Sampling Stable Marriages: Why Spouse-Swapping Won't Work*

Nayantara Bhatnagar[†] Sam Greenberg[‡] Dana Randall[§]

Abstract

We study the behavior of random walks along the edges of the stable marriage lattice for various restricted families of allowable preference sets. In the "k-attribute model," each man is valued in each of k attributes, and each woman's ranking of the men is determined by a linear function, representing her relative ranking of those attributes; men's rankings of the women are determined similarly. We show that sampling with a random walk on the marriage lattice can take exponential time, even when $k = 2$. Moreover, we show that the marriage lattices arising in the k-attribute model are more restrictive than in the general setting; previously such a restriction had only been shown for the sets of preference lists. The second model we consider is the "k-range model," where each person lies in a position in $[i, i + k - 1]$, for some i, on every preference list of the opposite sex. When $k = 1$ there is a unique stable marriage. When $k = 2$ there already can be an exponential number of stable marriages, but we show that a random walk on the stable marriage lattice always converges quickly to equilibrium. However, when $k \geq 5$, there are preference sets such that the random walk on the lattice will require exponential time to converge. Lastly, we show that in the extreme case where each gender's rankings of the other are restricted to one of just a constant k possible preference lists, there are still instances for which the Markov chain mixes exponentially slowly, even when $k = 4$. This oversimplification of the general model helps elucidate why Markov chains based on spouse-swapping are not good approaches to sampling, even in specialized scenarios.

1 Introduction

In the simplest formulation of the Stable Marriage Problem [5, 7, 11], there are two disjoint sets of n men and n women. Each man has a ranking of the n women, called a *preference list*, and each woman has an analogous ranking of the men. We call these collections of preference lists the *preference sets*. A marriage is a perfect matching, pairing men and women. It is considered *stable* if no man-woman pair simultaneously ranks each other higher than their respective spouses. Gale and Shapley [5] showed the remarkable result that every instance of preference lists has a stable marriage. Since then, many variants of the problem have been studied such as the non-bipartite version (the *stable roommates* problem) or case of more than two disjoint sets of people (the so-called *man-woman-dog* marriage problem). In addition, the stable marriage problem is studied in various other scenarios including colleges and applicants as well as hospitals and residents [5].

It is well-known that the set of stable marriages form a distributive lattice, known as the *stable marriage lattice*. This is a graph where vertices correspond bijectively to the stable marriages and a marriage is above another if every man is at least as happy with the first marriage as with the second [11, 7]. If one marriage is better for all men who have new partners, then it turns out to be worse for each of their partners. The lattice has unique male and female optimal marriages and the Gale-Shapley algorithm [5] efficiently produces these extremal stable marriages.

A common concern with the standard Gale-Shapley algorithm is that it unfairly favors one sex at the expense of the other. This gives rise to the problem of finding "fair" stable marriages. Previous work on finding fair marriages has focused on algorithms for optimizing an objective function that captures the "happiness" of either a group of people or the person that is worst off, without regard for the gender [6, 9, 11, 12, 13, 16]. This is in contrast to the Gale-Shapley algorithm, which is the special case maximizing the happiness of only one gender. Simple schemes such as flipping a coin to decide which group attains optimality, or walking along the edges of the lattice to find a stable marriage in the middle level, seem unsatisfactory; if the lattice has many stable marriages that are nearly female-optimal and very few that are nearly male-optimal, then it might be more "fair" to generate a marriage that favors the women.

The optimization approaches based on happiness

*Supported in part by NSF grants CCR-0515105 and DMS-0505505.

†Department of Statistics, University of California at Berkeley, Berkeley, CA 94720-3860.

‡School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332-0160.

§School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332-0765.

typically involve showing that an appropriately defined integer program can be efficiently solved. A different approach is to ask if one can generate a *random* stable marriage uniformly. The most natural method for random generation is to define a random walk along edges of the lattice so that it converges to the uniform distribution over all stable marriages. It known that the edges of the lattice correspond to "rotations," where a set of men swap partners so that either all of the men rank their new wives higher than before while the women rank their new husbands lower, or vice versa (see [7] or [11] for details). We can think of a move of the chain as forcing someone to accept a less favorable mate, setting off a series of "spouse swaps" whereby a set of people of one gender all improve in their preference lists while all of their previous partners do worse. The edges of the lattice dictate exactly how to formalize these moves so that there are at most a linear number of spouse-swapping moves at each step, these are easily identifiable, and they form reversible moves that connect the state space and converge to the uniform distribution. Unfortunately, it is possible to construct preferences sets such that the stable marriage lattice, when regarded as a Markov kernel for a random walk, has a "bad cut" in the state space, implying the walk will take exponential time to converge to equilibrium.

The existence of a bad cut follows from the realization that every distributive lattice is isomorphic to some marriage lattice [1, 8]. Since it is easy to construct examples of lattices on which the random walk will mix slowly, it is not surprising that this approach fails for the general problem of sampling stable marriages. In fact, the problem of counting the number of stable marriages can be reduced to the problem of counting the number of antichains of a partial order; for any partial order, there is an instance of the marriage problem (of size polynomial in the size of the partial order) whose rotation poset is isomorphic to the partial order and the number of stable marriages is in bijection with the number of antichains of its rotation poset. This puts the problem of counting stable marriages in a large class of problems studied in [4] which includes counting antichains of a poset and counting the number of independent sets in a bipartite graph, which are interreducible with respect to the existence of an FPRAS (fully polynomial randomized approximating scheme). The existence of an FPRAS for this class of problems is open. This apparent intractability motivates the consideration of the problem of sampling stable marriages for more restricted sets of instances that model situations that might arise in actual applications.

1.1 Specialized models

Each of the following models limits the allowable preference sets in different ways.

The k-attribute model: Our first model is motivated by on-line dating sites that ask many questions to compile extensive psychological profiles for each person. Some of these questions attempt to value a person in terms of some k attributes, such as athleticism, intelligence, salary, etc. Other questions try to determine the relative importance of these attributes in a potential mate.[1] Accordingly, we define the k-attribute model first by associating the men with points in R^k. Each woman's ranking of the men is defined by a linear function of these attributes, whereby her preference list is determined by projecting the men's points onto some line. Likewise, reversing the roles of men and women determines the men's preference lists. A similar model was considered previously where it was shown that not all preference sets can be realized in a k-attribute model when $k \leq n-2$ (see [2] and references therein).

The k-range model: In the setting where medical students are matched to hospital residencies, one might expect some amount of uniformity among the preference lists of each group. Of course this will not be true if students exhibit nonuniform preferences over geography or specialties, but we consider the extreme case where students are interested solely in hospitals' national ranking and hospitals are interested only in the students' academic performance. If the preference lists are completely uniform on each side, then there will be a unique stable marriage. Instead we weaken this slightly and consider rankings such that if one hospital ranks a student 25th, then she always falls in between positions, say, 20 and 30 on all the other hospitals' lists.

Formalizing this scenario, we say that a preference set is from a k-range model if, for each student, there exists an index i such that the student falls in position $i, i+1, ...,$ or $i+k-1$ on each hospital's list, and similarly each hospital falls in some range of width k on all of the students' lists. We say such an input has range k.

