Efficient, Local and Symmetric Markov Chains that Generate One-Factorizations

Maya Dotan* Nati Linial† Yuval Peled‡

Abstract

It is well known that for every even integer n, the complete graph K_n has a one-factorization, namely a proper edge coloring with $n-1$ colors. Unfortunately, not much is known about the possible structure of large one-factorizations. Also, at present we have only woefully few explicit constructions of large one-factorizations. In particular, we know essentially nothing about the *typical* properties of one-factorizations for large n. In this view, it is desirable to find rapidly mixing Markov Chains that generate one-factorizations uniformly and efficiently. No such Markov chain is currently known, and here we take a step this direction. We construct a Markov chain whose states are all edge colorings of K_n by $n-1$ colors. This chain is invariant under arbitrary renaming of the vertices or the colors. It reaches a one factorization in polynomial(n) steps from every starting state. In addition, at every transition only $O(n)$ edges change their color. We also raise some related questions and conjectures, and present results of numerical simulations of simpler variants of this Markov chain.

1 Introduction

It is a very old result (e.g. [6]) that for every even integer n, the edges of the complete graph K_n can be properly colored with $n-1$ colors. Such a coloring is called an order-n *one-factorization*, and there is a large body of research dedicated to its study, e.g., [14, 15]. A one-factorization is often viewed as a schedule for the games in a league of n teams. Clearly, every color class in a one-factorization is a perfect matching of $n/2$ edges. Accordingly, we speak of $n-1$ rounds of games in each of which the n teams are paired up to play. If the edge ij is colored k, this means that teams i and j meet at round k. We use the shorthand OF for one-factorization and denote by OF_n the *set* of all order-n one-factorizations.

Many questions about the extremal and typical properties of OFs suggest themselves, but we presently lack the necessary tools to attack such problems. In particular, we seek methods to *generate OFs randomly*, preferably uniformly, or at least with good control over the distribution

*Department of Computer Science, Hebrew University, Jerusalem 91904, Israel. email: maya.dotan1@mail.huji.ac.il .

†Department of Computer Science, Hebrew University, Jerusalem 91904, Israel. email: nati@cs.huji.ac.il .
Supported by ERC grant 339096 "High dimensional combinatorics".

‡Courant Institute, New York University, 251 Mercer Street, New York, NY 10012, USA. email: yuval.peled@courant.nyu.edu .

of the generated OFs. Unfortunately, this goal seems presently out of reach, and the main purpose of this article is to take a step in this direction.

There are several lines of work which are highly relevant. Older literature considers mostly *hill climbing* methods (e.g., [2]) for generating OF's. These studies are experimental in nature and do not suggest any rigorous approach to the problems at hand. A more recent body of work seeks to exactly enumerate and classify OF_n's for specific values of n (e.g., [9]). Although our interest is in the asymptotic behavior of OF_n as $n \to \infty$, many similar ideas and difficulties arise in both contexts, see, e.g., [8].

Furthermore, a one-factorization can also be viewed as a symmetric Latin square. There is an S_n-invariant Markov chain due to Jacobson and Matthews [7] that generates uniformly all order-n Latin squares. It is simple to modify this chain from Latin squares to OFs, but it is not even known whether the modified chain is connected (See [3]).

In recent years there has been remarkable progress in the theory of combinatorial designs, starting from the constructions of Keevash and of Glock, Kühn, Lo, and Osthus [10, 4]. For example, we now have the asymptotic estimate

$$|\mathrm{OF}_n| = \left((1+o(1)) \cdot \frac{n}{e^2}\right)^{(n^2/2)}.$$

The lower bound appears e.g., in [5], and the upper bound is from [13]. This progress has also led to non-trivial results about uniformly random Latin squares [11], and it is conceivable that a similar approach can yield interesting insights about random OFs as well.

In this paper we suggest a different approach to the study of large OFs. We investigate Markov chains that are random walks on graphs whose vertex set, $\mathcal{V}_n$, is comprised of all edge colorings of K_n by $n-1$ colors. Our shorthand for a Markov Chain is MC, and the relevant background on MC's, can be found, e.g., in [12]. Ideally, the MC should (i) Be rapidly mixing, (ii) Have limit distribution that is uniform on OF_n (iii) The limit probability of OF_n is at least $1/\mathrm{poly}(n)$ (iv) It should be possible to efficiently sample a transition.

