar{D}$ is at height $\ell \geq \lfloor d/2 \rfloor$ and let $D(u)$ be the nodes in S_i of height $\lfloor \ell/2 \rfloor$ through ℓ . By the mixing property, these nodes lie on at least $2\sqrt{\lfloor \ell/2 \rfloor}$ chains C . Moreover, by construction, on each subchain $S_j \in C$, at least $\lfloor \ell/2 \rfloor - 1$ nodes are dominated by some node in $D(u)$ and hence are dominated by u . Therefore, u dominates at least $(\lfloor \ell/2 \rfloor - 1) \cdot 2\sqrt{\lfloor \ell/2 \rfloor} + (\lfloor \ell/2 \rfloor - 1) = \Omega(d^{3/2})$ nodes. $\square$

Since the number of nodes in the DAG is dn , we conclude:

Corollary 3.3. *There is a set $\bar{D}$ of strictly more than $|V|/2$ nodes, that each dominate $\Omega(d^{3/2})$ nodes.*

A symmetric argument in which subchains are built starting from the sources of the DAG shows that there is a set $\underline{D}$ of strictly more than $|V|/2$ nodes that are dominated by $\Omega(d^{3/2})$ nodes. Therefore, there is some node v in the intersection of $\bar{D}$ and $\underline{D}$. This is the $c_0 d^{3/2}$ -critical node we seek.

Remark 3.4. *A crude analysis shows that $c_0 \geq 1/8$ when $d_0 > 25$.*

4 ROTATIONS AND THE ROTATION POSET

In this section we present a key theorem of Irving and Leather [12], and use it to prove [Theorem 2.6](#).

THEOREM 2.6. *For every stable matching instance $\mathcal{I}$ with a total of n men and women, there exists an n -mixing poset $(\mathcal{R}, <)$, called the rotation poset, such that the number of downsets of the poset is equal to the number of stable matchings of $\mathcal{I}$.*

We begin with the definitions needed to prove [Theorem 2.6](#).

Definition 4.1 (Rotation). *Let $k \geq 2$. A rotation ρ is an ordered list of pairs*

$$\rho = ((m_0, w_0), (m_1, w_1), \dots, (m_{k-1}, w_{k-1}))$$

that are matched in some stable matching M with the property that for every i such that $0 \leq i \leq k-1$, woman w_{i+1} (where the subscript is taken mod k) is the highest ranked woman on m_i 's preference list satisfying:

- i) *man m_i prefers w_i to w_{i+1} , and*
- ii) *woman w_{i+1} prefers m_i to m_{i+1} .*

In this case, we say ρ is exposed in M .

We will sometimes abuse notation and think of a rotation as the set containing those pairs. Also, we will need the following facts about rotations later.

Lemma 4.2 ([12, Lemma 4.7]). *A pair (m, w) can appear in at most one rotation.*

Lemma 4.3 ([10, Lemma 2.5.1]). *If ρ is a rotation with consecutive pairs (m_i, w_i) , and (m_{i+1}, w_{i+1}) , and w is a woman between w_i and w_{i+1} in m_i 's preference list, then there is no stable matching containing the pair (m_i, w) .*

Definition 4.4 (Elimination of a Rotation). *Let*

$\rho = ((m_0, w_0), \dots, (m_{k-1}, w_{k-1}))$ *be a rotation exposed in stable matching M . The rotation ρ is eliminated from M by matching m_i to $w_{(i+1) \bmod k}$, for all $0 \leq i \leq k-1$, leaving all other pairs in M unchanged, i.e., matching M is replaced with matching M' , where*

$$M' := M \setminus \rho \cup \{(m_0, w_1), (m_1, w_2), \dots, (m_{k-1}, w_0)\}.$$

*Note that when we eliminate a rotation from M , the resulting matching M' is stable.*⁶

Irving and Leather studied the following process: Fix a stable matching instance $\mathcal{I}$. Starting at the man-optimal matching⁷, choose a rotation in the current matching, and eliminate it. They show that for any stable matching M , there is a set of rotations, $R(M)$, one can eliminate (starting from the man-optimal matching) that will yield M .

⁶Switching from M to M' makes all the women in ρ happier and all the men in ρ less happy. It is easy to check that this switch cannot create a blocking pair inside the set ρ . The only other possibility for a blocking pair is a man in ρ with a woman outside ρ . For $(m_i \in \rho, w \notin \rho)$ to become a blocking pair, m_i would have to prefer w to w_{i+1} , but by the definition of rotation, w_{i+1} was the first woman on m_i 's list who would prefer to be matched to him, so he cannot prefer w to w_{i+1} . See also [10, Lemma 2.5.2].

⁷A fascinating fact about stable matching is that there is a matching, known as the man-optimal matching, in which each man is matched with his favorite attainable partner. Recall that woman w is attainable for man m if there is some stable matching in which they are matched.