The k-list model: In our third and final model, we further simplify the allowable preference lists to drive home how pervasive obstacles to fast mixing can be. Each gender is partitioned into at most k different groups, and all members of a group are restricted to have exactly the same preference list. Such a scenario depicts a world where, for instance, all the women on the chess club have the same preference list. This may be different from the list shared by the women on

[1] We restrict to a heterosexual context throughout this paper. Moreover, any stereotypes and/or politically incorrect statements are intended solely for clarity of exposition.

the basketball team, or the one shared by the drama club. Note that we are not making any assumption that women on the chess club prefer male chess players to male athletes; rather they can have any ordering based on factors describing individual men.

1.2 Results We give a general construction demonstrating that walks on the stable marriage lattice can take exponential time to converge to equilibrium, and then show how to modify it to fit the restricted contexts. We show that the k-attribute model can contain exponentially many stable marriages when $k \geq 2$, and that with as few as two attributes the Markov chain can converge exponentially slowly.

Moreover, we show that fewer stable marriage lattices are achievable in the k-attribute model than in the general model, thus necessitating the modified construction. This was shown previously only for the preference sets; Bogomalnaia and Laslier [2] show that there is a preference set that is not realizable with fewer than $n-1$ attributes. However, this does not immediately imply the corresponding result for stable marriage lattices. We show that there are marriage lattices for n people that can only realized with at least $n/2$ attributes.

Next, we show that for the k-range model, rates of convergence provide a more interesting landscape as we vary k, reminiscent of a phase transition. When $k = 1$, there is a unique stable marriage. When $k \geq 2$ there can be an exponential number of stable marriages. We show that when $k = 2$ the marriage lattice must be a hypercube, and hence the random walk will mix in polynomial time. (We conjecture that this remains true when $k = 3$, though the set of lattices have a much richer graph structure.) However, as soon as k is at least 5, there will be preference sets for which the random walk requires exponential time to converge to equilibrium.

The construction we used to show slow mixing for the k-range model can be extended to show slow mixing for the k-list model. It might be surprising to realize that there can be an exponential number of stable marriages, even when there are only a constant number of allowable lists for each gender. However, a careful construction demonstrates that even in this very restrictive model, already when $k = 4$ the structure of the stable marriage lattice can be rich enough to contain an exponentially small cut. We believe that this elucidates the inherent obstacles to sampling stable marriages by random walks on the marriage lattice.

In order to show fast mixing in the k-range model when $k = 2$, we first characterize the allowable stable marriage lattices and show that they are always binary hypercubes. Random walks on the edges of the hypercube are known to converge quickly to equilibrium.

The technical contributions underlying the proofs of slow mixing are two-fold. First we identify a general mechanism for realizing an exponentially small cut in the state space arising from the stable marriage lattice. There is a single stable marriage with exponentially many marriages above and below it in the lattice. The random walk must pass through this vertex to connect the two sets, and the bottleneck can be used to show that the random walk converges slowly to equilibrium. Second, we develop methods for adapting this general mechanism to each of the restricted models, each involving a more careful construction.

For any preference set, the allowable rotations form a poset under the relation that one rotation precedes another if the first must be performed before the second can be performed. We carefully construct this *rotation poset* such that there is a rotation $\mathcal{B}$ which depends on d different independent (in the sense that they can be rotated in any order) rotations $\{\mathcal{A}_0, \mathcal{A}_1, \ldots, \mathcal{A}_d\}$, but is depended on by a set of d independent rotations $\{\mathcal{C}_0, \mathcal{C}_1, \ldots, \mathcal{C}_d\}$. If $\mathcal{B}$ has not yet been rotated, then any subset of the *pre-rotations* $\mathcal{A}_i$ may be rotated independently, so there are at least 2^d marriages in this part of the state space. Similarly, if we have rotated $\mathcal{B}$, we may rotate any subset of the *post-rotations* $\mathcal{C}_i$ independently, so there are at least 2^d more marriages in this second part of the state space. There is a single edge connecting these two sets of marriages, thereby defining the bad cut.

While this construction is fairly straightforward when there are no restrictions on the preference set, it becomes much more challenging in each of our restricted scenarios. However, studying the allowable lattices in each of the restricted models revealed a rich structure worthy of further investigation.

2 Background

2.1 A Markov Chain on the Marriage lattice

In a given marriage $\mathcal{M}$, define a woman's *suitor* to be her favorite of the men that prefer her to their current wives. As the marriage is stable, every woman prefers her husband to her suitor. A *man-improving rotation* is a sequence $\rho = (M_1, w_1, \cdots, M_r, w_r)$ such that M_i is married to w_i and M_{i+1} is the suitor to w_i, where subscripts are modulo r. By dissolving the previous marriages and newly wedding each woman to her suitor, we obtain another stable marriage $\mathcal{M}/\rho$ that is preferred by the men. (see [7] or [11] for details). *Woman-improving rotation* are defined similarly by exchanging genders in all of these definitions.

It can be shown that after performing a rotation improving preferences of one gender, the inverse is a rotation improving the other. Also, the edges of the marriage lattice correspond exactly to allowable rotations [7]. To find all the rotations, we need only list the suitors for each person. Therefore we can use the rotations to define a Markov chain $\mathcal{R}$ on the set of stable marriages as follows. Suppose we are currently at the marriage $\mathcal{M}_t$.

- With probability 1/3, choose a man u.a.r. If he is part of a woman-improving rotation ρ, set $\mathcal{M}_{t+1} = \mathcal{M}_t/\rho$.
- With probability 1/3, choose a man u.a.r. If he is part of a man-improving rotation ρ, set $\mathcal{M}_{t+1} = \mathcal{M}_t/\rho$.
- With probability 1/3, set $\mathcal{M}_{t+1} = \mathcal{M}_t$.

The probability of picking a particular rotation is proportional to the number of couples it contains. Since a rotation and its inverse contain the same people, $\mathcal{R}$ is reversible. It is easy to see that the chain is aperiodic and connects the state space of stable marriages, so the Markov chain converges to the uniform distribution over stable marriages.

2.2 Mixing Time The convergence time of a Markov chain is captured by its mixing time (see, e.g., [14]). The *total variation distance* between distributions is 1/2 of their ℓ_1 distance. The *mixing time* $\tau(\varepsilon)$ of a Markov chain is the time to come within ϵ total variation distance of its stationarity. If the mixing time is upper bounded by a polynomial, we say the chain is *rapidly mixing*. If the mixing time is exponential, we say the chain is *slowly mixing*.

The *conductance* provides a way to upper and lower bound the mixing rate [10, 15], defined as

$$\Phi = \min_{S \subseteq \Omega : \pi(S) \leq 1/2} \frac{\sum_{x \in S, y \notin S} \pi(x) P(x,y)}{\pi(S)}.$$

The following conductance theorem is our main tool for showing slow mixing.

Theorem 2.1. *For any reversible Markov chain with conductance* Φ,

$$\tau(\varepsilon) \geq \frac{1 - 2\Phi}{2\Phi} \ln \varepsilon^{-1}.$$

Thus, to lower bound the mixing time, it is sufficient to show that the conductance is small.