A natural graph to consider over $\mathcal{V}_n$ is $\mathcal{L}_n = (\mathcal{V}_\mathrm{n}, \mathcal{E}_\mathrm{n})$, in which two edge colorings are adjacent if they differ on exactly one edge of K_n. Our walks on $\mathcal{L}_n$ are conducted with the help of a *potential function* $\Psi : \mathcal{V}_\mathrm{n} \to \mathbb{R}_+$ that vanishes on OF_n and only there. This potential function quantifies how far the present state is from being a OF. Concretely, the potential $\Psi(C)$ of $C \in \mathcal{V}_\mathrm{n}$ is defined to be *the number of pairs of incident edges in K_n that are equally colored* in C. We consider three classes of random walks: *strict, mild* and *weak*. A walk in each of the three classes proceeds in *turns*. Every turn starts by selecting a random neighbor C' of the current state C. Whether we move to C' or stay at C is determined differently in the three cases:

- *Strict:* Move to C' if and only if $\Psi(C) > \Psi(C')$.
- *Mild:* Same, with $\Psi(C) \geq \Psi(C')$.
- *Weak:* If $\Psi(C) \geq \Psi(C')$ we move to C', but some stochastic transition rule allows for occasional moves at which Ψ increases.

In what follows we present a *theorem* about strict random walks, a *conjecture* about mild random walks and a *problem* about the weak ones.

1.1 The strict random walk

We first consider the strict random walk. Such a walk on $\mathcal{L}_n$ can end up in a state that is a local minimum of Ψ but is not a OF. In order to reach a OF, we consider a supergraph $\mathcal{G}_n$ of $\mathcal{L}_n$ over the same vertex set. In addition to the edges in $\mathcal{E}_\mathrm{n}$ there is a second type of *directed* edges that correspond to *two-vertex rotation* steps that we define in Section 2. Our main result is that the strict random walk on $\mathcal{G}_n$ converges rapidly to a OF.

Theorem 1.1. *There is a strict walk on $\mathcal{G}_\mathrm{n}$ that arrives from every starting point to a one-factorization in $O(n^3)$ steps. The graph $\mathcal{G}_\mathrm{n}$ has the following properties:*

1. Symmetry: *It is invariant under permuting of names of K_n's vertices and the names of the $n-1$ colors.*
2. Locality: *For every possible transition C, C', the two edge-colorings C, C' of K_n differ on only $O(n)$ edges of K_n.*
3. Efficiency: *A random step in $\mathcal{G}_n$ can be sampled in time poly(n).*

Remark 1.2. *The symmetry, locality and efficiency properties distinguish the random walk on $\mathcal{G}_n$ from certain trivial and pointless Markov chains. For example, suppose that we are willing to accept a MC that is guaranteed to generate an OF quickly, provided that every OF_n is a possible outcome of the process. Here is a trivial method that accomplishes that goal: First consider a random (proper or improper) coloring of the edges with $n-1$ colors. If this happens to be a OF, we accept it. Otherwise, we walk directly, step by step, to some specific OF. Needless to say, this is far from being symmetric, and the limiting distribution is very far from being uniform. There are other similar trivial Markov chains of OFs that are ruled out by the requirements of locality and efficiency MC.*

Remark 1.3. *The bound of $O(n^3)$ on the number of steps the strict walk takes is self-evident. Indeed, the potential function Ψ takes integer values and is bounded by $O(n^3)$.*

1.2 The mild random walk

Unlike the case of the strict random walk on $\mathcal{L}_n$, all we know about the mild random walk is based on numerical experiments that we present in Section 3. Our simulations strongly supports the following conjecture.

Conjecture 1.4. *The mild random walk on $\mathcal{L}_\mathrm{n}$ started from a uniformly random starting point asymptotically almost surely reaches a one-factorization in $\tilde{O}(n^4)$ turns.*

While we conjecture that the mild random walk almost surely reaches a OF in $\tilde{O}(n^4)$ steps, we cannot even show that it almost surely *eventually* reaches a OF. In other words, that OF_n is reachable from *every* state of $\mathcal{L}_\mathrm{n}$ via a mild walk.

1.3 The weak random walk

What we ultimately wish for is an efficient method to *uniformly* sample OFs that is transparent enough to reveal the typical structure of large OFs. Unfortunately, there is no reason to believe that either the mild walk on $\mathcal{L}_{\mathrm{n}}$ or the strict walk on $\mathcal{G}_n$ reach a uniformly distributed OF, even when starting from a random vertex in $\mathcal{V}_n$. A sampling mechanism that could presumably work is a Markov chain equipped with a Metropolis filter. For a beautiful exposition of this subject, see [12], Chapter 3. Various approaches suggest themselves, and we illustrate this approach using the Markov chain $\mathcal{M}_{\mathrm{n},\epsilon}$, where $\epsilon > 0$ is some parameter.