Men's preferences:

m_1	w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
m_2	w_2	w_1	w_4	w_3	w_6	w_5	w_8	w_7
m_3	w_3	w_4	w_1	w_2	w_7	w_8	w_5	w_6
m_4	w_4	w_3	w_2	w_1	w_8	w_7	w_6	w_5
m_5	w_5	w_6	w_7	w_8	w_1	w_2	w_3	w_4
m_6	w_6	w_5	w_8	w_7	w_2	w_1	w_4	w_3
m_7	w_7	w_8	w_5	w_6	w_3	w_4	w_1	w_2
m_8	w_8	w_7	w_6	w_5	w_4	w_3	w_2	w_1

Women's preferences:

w_1	m_8	m_7	m_6	m_5	m_4	m_3	m_2	m_1
w_2	m_7	m_8	m_5	m_6	m_3	m_4	m_1	m_2
w_3	m_6	m_5	m_8	m_7	m_2	m_1	m_4	m_3
w_4	m_5	m_6	m_7	m_8	m_1	m_2	m_3	m_4
w_5	m_4	m_3	m_2	m_1	m_8	m_7	m_6	m_5
w_6	m_3	m_4	m_1	m_2	m_7	m_8	m_5	m_6
w_7	m_2	m_1	m_4	m_3	m_6	m_5	m_8	m_7
w_8	m_1	m_2	m_3	m_4	m_5	m_6	m_7	m_8

Figure 2: A size-8 stable matching instance with its rotation poset. This is part of the family of instances described by Irving and Leather, which produces the $\Omega(2.28^n)$ lower bound.

However, there is a partial order on the set of rotations – some must be eliminated before others.⁸ If ρ must be eliminated before ρ' , we write $\rho < \rho'$. Let $\mathcal{R}$ be the set of rotations for a stable matching instance, and let $<$ be that partial order on the rotations defined by elimination order. We call this poset the *rotation poset*. See Figure 2 for an example of a stable matching instance and the corresponding rotation poset.

For our purposes, the important result relating the rotation poset to stable matchings is the following.

THEOREM 4.5 ([12, THEOREM 4.1]). *For any stable matching instance, there is a one-to-one correspondence between downsets of the rotation poset and the set of stable matchings. In other words, the number of stable matchings is exactly equal to the number of downsets of $(\mathcal{R}, <)$.*

Indeed, the downset corresponding to M is exactly the set $R(M)$ discussed above.

Thus, to prove Theorem 2.6, it remains to show that the rotation poset associated to any stable matching instance with a total of n men and women is n -mixing. First we construct the chains $C_1, \dots, C_n$. We are going to have one chain for each agent (man or woman), where the corresponding chain contains all the rotations that include that agent. Call these sets $C_1, \dots, C_n$. To prove that $C_1, \dots, C_n$ is n -mixing, we first need to show that each C_i is indeed a chain, i.e., every pair of rotations where a specific agent appears are comparable and second we need to show that $C_1, \dots, C_n$ satisfy property (ii) of Definition 2.4.

Claim 4.6. *If two rotations share an agent, then they are comparable.*⁹

PROOF. First, suppose that the shared agent is a man. If it is a woman, we can just switch the designations of “men” and “women” and use the same proof on the “reversed” version of the rotation graph. Let ρ_1, ρ_2 be rotations sharing an agent m , and let m be matched to w_1 in ρ_1 and w_2 in ρ_2 , where m prefers w_1 to w_2 .

For the sake of contradiction assume that ρ_1 and ρ_2 are incomparable. We show that there exists a rotation ρ which causes m to skip over w_1 . This would contradict Lemma 4.3 as it implies that (m, w_1) belongs to no stable matching.

Suppose we start from the man-optimal stable matching and eliminate all rotations dominated by ρ_2 and let M be the resulting stable matching. By the correspondence in Theorem 4.5, $(m, w_2) \in M$. Since m prefers w_1 to w_2 , there must be a rotation ρ that we eliminated which caused m to be matched to someone worse than w_1 for the first time. Since ρ_2 and ρ_1 are incomparable, $\rho_1 \neq \rho$. Therefore, by Lemma 4.2 $(m, w_1) \notin \rho$; so, ρ caused m to skip over w_1 . This is a contradiction. $\square$

Claim 4.7. *Every set of k rotations contains at least $2\sqrt{k}$ agents.*

⁸For example, a rotation containing a pair (m, w) is not exposed (and thus cannot be eliminated) until that pair is matched, so a rotation with consecutive pairs (m, w') , (m', w) must be eliminated first. The details of exactly when one rotation must be eliminated before another are not of direct use to us (we only require the rather coarse description in Claim 4.6), so we do not describe them here. See [12] or [10] for a full description of the poset.

⁹This observation is not novel; for example, it is implicit in the discussion of [10], but we have not seen the statement explicitly written down, so we prove it here.

PROOF. We argue by contrapositive. Suppose we have a set of rotations involving fewer than $2\sqrt{k}$ agents. Every rotation contains a (man, woman) pair, who (by Lemma 4.2) have not appeared together before. With fewer than $2\sqrt{k}$ agents, there are strictly less than k (man, woman) pairs which can appear, and thus fewer than k rotations in the set. $\square$

5 CONCLUSION

We have shown there is some constant c such that $f(n) \leq c^n$. We have not made a significant effort to optimize the constants in our argument, favoring ease of exposition over the exact result. By making a few minor changes to the argument, we obtain $f(n) \leq 2^{17n}$ for sufficiently large n . A more careful argument could probably improve this constant somewhat, but this approach will not get a constant c close to the (approximately) 2.28 we would need to match the best known lower bound. Determining the precise asymptotic behavior of $f(n)$ remains an interesting open problem.