3 The k-attribute model

Recall that in the k-attribute model, men and women are represented by points in R^k, and preference lists are determined by linear objective functions. We show that even in the case that there are only two attributes, there are instances for which the Markov chain $\mathcal{R}$ mixes slowly.

Before looking into the mixing time of $\mathcal{R}$ in the attribute model, we show that, in a robust sense, this model is indeed more restrictive than the general case of arbitrary preference sets. Bogomolnaia and Laslier [2] showed that there exist preference sets on n men which cannot be realized with fewer than $n-1$ attributes. This *cyclic* preference set is shown in Figure 1 for $n = 4$.

$A:$	a	b	c	d	$a:$	D	C	B	A
$B:$	b	c	d	a	$b:$	A	D	C	B
$C:$	c	d	a	b	$c:$	B	A	D	C
$D:$	d	a	b	c	$d:$	C	B	A	D

Figure 1: A preference set, either side of which cannot be realized in the 3-attribute model.

However, while this shows that there are preference sets in the general model which cannot be realized in the k-attribute model, it does not show the same about the rotation posets arising from these preference sets. For instance, if the men and women have the disallowed preference set as in Figure 1, the marriage lattice is simply a path of n marriages, each dependent on the previous, as shown in Figure 2. This marriage lattice (and the equivalent rotation poset) can be easily realized in the k-attribute model, even when $k = 2$.

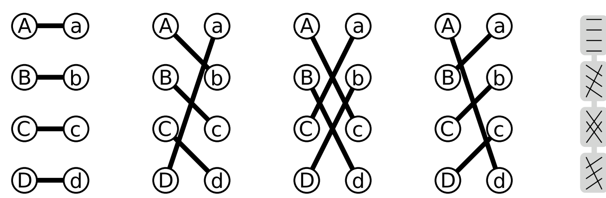

Figure 2: The range of stable marriages of Figure 1, with the marriage lattice to the far right

We show that the rotation posets realized in the k-attribute model are, in fact, strictly contained within the rotation posets arising in the general model. The unrealizable preference set in Figure 1 corresponded to a realizable rotation poset because there was an alternate preference set for the same rotation poset. To create a rotation poset with no realizable preference set, we remove this flexibility by creating a poset with the maximum number of rotations. We can then show the following.

Theorem 3.1. *There exist rotation posets on n men which cannot be realized by any preference set in the k-attribute model when $k < n/2$.*

Proof. The maximum number of rotations for any n men is $\frac{n(n-1)}{2}$. To see this, define the *improvement* of a rotation to be the total distance the wives move on the men's preference lists. For instance, if a rotation on three men improves their allotment from their least favorite to second-to-least favorite, the rotation has improvement 3. Clearly the minimum improvement of any rotation is 2, as a rotation involves at least two men and they must improve at least one position each. On the other hand, the maximum sum of the improvements of all rotations is $n(n-1)$, as the n men can only improve $n-1$ positions each. Therefore there can be at most $\frac{n(n-1)}{2}$ rotations.

A rotation poset with $\frac{n(n-1)}{2}$ rotations is illustrated in Figure 3. It consists of $n-1$ rows of $n/2$ independent rotations, where the edges between any two rows forms a $2n$ cycle. We will show that, up to permutation of labels, this rotation poset has only one possible preference set.

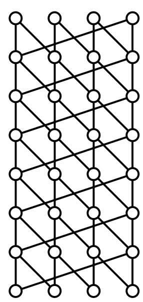

Figure 3: A rotation poset on 8 men which cannot be realized in the 3-attribute model.

We know that the male-optimal marriage weds each man to his favorite woman or else the number of rotations would not be maximal. Without loss of generality, let these favorites be (M_i, w_i). Again by the maximality, the n independent initial rotations must involve only pairs of men swapping their two wives, and decrease the men's ranking by only one position each. We may assume these pairs are (M_1, M_2), (M_3, M_4), $\ldots$, (M_{n-1}, M_n). Therefore the men's preference set begins as in the upper left of Figure 4.

The next round of rotations must decrease the men from their second choice to their third, and again involve only pairs of men swapping wives. The only way for a pair-rotation to depend on both the (M_1, M_2) rotation and the (M_3, M_4) rotation, is for it to involve one man from each, for instance (M_2, M_3). If a pair-rotation depends on (M_3, M_4) and (M_5, M_6), it must then involve (M_4, M_5). This allows us to progressively build the preference list up to the third choice, as in the upper right of Figure 4.

This pattern continues, until we have finished the

M_1 :	w_1	w_2
M_2 :	w_2	w_1
M_3 :	w_3	w_4
M_4 :	w_4	w_3
M_5 :	w_5	w_6
M_6 :	w_6	w_5
M_7 :	w_7	w_8
M_8 :	w_8	w_7

M_1 :	w_1	w_2	w_7
M_2 :	w_2	w_1	w_4
M_3 :	w_3	w_4	w_1
M_4 :	w_4	w_3	w_6
M_5 :	w_5	w_6	w_3
M_6 :	w_6	w_5	w_8
M_7 :	w_7	w_8	w_5
M_8 :	w_8	w_7	w_1

M_1 :	w_1	w_2	w_7	w_4	w_5	w_6	w_3	w_8
M_2 :	w_2	w_1	w_4	w_7	w_6	w_5	w_8	w_3
M_3 :	w_3	w_4	w_1	w_6	w_7	w_8	w_5	w_2
M_4 :	w_4	w_3	w_6	w_1	w_8	w_7	w_2	w_5
M_5 :	w_5	w_6	w_3	w_8	w_1	w_2	w_7	w_4
M_6 :	w_6	w_5	w_8	w_3	w_2	w_1	w_4	w_7
M_7 :	w_7	w_8	w_5	w_2	w_3	w_4	w_1	w_6
M_8 :	w_8	w_7	w_2	w_5	w_4	w_3	w_6	w_1

Figure 4: The incremental construction of the men's preference set from Figure 3.

preference set as on be bottom of Figure 4. However, the men with even subscripts form a cyclic preference set on $n/2$ people which, as stated above, can not be realized in fewer than $n/2$ attributes.

We show that although the k-attribute model is a real restriction on the possible marriage lattices, $\mathcal{R}$ may still mix slowly. First it will be useful to give a general construction for slow mixing that will be the starting point for each of our subsequent constructions.

3.1 Slow mixing in the general model We begin by constructing a specific instance of the *general* stable marriage problem for which the Markov chain mixes slowly. It is known that every distributive lattice is isomorphic to some marriage polytope [1, 8]. Since it is easy to construct examples of lattices where a random walk on the edges will mix slowly, it is not surprising that this approach fails for the problem of sampling stable marriages in general. However, our particular construction is the crux to showing that even for very restricted marriage lattices, the chain will mix slowly.

Theorem 3.2. *There is a preference set on n men and n women on which, for some constant $c_1 > 0$, the mixing time of $\mathcal{R}$ is $e^{c_1 n}$.*

Proof. We design preference sets so that there is one stable marriage that forms a bottleneck in the stable marriage lattice, with an exponential number of stable marriages above and below it. This establishes that the conductance is exponentially small and consequently, by Theorem 2.1, the Markov chain $\mathcal{R}$ is slow mixing.