The underlying graph of $\mathcal{M}_{\mathrm{n},\epsilon}$ is $\mathcal{L}_n$. In every turn, we pick a random neighbor C' in $\mathcal{L}_n$ of the present state C. We take the step from C to C' with certainty when $\Psi(C') \leq \Psi(C)$, but if $\Psi(C') > \Psi(C)$, we take this step only with probability $\epsilon^{2(\Psi(C')-\Psi(C))}$. The main feature of this walk is that the limit distribution of every state C is proportionate to $\epsilon^{2\Psi(C)}$. In particular, all OF_n are equally likely in the limit.

There are two features that we desire: (i) That $\mathcal{M}_{\mathrm{n},\epsilon}$ mixes rapidly, and (ii) That the limit probability of OF_n is at least $\frac{1}{\mathrm{poly}(n)}$. Whether it is possible to achieve these two goals simultaneously is presently unknown, and we raise the following problem:

Problem 1.5. *Is there a function $\epsilon = \epsilon(n)$ under which the resulting Markov Chain $\mathcal{M}_{\mathrm{n},\epsilon}$ is rapidly mixing and the stationary probability of the one-factorizations is at least $n^{-O(1)}$?*

The remainder of the paper is organized as following. In Section 2 we prove Theorem 1.1. In the subsequent section we present numerical evidence that shed some light on Conjecture 1.4 and Problem 1.5. Finally, in Section 4 we present several open problems.

2 A way out of locally optimal states

Consider the strict random walk on $\mathcal{L}_{\mathrm{n}}$. This walk ends up most likely in a sink, or, what we call a *locally optimal* coloring, namely a state that is a local minimum of Ψ in $\mathcal{L}_n$ which is not a OF. Here is our main result on strict walks:

Theorem 2.1. *If C is a locally optimal coloring of $E(K_n)$ that is not a one-factorization, then there are two vertices $x, y \in [n]$ and a recoloring of the edges incident with x, y such that the resulting coloring C' satisfies $\Psi(C) > \Psi(C')$. Furthermore, the recoloring is invariant under permuting the names of the vertices and colors, and it is possible to find the vertices x, y and the said recoloring in time poly(n).*

The recolorings $C \to C'$ of a locally optimal coloring C that we describe in the proof of Theorem 2.1 are called *two-vertex rotations*. The graph $\mathcal{G}_n$ contains the edges of $\mathcal{L}_n$ and the directed edges

$$\{(C, C') : C \text{ is locally optimal, } C \to C' \text{ is a two-vertex rotation}\}.$$

Thus, Theorem 1.1 follows directly from Theorem 2.1.

The first step in proving Theorem 2.1 is to show, in the next lemma, that locally optimal colorings have a quite restricted structure. In such a coloring every monochromatic connected component is either a single edge or a path with two edges. We refer to such a coloring as an *IV coloring*. We call a two-edge path a *Vee*. A Vee has a *center* and two *ends*. Unless otherwise stated, when we speak of a Vee, we implicitly assume that it is *monochromatic*. A Vee whose two edges are colored α is denoted by Vee_α. Clearly, some colors are missing at the Vee's center vertex and if β is such a missing color we may refer to it as Vee_α^β. (Note that β need not be uniquely defined. There may be several colors missing at a Vee's center). We also use the notation Vee^β to indicate only that color β is missing at the Vee's center.

It is convenient to have an alternative description of our potential function. If $a_{u,\mu}$ is the number of μ-colored edges that are incident with u in some edge-coloring C of K_n, then $\Psi(C) = \sum_{u,\mu} \binom{a_{u,\mu}}{2} = \frac{1}{2}\sum_{u,\mu} a_{u,\mu}^2 - \binom{n}{2} = \frac{\Phi(C)}{2} - \binom{n}{2}$, where $\Phi(C) = \sum_{u,\mu} a_{u,\mu}^2$. Consequently it is immaterial whether we use Φ or Ψ as our potential function and we freely switch between the two. We also use the notation $\Phi(u) = \sum_\mu \binom{a_{u,\mu}}{2}$.

Here is the description of a locally-optimal coloring.