ACKNOWLEDGMENTS

The first author gratefully acknowledges the support of NSF grant CCF 1420381. The second author gratefully acknowledges the support of NSF grant CCF-1552097 and ONR-YIP grant N00014-17-1-2429.

REFERENCES

- [1] Itai Ashlagi, Yash Kanoria, and Jacob D Leshno. 2017. Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy* 125, 1 (2017), 69–98.
- [2] Arthur T Benjamin, Cherlyn Converse, and Henry A Krieger. 1995. How do I marry thee? Let me count the ways. *Discrete Applied Mathematics* 59, 3 (1995), 285–292.
- [3] Nayantara Bhatnagar, Sam Greenberg, and Dana Randall. 2008. Sampling stable marriages: why spouse-swapping won't work. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1223–1232.
- [4] Prasad Chebolu, Leslie Ann Goldberg, and Russell Martin. 2012. The complexity of approximately counting stable matchings. *Theoretical Computer Science* 437 (2012), 35–68.
- [5] Brian C Dean and Siddharth Munshi. 2010. Faster algorithms for stable allocation problems. *Algorithmica* 58, 1 (2010), 59–81.
- [6] Ewa Drgas-Burchardt and Zbigniew Świtalski. 2013. A number of stable matchings in models of the Gale–Shapley type. *Discrete Applied Mathematics* 161, 18 (2013), 2932–2936.
- [7] Martin Dyer, Leslie Ann Goldberg, Catherine Greenhill, and Mark Jerrum. 2004. The relative complexity of approximate counting problems. *Algorithmica* 38, 3 (2004), 471–500.
- [8] David Gale and Lloyd S Shapley. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.
- [9] Leslie Ann Goldberg and Mark Jerrum. 2012. Approximating the partition function of the ferromagnetic Potts model. *Journal of the ACM (JACM)* 59, 5 (2012), 25.
- [10] Dan Gusfield and Robert W Irving. 1989. *The stable marriage problem: structure and algorithms*. MIT press.
- [11] Nicole Immorlica and Mohammad Mahdian. 2005. Marriage, honesty, and stability. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 53–62.
- [12] Robert W Irving and Paul Leather. 1986. The complexity of counting stable marriages. *SIAM J. Comput.* 15, 3 (1986), 655–667.
- [13] Donald Ervin Knuth. 1976. *Mariages stables et leurs relations avec d'autres problèmes combinatoires: introduction à l'analyse mathématique des algorithmes*. Montréal: Presses de l'Université de Montréal.
- [14] Donald Ervin Knuth. 1997. *Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms*. Vol. 10. American Mathematical Soc. English translation of *Mariages stables*.
- [15] Donald E Knuth, Rajeev Motwani, and Boris Pittel. 1990. Stable husbands. *Random Structures & Algorithms* 1, 1 (1990), 1–14.
- [16] Avi Levi. 2017. *Novel uses of the Mallows model in coloring and matching*. Ph.D. Dissertation. University of Washington.
- [17] Colin L Mallows. 1957. Non-null ranking models. I. *Biometrika* 44, 1/2 (1957), 114–130.
- [18] David F Manlove. 2013. *Algorithmics of matching under preferences*. Vol. 2. World Scientific.
- [19] Boris Pittel. 1989. The average number of stable matchings. *SIAM Journal on Discrete Mathematics* 2, 4 (1989), 530–549.
- [20] Boris Pittel. 2017. On likely solutions of the stable matching problem with unequal numbers of men and women. *arXiv preprint arXiv:1701.08900* (2017).
- [21] Boris Pittel. 2017. On random exchange-stable matchings. *arXiv preprint arXiv:1707.01540* (2017).
- [22] Boris Pittel. 2017. On random stable partitions. *arXiv preprint arXiv:1705.08340* (2017).
- [23] Catherine Rampell. October 15, 2012. 2 From U.S. Win Nobel in Economics. *The New York Times* (October 15, 2012). <http://www.nytimes.com/2012/10/16/business/economy/alvin-roth-and-lloyd-shapley-win-nobel-in-economic-science.html>
- [24] Alvin E Roth. 2015. *Who Gets What—and Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt.
- [25] Alvin E Roth and Marilda A Oliveira Sotomayor. 1992. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press.
- [26] John M Stalnaker. 1953. The matching program for intern placement. *Academic Medicine* 28, 11 (1953), 13–19.
- [27] Georgios K. Stathopoulos. 2011. *Variants of Stable Marriage Algorithms, Complexity and Structural Properties*. Master's thesis. University of Athens, Department of Mathematics.
- [28] Edward G Thurber. 2002. Concerning the maximum number of stable matchings in the stable marriage problem. *Discrete Mathematics* 248, 1-3 (2002), 195–219.