The construction involves $5d$ people of each gender. The construction's rotation poset will be arranged such that there is a rotation $\mathcal{B}$ which depends on d different independent rotations $\{\mathcal{A}_0, \mathcal{A}_1, \ldots, \mathcal{A}_d\}$, but is depended on by a set of d independent rotations $\{\mathcal{C}_0, \mathcal{C}_1, \ldots, \mathcal{C}_d\}$. If we have not rotated $\mathcal{B}$, we may rotate any subset of these *pre-rotations* $\mathcal{A}_i$, so there are 2^d marriages. Similarly, if we have rotated $\mathcal{B}$ we may rotate any subset of these *post-rotations* $\mathcal{C}_i$, so there are 2^d more marriages. Such a construction is illustrated in Figure 5.

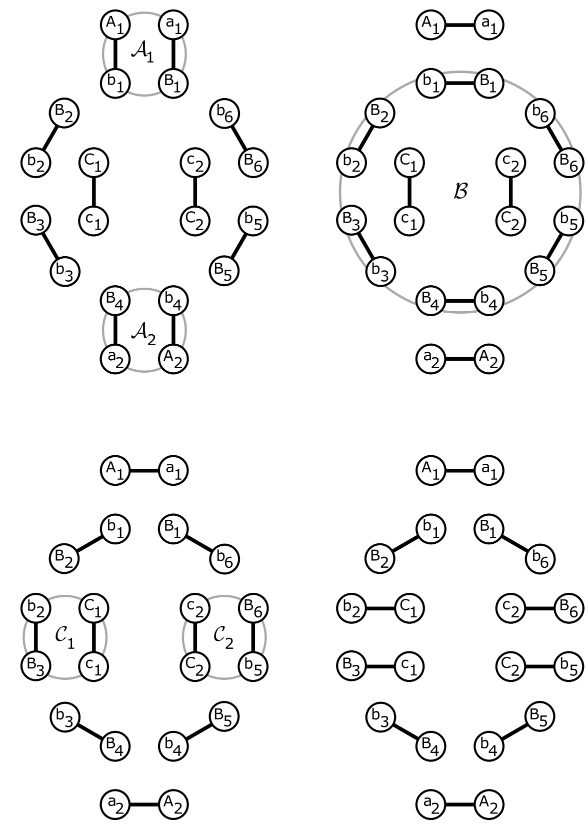

Figure 5: Four stable marriages, ranging from man-optimal to woman-optimal. The suitor-rotations are labeled before they are rotated.

We use capital letters to represent men and lowercase for women throughout. For instance, our large rotation $\mathcal{B} = (B_1, b_1, B_2, \ldots, b_{3d})$. For the men and women used exclusively in pre-rotation $\mathcal{A}_i$, we use A_i and a_i, so $\mathcal{A}_i = (A_i, a_i, B_{4i-3}, b_{4i-3})$. Similarly, we use C_i and c_i for those exclusively in post-rotation $\mathcal{C}_i$, so $\mathcal{C}_i = (C_i, c_i, B_{3i}, b_{3i-1})$. This is illustrated in Figure 5.

Building preference sets in the general model that gives rise to such a rotation poset is straightforward; we need only go through our various marriages and list the spouses for each person, then move those spouses to the front of that person's preference list. For instance, if man B_1 is married to first a_1, then b_1, and then b_{3d}, his preference list is $(a_1, b_1, b_{3d}, \ldots)$. Note that the women after these spouses are irrelevant, as he will always have one of his first three choices available. Woman b_1 is married to A_1, B_1, then B_2. However, these rotations sequentially decrease b_1's happiness, so her preference list is $(B_2, B_1, A_1, \ldots)$. Again, the remaining preference lists are unimportant. The partial preference set for Figure 5 is depicted in Figure 6.

A_1 :	$\mathbf{b_1}$	$\mathbf{a_1}$	
A_2 :	$\mathbf{b_4}$	$\mathbf{a_2}$	
C_1 :	$\mathbf{c_1}$	$\mathbf{b_2}$	
C_2 :	$\mathbf{c_2}$	$\mathbf{b_5}$	
B_1 :	$\mathbf{a_1}$	$\mathbf{b_1}$	$\mathbf{b_6}$
B_2 :	$\mathbf{b_2}$	$\mathbf{b_1}$	
B_3 :	$\mathbf{b_3}$	$\mathbf{b_2}$	$\mathbf{c_1}$
B_4 :	$\mathbf{a_2}$	$\mathbf{b_4}$	$\mathbf{b_3}$
B_5 :	$\mathbf{b_5}$	$\mathbf{b_4}$	
B_6 :	$\mathbf{b_6}$	$\mathbf{b_5}$	$\mathbf{c_2}$

b_1 :	$\mathbf{B_2}$	$\mathbf{B_1}$	$\mathbf{A_1}$
b_2 :	$\mathbf{C_1}$	$\mathbf{B_3}$	$\mathbf{B_2}$
b_3 :	$\mathbf{B_4}$	$\mathbf{B_3}$	
b_4 :	$\mathbf{B_5}$	$\mathbf{B_4}$	$\mathbf{A_2}$
b_5 :	$\mathbf{C_2}$	$\mathbf{B_6}$	$\mathbf{B_5}$
b_6 :	$\mathbf{B_1}$	$\mathbf{B_6}$	
a_1 :	$\mathbf{A_1}$	$\mathbf{B_1}$	
a_2 :	$\mathbf{A_2}$	$\mathbf{B_4}$	
c_1 :	$\mathbf{B_3}$	$\mathbf{C_1}$	
c_2 :	$\mathbf{B_6}$	$\mathbf{C_2}$	

Figure 6: The preference set corresponding to Figure 5

This construction of the preference lists guarantees stability, as the only people someone might prefer to their spouse in one marriage are their spouses in better (from their perspective) marriages. The object of affection therefore prefers their current spouse; they do not favor switching and we have stability.

On this preference set, the mixing time of $\mathcal{R}$ is exponential in n. The above construction helps display a general mechanism for showing slow-mixing in the restricted models.

3.2 Slow mixing in the k-attribute Model We now proceed to adapt the construction from Section 3.1 to the k-attribute model.

THEOREM 3.3. *If $k \geq 2$, there are k-attribute preference sets for which there is a constant $c_2 > 0$ such that the mixing time of $\mathcal{R}$ is $e^{c_2 n}$.*

Here we first prove the case when $k \geq 3$, and then extend this to the case when $k = 2$.

Proof. We build the construction in Figure 5 in R^3. In doing so, we will place all women on the quarter unit ball. That is, $\forall w$, $||\mathbf{v}_w|| = 1$. For any three equidistant neighboring women on the ball, a man may choose these to be his favorite; he simply finds a vector that is a convex combination of the three women's locations, and lets his function correspond to that vector. Furthermore, he can arbitrarily choose the order of those three points by letting the average weigh more heavily towards his first and then second choice. As Figure 6 needs only the first three people in any preference list, will can recreate the preference sets exactly.