Lemma 2.2. *In a locally optimal coloring every color class is the vertex-disjoint union of edges and Vees.*

Proof. Suppose that some edge e of color i is incident to more than one other edge of the same color. There are $2n-4$ edges incident to e and $n-1$ colors. Thus, some color j appears at most once among these edges, and recoloring e by j decreases the number of monochromatic Vees. □

The number of Vees in a locally optimal coloring C is $\Psi(C)$. Therefore reducing Φ for a locally optimal coloring, is synonymous with a reduction in its number of Vees. Let us consider an IV coloring C that is not an OF_n. Since C is not an OF, it contains some monochromatic Vee, say Vee_α. But n is even, so if C has exactly one Vee_α, then necessarily there must be a vertex that is incident to no α-colored edges. Such a vertex must, in turn, be the center of some Vee, say a Vee_γ for some $\gamma \neq \alpha$. To sum up, if C is IV but not an OF_n, and if C has a monochromatic Vee_α, then C must

1. Contain a Vee_γ^α for some $\gamma \neq \alpha$. (See Figure: 1), or
2. Contain an additional Vee_α (See Figure: 2)

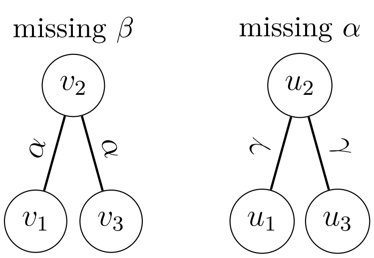

Figure 1: There is a Vee_γ^α in C

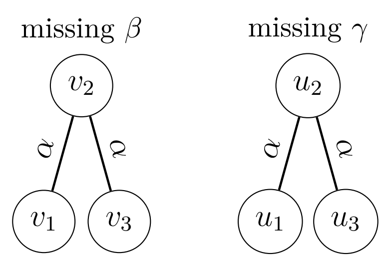

Figure 2: C contains an additional Vee_α

Fix an IV coloring C that is not an OF_n. As it turns out, one of the following happens now. Either we can find a two-vertex recoloring step that reduces Φ, or we can recolor an edge in a way that does not change Φ and then do a two-vertex recoloring step that reduces the potential. In our search for a two-vertex recoloring we actually restrict ourselves to a *flip*, a special kind of a two-vertex recoloring step as we now describe. Let u and v be the two vertices whose edges we intend to recolor. In a (u, v)-flip, for every vertex $w \neq u, v$ either the edges wu and wv retain their original colors or they exchange their colors between them (See Figure: 3). We turn to discuss next such steps that decrease Φ. It is interesting to note that the very same step has been used in [8] toward the exhaustive listing and enumeration of OF_n for small values of n.

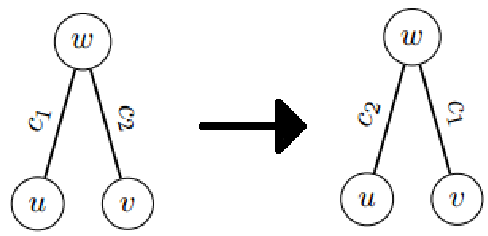

Figure 3: A (u, v)-flip as seen from vertex w

Lemma 2.3. *If $|a_{u,\lambda} - a_{v,\lambda}| \geq 2$ for some color λ, then there is a (u, v)-flip that decreases Φ.*

Proof. We use, as above, the notation $x_\mu := a_{u,\mu}$ and $y_\mu := a_{v,\mu}$. Recall that $\Phi(C) = \sum_\nu \Phi(\nu) = \sum_{\nu,\mu} (a_{\nu,\mu}(C))^2$. For every $w \neq u, v$, the term $\Phi(w)$ does not change as a result of a (u, v)-flip. Let x'_μ, y'_μ be the values of $a_{u,\mu}$ resp. $a_{v,\mu}$ after recoloring. The lemma claims that there is a (u, v)-flip for which

$$\sum x_\mu^2 + y_\mu^2 > \sum (x'_\mu)^2 + (y'_\mu)^2.$$

This inequality can be equivalently stated as

$$\sum (x_\mu + y_\mu)^2 + (x_\mu - y_\mu)^2 > \sum (x'_\mu + y'_\mu)^2 + (x'_\mu - y'_\mu)^2.$$

But it is clear that for every μ there holds $x_\mu + y_\mu = x'_\mu + y'_\mu$ (since a flip preserves the number of edges of any given color that are incident with u or v) and so, the lemma would follow if we could prove

Claim 2.4. *Assuming that $|x_\lambda - y_\lambda| \geq 2$ for some color λ, there is a (u, v)-flip such that $|x_\mu - y_\mu| \geq |x'_\mu - y'_\mu|$ holds for every color μ and at least one of these inequalities is strict.*