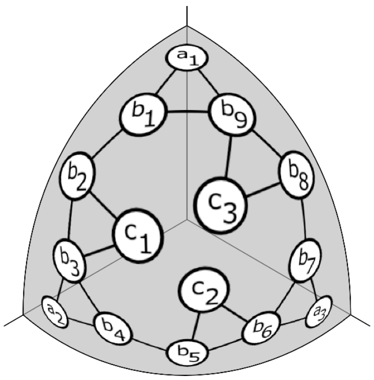

Figure 7: The arrangement of women in 3-space when $k = 3$.

In actually placing the women, we need only put the large rotation in a circle, and the women for the other rotations near their neighbors in Figure 5. This is illustrated in Figure 7. Note that, for each man, there is some face or edge to which all his wives are incident. This way they may be placed at the very beginning of his preference list, and our marriages are stable, proving Theorem 3.3 in three or more dimensions.

When $k = 2$, again we place the women on the unit ball; this time in two dimensions. Notice that a man can choose any adjacent pair as his favorites (and their order). He can then choose a woman incident to that pair as his third, with a few exceptions. (Figure 8 shows such an exception. The only way for a preference list to begin $\{x, y, z\}$ would be if the distance from y to z was strictly less than the distance from w to x.)

There is no way to place all spouses at the very beginning of the preference lists as we did in Figure 6, even if we're simply building a long rotation (with no pre- or post-rotations). However, if we arrange the women in order on the unit ball $\{b_1, b_2, \ldots, b_{3k}\}$, and do the same with the men, then all but one of the preference lists can place the spouses at the very beginning, and we still have stability (as in Figure 9).

Inserting the women for pre- and post-rotations, we place a_i just before b_{3i-2} and place c_i just before b_{3i-1}. This will place even more women in front of the spouses,

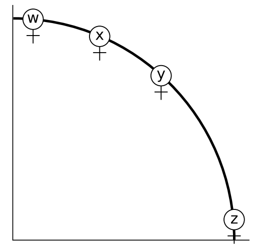

Figure 8: An example where, although adjacent, no preference list may begin with $\{x, y, z\}$.

which might create instability. Luckily, these additional women are all married to their first or second choice, so stability is preserved. (The extra men are inserted similarly, with A_i just behind B_{3i-2} and C_i just in front of B_{3i-1}.) This is illustrated in Figure 10. This allows us to create a preference set, as illustrated in Figure 11.

B_1 : $\mathbf{b_1}$ b_2 b_3 b_4 b_5 $\mathbf{b_6}$	b_1 : $\mathbf{B_2}$ $\mathbf{B_1}$
B_2 : $\mathbf{b_2}$ $\mathbf{b_1}$	b_2 : $\mathbf{B_3}$ $\mathbf{B_2}$
B_3 : $\mathbf{b_3}$ $\mathbf{b_2}$	b_3 : $\mathbf{B_4}$ $\mathbf{B_3}$
B_4 : $\mathbf{b_4}$ $\mathbf{b_3}$	b_4 : $\mathbf{B_5}$ $\mathbf{B_4}$
B_5 : $\mathbf{b_5}$ $\mathbf{b_4}$	b_5 : $\mathbf{B_6}$ $\mathbf{B_5}$
B_6 : $\mathbf{b_6}$ $\mathbf{b_5}$	b_6 : $\mathbf{B_1}$ B_2 B_3 B_4 B_5 $\mathbf{B_6}$

Figure 9: The preference set from a 6-rotation in 2-dimensions.

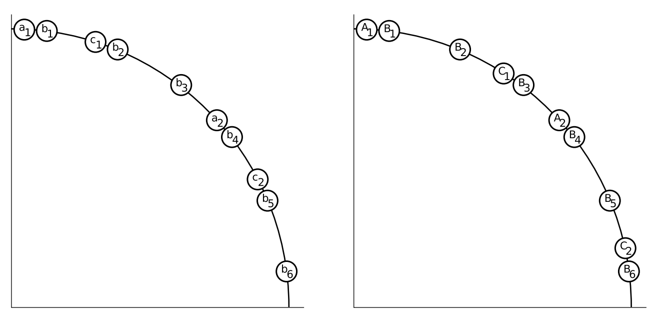

Figure 10: The arrangement of the women in 2-space, when $k = 2$.

Note that the construction above barely avoids the problem illustrated in Figure 8. For instance, b_2's preference list begins $\{C_1, B_3, B_2\}$. If she preferred these in a different order, for example $\{B_3, C_1, B_2\}$, the arrangement in 2-space must be changed, as the distance from B_2 to C_1 is exactly the same as B_3 to A_2.

$A_1 : \mathbf{b_1}\ \mathbf{a_1}$	$B_1 : \mathbf{a_1}\ \mathbf{b_1}\ c_1\quad b_2\ b_3\ a_2\ b_4\ c_2\ b_5\ \mathbf{b_6}$
$A_2 : \mathbf{b_4}\ \mathbf{a_2}$	$B_2 : \mathbf{b_2}\ c_1\ \mathbf{b_1}$
	$B_3 : \mathbf{b_3}\ \mathbf{b_2}\ \mathbf{c_1}$
$C_1 : \mathbf{c_1}\ \mathbf{b_2}$	$B_4 : \mathbf{a_2}\ \mathbf{b_4}\ \mathbf{b_3}$
$C_2 : \mathbf{c_2}\ \mathbf{b_5}$	$B_5 : \mathbf{b_5}\ c_2\ \mathbf{b_4}$
	$B_6 : \mathbf{b_6}\ \mathbf{b_5}\ \mathbf{c_2}$
$a_1 : \mathbf{A_1}\ \mathbf{B_1}$	$b_1 : \mathbf{B_2}\ \mathbf{B_1}\ \mathbf{A_1}$
$a_2 : \mathbf{A_2}\ \mathbf{B_4}$	$b_2 : \mathbf{C_1}\ \mathbf{B_3}\ \mathbf{B_2}$
	$b_3 : \mathbf{B_4}\ A_2\ \mathbf{B_3}$
$c_1 : \mathbf{B_3}\ \mathbf{C_1}$	$b_4 : \mathbf{B_5}\ \mathbf{B_4}\ \mathbf{A_2}$
$c_2 : \mathbf{B_6}\ \mathbf{C_2}$	$b_5 : \mathbf{C_2}\ \mathbf{B_6}\ \mathbf{B_5}$
	$b_6 : A_1\ \mathbf{B_1}\ B_2\quad C_1\ B_3\ A_2\ B_4\ B_5\ C_2\ \mathbf{B_6}$

Figure 11: The people from Figure 5 arranged in 2-dimensions and the corresponding preference lists.

4 The k-range model

Recall that in the k-range model, each person appears within k possible positions on the other genders' preference lists. If either gender has range 1, then every member of the opposite gender has the same preference list, and there is only one stable marriage. When $k \geq 2$ the number of stable marriages can be exponential. The following two theorems show a type of "phase transition" in the mixing rate as k increases.

THEOREM 4.1. *When $k = 2$, there are k-range preference sets that allow exponentially many stable marriages, but the mixing time of $\mathcal{R}$ is always bounded by a polynomial in n.*

Interestingly, increasing the range slightly yields very different behavior in the mixing rate.

THEOREM 4.2. *For $k \geq 5$, there is a k-range preference set so that, for some constant c, the mixing time of $\mathcal{R}$ is at least e^{cn}.*

We now outline proofs of these two theorems to explain the dichotomy in the mixing rates as we vary k.

Proof of Theorem 4.1: To prove $\mathcal{R}$ mixes in polynomial time on 2-range preference sets, we use the fact that if a preference set has range 2, removing an individual from the lists never increases the range. We claim that the marriage lattice arising from a preference set with range 2 must be a hypercube, and therefore a random walk on its edges is fast-mixing. The proof of this claim that the lattice is a hypercube is by induction.

If all the men agree that some woman is best, she gets her first choice in any stable marriage. We can remove the newlyweds from the other preference lists and use induction.

If there is a pair of people who are both first on the other's list, then any stable marriage will wed these two. Again, we can remove both from the other preference list and use induction.

If neither of these cases occur, then there are two men at the top of the women's lists, and two women at the top of the men's lists, and these four people form a rotation. There are two ways to wed this group, but after choosing such a pairing, we can remove all four from others' lists. The remaining preferences will be the same regardless of the pairing we chose, and we again use induction.

Therefore the marriage lattice is simply the binary hypercube with dimension equal to the number of rotations we found. Walks on the hypercube are well known to be fast mixing (see, e.g., [3]).

B_3 :	$\mathbf{b_3}$	$\mathbf{b_1}$	b_2	b_4	b_5	b_6	b_7	b_8	b_9	b_{10}	b_{11}	b_{12}
B_5 :	b_1	b_2	$\mathbf{b_5}$	$\mathbf{b_3}$	b_4	b_6	b_7	b_8	b_9	b_{10}	b_{11}	b_{12}
B_7 :	b_1	b_2	b_3	b_4	$\mathbf{b_7}$	$\mathbf{b_5}$	b_7	b_8	b_9	b_{10}	b_{11}	b_{12}
B_9 :	b_1	b_2	b_3	b_4	b_5	b_6	$\mathbf{b_9}$	$\mathbf{b_7}$	b_8	b_{10}	b_{11}	b_{12}
B_{11} :	b_1	b_2	b_3	b_4	b_5	b_6	b_7	b_8	$\mathbf{b_{11}}$	$\mathbf{b_9}$	b_{10}	b_{12}
B_{12} :	b_1	b_2	b_3	b_4	b_5	b_6	b_9	b_7	b_8	b_{10}	$\mathbf{b_{12}}$	$\mathbf{b_{11}}$
B_{10} :	b_1	b_2	b_3	b_4	b_5	b_6	b_7	b_8	b_9	$\mathbf{b_{12}}$	$\mathbf{b_{10}}$	b_{11}
B_8 :	b_1	b_2	b_3	b_4	b_5	b_6	b_7	$\mathbf{b_8}$	$\mathbf{b_{10}}$	b_9	b_{11}	b_{12}
B_6 :	b_1	b_2	b_3	b_4	b_5	$\mathbf{b_6}$	$\mathbf{b_8}$	b_7	b_9	b_{10}	b_{11}	b_{12}
B_4 :	b_1	b_2	b_3	$\mathbf{b_4}$	$\mathbf{b_6}$	b_5	b_7	b_8	b_9	b_{10}	b_{11}	b_{12}
B_2 :	b_1	$\mathbf{b_2}$	$\mathbf{b_4}$	b_3	b_5	b_6	b_7	b_8	b_9	b_{10}	b_{11}	b_{12}
B_1 :	$\mathbf{b_1}$	$\mathbf{b_2}$	b_3	b_4	b_5	b_6	b_7	b_8	b_9	b_{10}	b_{11}	b_{12}

Figure 12: The male preferences for a 12-rotation with range 4.

Proof of Theorem 4.2: In order to show that the random walk on the stable marriage lattice is slowly mixing when $k \geq 5$, we would like to create a k-range example as in Section 3.1. Unlike in the previous examples, the difficulty here is creating an arbitrarily long rotation. If there is a rotation on $(B_1, B_2, \ldots, B_{3d})$, then some woman will be married to a man of high value before the rotation (B_1), but a man of low value after (B_{3d}). No matter her own worth, this cannot be stable for large d.

Instead, our construction of the long cycle proceeds up along the odd numbered men and then back along the even, so their quality never varies too much along the cycle; with a rotation on $(B_1, B_3, \ldots, B_{3d}, B_{3d-1}, B_{3d-3}, \ldots, B_2)$, the value of the men slowly increases and then slowly decreases.

To build the long cycle, consider the worth of the men ordered as $\{B_1, \ldots, B_{3d}\}$. A woman's preference list will be a copy of this ranking, but with her own husband pushed slightly earlier. More formally, for each woman w, we define her *stratum* as the set S_w that includes all of her potential husbands and all of the men between them in worthiness. For instance, if b_3 is married to B_5 and then B_3, $\mathcal{S}_{b_3} = \{B_3, B_4, B_5\}$. To build w's preference list, we use the natural order on the men above and below $\mathcal{C}_w$, but order her stratum such that her husbands are first. (Of course the same process is done with the genders reversed for the women's preference lists.) This construction is illustrated in Figure 13.

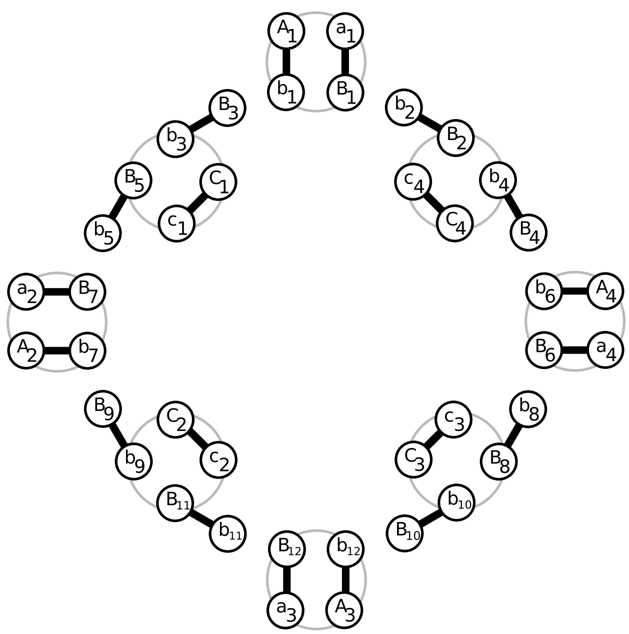

Figure 13: The 5-range construction when $k = 2$.

We build the full construction, with pre- and post-rotations $\mathcal{A}_i$ and $\mathcal{C}_i$, in much the same way. However, to ensure that adding the new men and women does not add too much to the range, we only insert the A_i and C_i on the odd side of the cycle. This will cut the number of independent cycles in half, but still give us exponentially many marriages on each side of the cut.

Inserting the remaining people into the ordering $\{B_1, B_2, \ldots\}$, we put A_i next to the other man in $\mathcal{A}_i$ and C_i next to the other man in $\mathcal{C}_i$. As seen in Figure 13, A_i is before B_{6i-2} and C_i is before B_{6i-1}. The full set of strata is illustrated in Figure 14. To fill in the preference lists, one needs only insert the missing women in order, as in Figure 15

It is clear that such a construction will give k-range for some small constant k, as each of the strata are of constant size. Analyzing this specific construction, we find that this example can be constructed with range 5.