Claim 2.4 is a consequence of the following simple variation on Euler's Theorem:

Claim 2.5. *Every multigraph H (possibly with loops) can be oriented in such a way that for every vertex the indegree and outdegree differ at most by one. Such an orientation can be found in polynomial time.*

Proof. As long as the set of edges of H that has not yet been oriented contains a cycle C, we orient the edges of C cyclically. This yields an equal indegree and outdegree for every vertex. Afterwards, the set of edges of H that has not been oriented forms an acyclic graph. As long as this acyclic graph is not empty, we select a simple path between two leaves and orient its edges along it. This yields an equal indegree and outdegree for every vertex except a difference of one for the two leaves. The proof is concluded since every vertex can appear as a leaf in this process at most once. □

We introduce next the directed multigraph H with vertex set $\{1, \ldots, n-1\}$ whose vertices correspond to colors. Associated with every vertex $w \neq u, v$ in K_n is a directed edge in H. If in the original edge-coloring of K_n the edges wu, wv are colored α and β, then the edge in H that corresponds to w goes from α to β. Consider a reorientation of H as given by Claim 2.5. Accordingly, if the directed edge $\alpha \to \beta$ gets reversed, then we switch colors between the edges wu and wv in K_n, that are presently recolored β and α respectively.

We can now complete the proof of Lemma 2.3. Let us interpret the conclusion of Claim 2.5 in the language of edge coloring of K_n. The conclusion says that $|x'_\mu - y'_\mu| \leq 1$ for all colors μ. However, as mentioned, there also holds $x_\mu + y_\mu = x'_\mu + y'_\mu$, so that $|x_\mu - y_\mu| \geq |x'_\mu - y'_\mu|$ for all colors μ. Also, by assumption, $|x_\lambda - y_\lambda| \geq 2$, so that for $\mu = \lambda$ the inequality is strict. The claim follows. □

Flips thus deal successfully with the case of an IV coloring C that has a Vee_α as well as a Vee^α. Just apply Lemma 2.3 with u, v the centers of these Vees and $\lambda = \alpha$.

There is one last remaining case to consider. Namely, an IV coloring C that is not an OF_n, which contains, for no λ, both a Vee_λ and a Vee^λ. By the parity argument mentioned above, it must have at least two Vee_α's, say a $\text{Vee}_\alpha^\beta = \{v_1, v_2, v_3\}$ centered at v_2, and a Vee_α^γ. However, in this last remaining case there are no Vee^α, Vee_β or Vee_γ. (We do not care whether $\beta = \gamma$ or not). The situation can be resolved by either a single edge recoloring or a single edge recoloring followed by a flip. Let us recolor the edge v_1v_2 from α to β. This clearly decreases $\Phi(v_2)$ by 2, and note that $\Phi(v_1)$ increases by at most 2, since there are no Vee_β's in C. This single edge change did not change $\Phi(C)$ since it cannot decrease, and this means that there is now a Vee_β^α centered at v_1. Together with the Vee_α^γ mentioned above, we are now in a position to apply a flip that reduces Φ. This completes the description of the algorithm and the proof of Theorem 2.1.

3 Numerical simulations

In this Section we discuss numerical simulations that we carried out with the Markov chain $\mathcal{M}_{\mathrm{n},\epsilon}$ as defined in Section 1. As formulated in Problem 1.5, we wish to find values of $\epsilon = \epsilon(n)$ for which the chain mixes rapidly and the limit probability of OFs is not too small.

In order to see why these two properties are in conflict, it is helpful to consider what happens for extreme values of ϵ. When $\epsilon = 1$, we are taking a walk on a discrete cube $\{1, \ldots, n-1\}^{\binom{n}{2}}$. As is well-known, this walk mixes in time polynomial(n), and the limit distribution is uniform. However, the OF_n occupy only a tiny fraction of $\exp(-\Omega(n^2))$ of the space.

At the other extreme, when $\epsilon = 0$, this walk coincides with the mild walk on $\mathcal{L}_\mathrm{n}$ with potential function Ψ as described in Section 1. Now the walk stops when it reaches a sink, so there is no mixing to speak of. As mentioned, we do not know that OFs are the only sinks and, if not, whether the mild walk reaches a OF with considerable probability. Our simulations suggest that starting from a random coloring, this walk does converge to a OF after $\tilde{O}(n^4)$ turns (See Figure 4), but there is no reason to believe that the induced distribution on the OFs is close to uniform.