A_1 :	$\mathbf{b_1}\ \mathbf{a_1}$
A_2 :	$\mathbf{b_7}\ \mathbf{a_2}$
C_1 :	$\mathbf{c_1}\ \mathbf{b_3}$
C_2 :	$\mathbf{b_9}\ \ \mathbf{c_2}$
B_1 :	$\mathbf{a_1}\ \mathbf{b_1}\ \mathbf{b_2}$
B_3 :	$\mathbf{b_3}\ \mathbf{b_1}\ b_2$
B_5 :	$\mathbf{b_5}\ \mathbf{b_3}\ \mathbf{c_1}\ b_4$
B_7 :	$\mathbf{a_2}\ \mathbf{b_7}\ \mathbf{b_5}\ b_6$
B_9 :	$\mathbf{b_9}\ \mathbf{b_7}\ b_8$
B_{11} :	$\mathbf{b_{11}}\ \mathbf{b_9}\ \mathbf{c_2}\ \ b_{10}$
B_{12} :	$\mathbf{b_{12}}\ \mathbf{b_{11}}$
B_{10} :	$\mathbf{b_{10}}\ \mathbf{b_{12}}\ b_{11}$
B_8 :	$\mathbf{b_8}\ \mathbf{b_{10}}\ b_9\ \ c_2$
B_6 :	$\mathbf{b_6}\ \mathbf{b_8}\ a_2\ \ b_7$
B_4 :	$\mathbf{b_4}\ \mathbf{b_6}\ b_5$
B_2 :	$\mathbf{b_2}\ \mathbf{b_4}\ \mathbf{b_3}\ c_1$

Figure 14: The set of strata for a 5-range construction.

A_1 :	$\mathbf{b_1}$	$\mathbf{a_1}$	b_2	b_3	c_1	b_4	b_5	b_6	a_2	b_7	b_8	b_9	c_2	b_{10}	b_{11}	b_{12}
A_2 :	a_1	b_1	b_2	b_3	c_1	b_4	b_5	b_6	$\mathbf{b_7}$	$\mathbf{a_2}$	b_8	b_9	c_2	b_{10}	b_{11}	b_{12}
C_1 :	a_1	b_1	b_2	$\mathbf{c_1}$	$\mathbf{b_3}$	b_4	b_5	b_6	a_2	b_7	b_8	b_9	c_2	b_{10}	b_{11}	b_{12}
C_2 :	a_1	b_1	b_2	b_3	c_1	b_4	b_5	b_6	a_2	b_7	b_8	$\mathbf{b_9}$	$\mathbf{c_2}$	b_{10}	b_{11}	b_{12}
B_1 :	$\mathbf{a_1}$	$\mathbf{b_1}$	$\mathbf{b_2}$	b_3	c_1	b_4	b_5	b_6	a_2	b_7	b_8	b_9	c_2	b_{10}	b_{11}	b_{12}
B_3 :	a_1	$\mathbf{b_3}$	$\mathbf{b_1}$	b_2	c_1	b_4	b_5	b_6	a_2	b_7	b_8	b_9	c_2	b_{10}	b_{11}	b_{12}
B_5 :	a_1	b_1	b_2	$\mathbf{b_5}$	$\mathbf{b_3}$	$\mathbf{c_1}$	b_4	b_6	a_2	b_7	b_8	b_9	c_2	b_{10}	b_{11}	b_{12}
B_7 :	a_1	b_1	b_2	b_3	c_1	b_4	$\mathbf{a_2}$	$\mathbf{b_7}$	$\mathbf{b_5}$	b_6	b_8	b_9	c_2	b_{10}	b_{11}	b_{12}
B_9 :	a_1	b_1	b_2	b_3	c_1	b_4	b_5	b_6	a_2	$\mathbf{b_9}$	$\mathbf{b_7}$	b_8	c_2	b_{10}	b_{11}	b_{12}
B_{11} :	a_1	b_1	b_2	b_3	c_1	b_4	b_5	b_6	a_2	b_7	b_8	$\mathbf{b_{11}}$	$\mathbf{b_9}$	$\mathbf{c_2}$	b_{10}	b_{12}
B_{12} :	a_1	b_1	b_2	b_3	c_1	b_4	b_5	b_6	a_2	b_7	b_8	b_9	c_2	b_{10}	$\mathbf{b_{12}}$	$\mathbf{b_{11}}$
B_{10} :	a_1	b_1	b_2	b_3	c_1	b_4	b_5	b_6	a_2	b_7	b_8	b_9	c_2	$\mathbf{b_{10}}$	$\mathbf{b_{12}}$	b_{11}
B_8 :	a_1	b_1	b_2	b_3	c_1	b_4	b_5	b_6	a_2	b_7	$\mathbf{b_8}$	$\mathbf{b_{10}}$	b_9	c_2	b_{11}	b_{12}
B_6 :	a_1	b_1	b_2	b_3	c_1	b_4	b_5	$\mathbf{b_6}$	$\mathbf{b_8}$	a_2	b_7	b_9	c_2	b_{10}	b_{11}	b_{12}
B_4 :	a_1	b_1	b_2	b_3	c_1	$\mathbf{b_4}$	$\mathbf{b_6}$	b_5	a_2	b_7	b_8	b_9	c_2	b_{10}	b_{11}	b_{12}
B_2 :	a_1	b_1	$\mathbf{b_2}$	$\mathbf{b_4}$	$\mathbf{b_3}$	c_1	b_5	b_6	a_2	b_7	b_8	b_9	c_2	b_{10}	b_{11}	b_{12}

Figure 15: The 5-range preference set when $k = 2$.

5 The k-list model

THEOREM 5.1. *If $k \geq 4$, there are k-list preference sets such that, for some constant c', the mixing time of $\mathcal{R}$ is at least $e^{c'n}$.*

$$
\begin{array}{ll}
A_1: & \mathbf{b_1}\ \mathbf{a_1}\ b_2\ b_4\ b_3\ c_1\ b_5\ b_6\ b_8\ a_2\ b_7\ b_9\ c_2\ b_{10}\ b_{12}\ b_{11} \\
A_2: & a_1\ b_1\ b_2\ c_1\ b_3\ b_4\ b_6\ b_5\ \mathbf{b_7}\ \mathbf{a_2}\ b_8\ b_{10}\ b_9\ c_2\ b_{12}\ b_{11} \\
C_1: & a_1\ b_1\ b_2\ \mathbf{c_1}\ \mathbf{b_3}\ b_4\ b_6\ b_5\ b_7\ a_2\ b_8\ b_{10}\ b_9\ c_2\ b_{12}\ b_{11} \\
C_2: & b_1\ a_1\ b_2\ b_4\ b_3\ c_1\ b_5\ b_6\ b_8\ a_2\ b_7\ \mathbf{b_9}\ \mathbf{c_2}\ b_{10}\ b_{12}\ b_{11} \\
B_1: & \mathbf{a_1}\ \mathbf{b_1}\ \mathbf{b_2}\ b_5\ b_3\ c_1\ b_4\ b_6\ a_2\ b_9\ b_7\ b_8\ c_2\ b_{10}\ b_{12}\ b_{11} \\
B_3: & a_1\ \mathbf{b_3}\ \mathbf{b_1}\ \mathbf{b_2}\ c_1\ b_4\ a_2\ b_7\ b_5\ b_6\ b_8\ b_{11}\ b_9\ c_2\ b_{10}\ b_{12} \\
B_5: & a_1\ b_1\ b_2\ \mathbf{b_5}\ \mathbf{b_3}\ \mathbf{c_1}\ b_4\ b_6\ a_2\ b_9\ b_7\ b_8\ c_2\ b_{10}\ b_{12}\ b_{11} \\
B_7: & a_1\ b_3\ b_1\ b_2\ c_1\ b_4\ \mathbf{a_2}\ \mathbf{b_7}\ \mathbf{b_5}\ b_6\ b_8\ b_{11}\ b_9\ c_2\ b_{10}\ b_{12} \\
B_9: & a_1\ b_1\ b_2\ b_5\ b_3\ c_1\ b_4\ b_6\ a_2\ \mathbf{b_9}\ \mathbf{b_7}\ b_8\ c_2\ b_{10}\ b_{12}\ b_{11} \\
B_{11}: & a_1\ b_3\ b_1\ b_2\ c_1\ b_4\ a_2\ b_7\ b_5\ b_6\ b_8\ \mathbf{b_{11}}\ \mathbf{b_9}\ \mathbf{c_2}\ b_{10}\ b_{12} \\
B_{12}: & a_1\ b_1\ b_2\ c_1\ b_3\ b_4\ b_6\ b_5\ b_7\ a_2\ b_8\ b_{10}\ b_9\ c_2\ \mathbf{b_{12}}\ \mathbf{b_{11}} \\
B_{10}: & b_1\ a_1\ b_2\ b_4\ b_3\ c_1\ b_5\ b_6\ b_8\ a_2\ b_7\ b_9\ c_2\ \mathbf{b_{10}}\ \mathbf{b_{12}}\ b_{11} \\
B_8: & a_1\ b_1\ b_2\ c_1\ b_3\ b_4\ b_6\ b_5\ b_7\ a_2\ \mathbf{b_8}\ \mathbf{b_{10}}\ b_9\ c_2\ b_{12}\ b_{11} \\
B_6: & b_1\ a_1\ b_2\ b_4\ b_3\ c_1\ b_5\ \mathbf{b_6}\ \mathbf{b_8}\ a_2\ b_7\ b_9\ c_2\ b_{10}\ b_{12}\ b_{11} \\
B_4: & a_1\ b_1\ b_2\ c_1\ b_3\ \mathbf{b_4}\ \mathbf{b_6}\ b_5\ b_7\ a_2\ b_8\ b_{10}\ b_9\ c_2\ b_{12}\ b_{11} \\
B_2: & b_1\ a_1\ \mathbf{b_2}\ \mathbf{b_4}\ b_3\ c_1\ b_5\ b_6\ b_8\ a_2\ b_7\ b_9\ c_2\ b_{10}\ b_{12}\ b_{11}
\end{array}
$$

Figure 16: A slow construction in 4-list when $k = 2$

Proof. In the construction in the proof of Theorem 4.2, the preference lists only needed a certain order in the individual strata. Above or below the strata, the lists had a great deal of flexibility. In fact, if two strata are disjoint, a single preference list could contain both. Combining strata in this way will give us a slow construction with a constant number of lists. It happens that our construction can be consolidated into four lists, as illustrated in Figure 16.

6 Conclusions

We have demonstrated that random walks on the stable marriage lattice are slowly mixing, even in very restricted scenarios. Finding alternative methods for sampling stable marriages in any of these restrictive settings is an interesting open problem that might shed light on the general problem of sampling on distributive lattices.

It remains open whether $\mathcal{R}$ is fast for the k-range model when $k = 3$ or 4. We conjecture that it is when $k = 3$. In this case we believe the allowable lattices can be expressed as a product of a small family of allowable graphs, which would imply that random walks on all the allowable lattices mix in polynomial time. We leave $k = 4$ as an open problem.

Finally, we point out that the models defined in this paper, in particular the k-attribute and k-range models, are worthy of further study. For instance, given a set of preference set, can we efficiently decide whether it could have arisen from a k-attribute model? I.e., can we associate the members of the one gender with points in R^k so that the preference lists arise as projections of these points onto a line. It also seems worthwhile to ask whether the algorithms for constructing stable marriages can be implemented more efficiently if the the preference lists come from one of these restricted scenarios. For example, we can show that we can find the male- and female-optimal marriage more efficiently in the k-range model.

References

[1] C. Blair. Every Finite Distributive Lattice is a set of Stable Matchings. *Journal of Combinatorial Theory (A)*, 37: 353-356, 1984.

[2] A. Bogomolnaia and J-F. Laslier. Euclidean Preferences. Preprint: submitted to *the Journal of Economic Theory*, 2004.

[3] P. Diaconis, R.L. Graham and J.A. Morrison. Asymptotic Analysis of a Random Walk on a Hypercube with Many Dimensions. *Random Structures and Algorithms*, 1: 51–72, 1990.

[4] M. Dyer, L.A. Goldberg, C. Greenhill, M. Jerrum. On the Relative Complexity of Approximate Counting Problems *Algorithmica* 38: 471 - 500, 2003.

[5] D. Gale and L.S. Shapley. College admissions and the Stability of Marriage. *American Mathematical Monthly*, 69: 9-15, 1962

[6] D. Gusfield. Three Fast Algorithms for Four Problems in Stable Marriage. *SIAM Journal on Computing*, 16: 111-128, 1987.

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

[8] D. Gusfield, R. Irving, P. Leather, and M. Saks. Every Finite Distributive Lattice is a Set of Stable Matchings for a Small Stable Marriage Instance. *Journal of Combinatorial Theory (A)*, 44: 304-309, 1987.

[9] R. Irving, P. Leather, D. Gusfield. An Efficient Algorithm for the "Optimal" Stable Marriage. *Journal of the A.C.M.*, 34:532-543, 1987.

[10] G.F. Lawler and A.D. Sokal, A.D. Bounds on the L2 spectrum for Markov chain and Markov processes: a generalization of Cheeger's inequality. *Trans. Amer. Math. Soc.*, 309: 557-580, 1988.

[11] D.E. Knuth. *Marriages Stables. Les Presses de l'Universite de Montreal, Montreal*, 1976.

[12] A. Mehrez, Y. Yuan, A. Gafni. Stable Solutions vs. Multiplicative Utility Solutions for the Assignment Problem. *Operations Research Letters*, 7:131-139, 1988.

[13] S. Menju. The Stable Marriage Assignment and Related problems. Technical Report C-83, Tokyo Institute of Technology, 1987.

[14] A. Sinclair. Algorithms for Random Generation and Counting. *Birkhauser*, 1993

[15] A. Sinclair and M. Jerrum. Approximate Counting, Uniform Generation and Rapidly Mixing Markov Chains *Inf. Comput.*, 82: 93-133, 1989.

[16] C-P. Teo and J. Sethuraman. The Geometry of Fractional Stable Matchings and its Applications. *Mathematics of Operations Research.*, 23(4): 874-891, 1998.