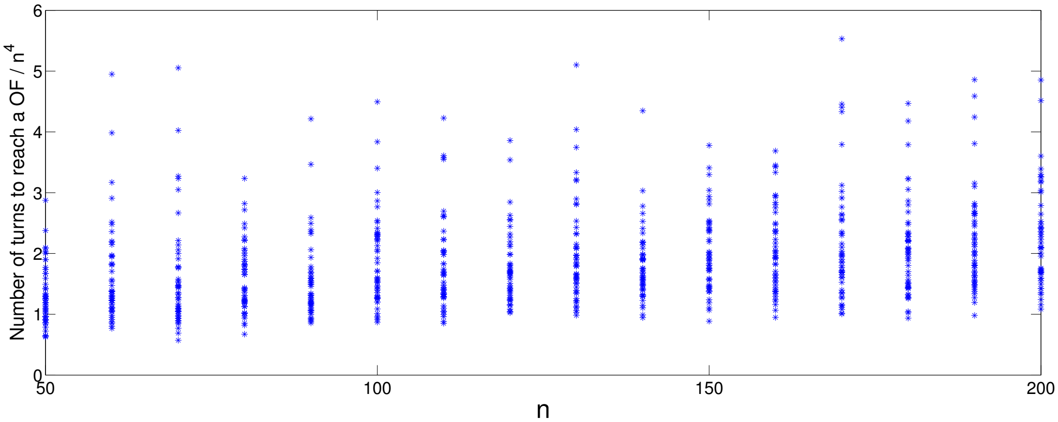

Figure 4: The number of turns in a mild walk on $\mathcal{L}_{\mathrm{n}}$ with respect to Ψ from a random coloring to a OF, normalized by n^4. There are 50 runs for each $n \in \{50, 60, ..., 200\}$.

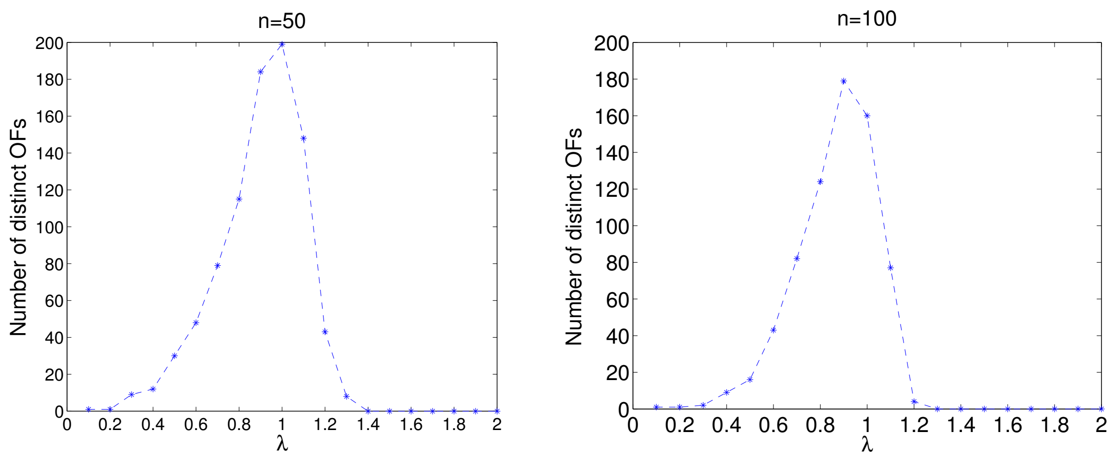

Figure 5: The number of distinct OFs visited by a walk of $m = 1000 \cdot n^4$ turns on $\mathcal{M}_{\mathrm{n},\lambda/\mathrm{n}}$ starting from a random coloring, as a function of $\lambda \in \{0.1, 0.2, ..., 2\}$. Left: $n = 50$. Right: $n = 100$.

Roughly speaking, Problem 1.5 asks whether there exists an $1 > \epsilon(n) > 0$ that is both small enough so that the walk on $\mathcal{M}_{\mathrm{n},\epsilon}$ spends a considerable fraction of its time in one-factorizations, and large enough to guarantee rapid mixing. In order to detect the values of $\epsilon(n)$ in which this is likely to occur, we counted, for a variety of values of ϵ, the number of *distinct* OFs visited by a walk on $\mathcal{M}_{\mathrm{n},\epsilon}$ of some predetermined length. Experimental evidence indicate that this number is maximized around $\epsilon = 1/n$ (See Figure 5).

Our experiments suggest that in $\mathcal{M}_{n,1/n}$ a new OF is reached every $\tilde{O}(n^4)$ turns (where $\tilde{O}(n^2)$ of them are not-lazy). It is hard to test numerically whether the chain $\mathcal{M}_{\mathrm{n},1/\mathrm{n}}$ is rapidly mixing, but we did observe the following supporting evidence. Consider a segment of a walk on $\mathcal{M}_{\mathrm{n},1/\mathrm{n}}$ beginning with a one factorization $C \in \mathrm{OF}_n$, proceeding through non-OFs until for the first time a OF $C' \in \mathrm{OF}_n$ is reached. Let t be the number of steps *taken* (i.e., state changes) en route from C to C'. As Figure 6 shows, the fraction of edges on which C and C' agree tends to decay like $\exp\left(-\Theta(t/\binom{n}{2})\right)$. Consider for comparison, a process where we generate

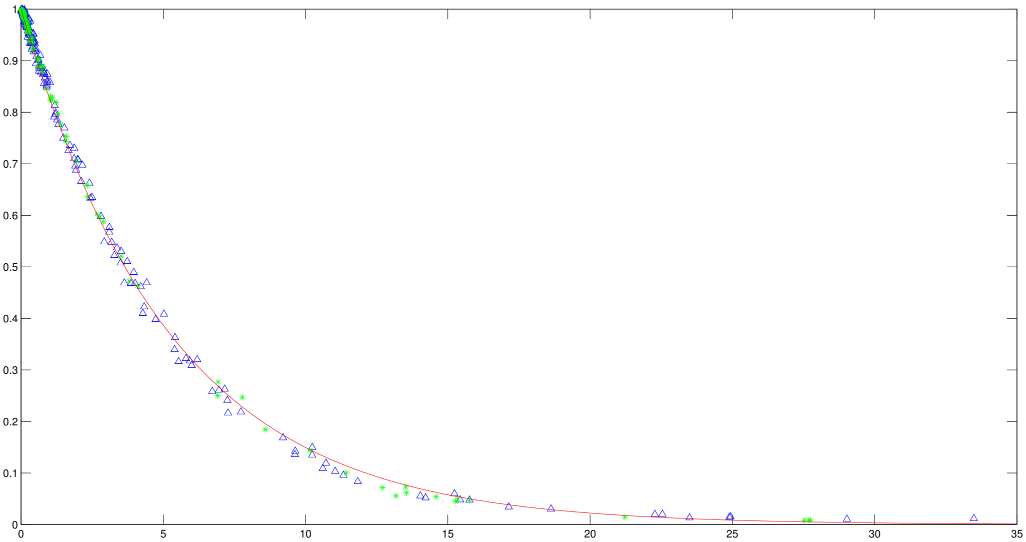

Figure 6: Consider two consecutively visited OFs in $\mathcal{M}_{\mathrm{n},1/\mathrm{n}}$. This plot shows the fraction of edges on which the two OFs agree as a function of *the number of steps taken* between the two visits (normalized by $\binom{n}{2}$) , for $n = 100$ (triangles) and $n = 200$ (stars). The red line shows the function $t \mapsto 0.82^t$.

C' in t steps in each of which we randomly recolor an edge in C that is chosen independently, uniformly with repetition. In this model the agreement between C and C' would, with high probability, be roughly $\exp\left(-t/\binom{n}{2}\right)$. This may be indicative of rapid mixing.

4 Open Problems

- Conjecture 1.4 raises several questions about the structure of $\mathcal{L}_n$. Are there non-OFs sinks in the mild walk on $\mathcal{L}_n$? Even if the answer is negative, we conjecture that there is a path from every state in $\mathcal{L}_n$ to a OF in which Ψ never increases. It is also interesting to understand the distribution on OF_n that the mild walk induces on $\mathcal{L}_n$.

- We are hoping that in the long time limit the chain $\mathcal{M}_{n,1/n}$ spends considerable time, i.e., $\frac{1}{\mathrm{poly}(n)}$ in OF_n. A proof of this would require a fairly precise understanding of the way in which Ψ is distributed on the vertices of $\mathcal{L}_\mathrm{n}$. For every integer k, let $L_{n,k}$ be the set of vertices $\{C \in \mathcal{L}_\mathrm{n} \ : \ \Psi(\mathrm{C}) = \mathrm{k}\}$. The limit probability of $L_{n,k}$ in $\mathcal{M}_{n,1/n}$ is proportional to $|L_{n,k}| \cdot (1/n)^{2k}$. We know that the OF_n are characterized by the condition $\Psi = 0$ and that their number is $|\mathrm{OF}_n| = ((1+o(1))\frac{n}{e^2})^{n^2/2}$. It would be needed to tighten these asymptotic estimates, and get similar tight estimates for $L_{n,k}$ in the range $0 < k \leq O(n^2/\log n)$.

- The previous question raises a different problem concerning the distribution of $L_{n,k}$ for $k = \Theta(n^2)$. This problem is best formulated in the terminology of large deviations theory:

Problem 4.1. *Let C be a random vertex of $\mathcal{L}_{\mathrm{n}}$, and $0 \le x \le 1/2$ an integer. Show that the sequence $\frac{2}{n^2} \log P[\Psi(C) \le xn^2]$ converges and estimate the limit.*

The values of this rate function at $x = 0$ and $x = 1/2$ are known. It is 0 at $x = 1/2$, since the average of Ψ over all the vertices of $\mathcal{L}_{\mathrm{n}}$ is $\frac{n(n-2)}{2}$ and Ψ is concentrated about the mean, i.e., $|L_{n,n^2/2}| = ((1 - o(1))n)^{n^2/2}$. In addition, the value of the rate function at $x = 0$ is -2 by the asymptotic enumeration of OF_n.

To the best of our knowledge, this is a new type of large deviation problem, but there may be hope that new advances in large deviations theory (see the survey [1]) can help resolve it.

- Other combinatorial designs such as Latin Squares, Steiner Triple Systems and high-dimensional permutations can presumably be generated by Markov Chains of a similar flavor. This immediately raises analogous questions concerning the corresponding strict, mild and weak random walks, as well as additional problems in large deviations.

Acknowledgement. The authors would like to thank Amir Dembo and Michael Simkin for useful discussions, and the anonymous referee for their helpful suggestions.

References

[1] Sourav Chatterjee. An introduction to large deviations for random graphs. *Bulletin of the American Mathematical Society*, 53(4):617–642, 2016.

[2] JH Dinitz and DR Stinson. A hill-climbing algorithm for the construction of one-factorizations and room squares. *SIAM Journal on Algebraic Discrete Methods*, 8(3):430–438, 1987.

[3] Andy L Drizen. *Generating Uniformly-Distributed Random Generalised 2-designs with Block Size 3*. PhD thesis, Queen Mary University of London, 2012.

[4] Stefan Glock, Daniela Kühn, Allan Lo, and Deryk Osthus. The existence of designs via iterative absorption. *arXiv preprint arXiv:1611.06827*, 2016.

[5] Stefan Glock, Daniela Kühn, Allan Lo, and Deryk Osthus. On a conjecture of Erdős on locally sparse Steiner triple systems. *arXiv preprint arXiv:1802.04227*, 2018.

[6] Marshall Hall. *Combinatorial theory*, volume 71. John Wiley & Sons, 1967.

[7] Mark T Jacobson and Peter Matthews. Generating uniformly distributed random latin squares. *Journal of Combinatorial Designs*, 4(6):405–437, 1996.

[8] Petteri Kaski, André de Souza Medeiros, Patric RJ Östergård, and Ian M Wanless. Switching in one-factorisations of complete graphs. *The Electronic Journal of Combinatorics*, 21(2):2–49, 2014.

[9] Petteri Kaski and Patric R. J. Östergård. There are $1,132,835,421,602,062,347$ nonisomorphic one-factorizations of K_{14}. *J. Combin. Des.*, 17(2):147–159, 2009.

[10] Peter Keevash. The existence of designs. *arXiv preprint arXiv:1401.3665*, 2014.

[11] Matthew Kwan. Almost all Steiner triple systems have perfect matchings. *arXiv preprint arXiv:1611.02246*, 2016.

[12] David Asher Levin, Yuval Peres, and Elizabeth Lee Wilmer. *Markov chains and mixing times*. American Mathematical Soc., 2009.

[13] Nathan Linial and Zur Luria. An upper bound on the number of Steiner triple systems. *Random Structures & Algorithms*, 43(4):399–406, 2013.

[14] Eric Mendelsohn and Alexander Rosa. One-factorizations of the complete graph—a survey. *Journal of Graph Theory*, 9(1):43–65, 1985.

[15] Walter D Wallis. *One-factorizations*, volume 390. Springer Science & Business Media, 2013.