

Cycle-factors of regular graphs via entropy

Micha Christoph[†] Nemanja Draganić[‡] António Girão[‡] Eoin Hurley[‡]
Lukas Michel[‡] Alp Müyesser[‡]

28 July 2025

Abstract

It is a classical result that a random permutation of n elements has, on average, about $\log n$ cycles. We generalise this fact to all directed d -regular graphs on n vertices by showing that, on average, a random cycle-factor of such a graph has $\mathcal{O}((n \log d)/d)$ cycles. This is tight up to the constant factor and improves the best previous bound of the form $\mathcal{O}(n/\sqrt{\log d})$ due to Vishnoi. Our results also yield randomised polynomial-time algorithms for finding such a cycle-factor and for finding a tour of length $(1 + \mathcal{O}((\log d)/d)) \cdot n$ if the graph is connected. This makes progress on a conjecture of Magnan and Martin and on a problem studied by Vishnoi and by Feige, Ravi, and Singh. Our proof uses the language of entropy to exploit the fact that the upper and lower bounds on the number of perfect matchings in regular bipartite graphs are extremely close.

1 Introduction

The probability that a random permutation has a fixed point was considered as early as 1704 in Pierre-Rémond de Montmort's *Essai d'Analyse sur les Jeux de Hasard*. Random permutations have since become one of the most classical topics in discrete probability. Cauchy was the first to study permutations for their own sake. In the 1840s, he derived a formula for the number of permutations of n elements with a given cycle structure [Cau40]. This formula implies that, on average, a random permutation contains $\sum_{k=1}^n 1/k \simeq \log n$ cycles. The latter result is now a standard exercise in discrete probability.

In this paper, we generalise this result to all directed d -regular graphs¹ on n vertices. A *cycle-factor* of such a graph is a partition of its vertices into directed cycles. Equivalently, this corresponds to a permutation of the vertices such that every vertex is mapped to one of its d out-neighbours. In particular, a random permutation of n elements is the same as a random cycle-factor of the complete directed graph on n vertices with a loop at every vertex.

It follows from Hall's theorem that every directed d -regular graph has at least one cycle-factor. How many cycles are there in a typical cycle-factor of such a graph? By the discussion above, a random cycle-factor of the disjoint union of n/d complete directed graphs on d vertices has, on average, about $(n \log d)/d$ cycles. We prove that, up to constant factors, this is an upper bound on the expected number of cycles in a typical cycle-factor for any directed d -regular graph.

[†]Department of Mathematics, ETH Zürich, Switzerland (micha.christoph@math.ethz.ch). Research supported by SNSF Ambizione Grant No. 216071.

[‡]Mathematical Institute, University of Oxford, United Kingdom ({nemanja.draganic,antonio.girao,eoin.hurley,lukas.michel,alp.muyesser}@maths.ox.ac.uk). Research of Nemanja Draganić supported by SNSF project 217926. Research of Eoin Hurley supported by ERC Advanced Grant 883810.

¹A directed graph can have loops and directed cycles of length 2, but no parallel edges. A directed graph is *d*-regular if the in-degree and out-degree of every vertex is exactly d .

Theorem 1.1. *The expected number of cycles in a uniformly random cycle-factor of a directed d -regular graph on n vertices is $\mathcal{O}((n \log d)/d)$.*

This recovers the aforementioned fact that, on average, uniformly random permutations of n elements have $\mathcal{O}(\log n)$ cycles. The constant factor in Theorem 1.1 can be taken to be 4. Note that Theorem 1.1 also applies to undirected d -regular graphs if we direct each edge in both directions². Previously, the best upper bound known on the number of cycles in a typical cycle-factor of a d -regular graph was $\mathcal{O}(n/\sqrt{\log d})$ due to Vishnoi [Vis12].

While Theorem 1.1 only proves the existence of cycle-factors with few cycles, we can also find such cycle-factors efficiently. This is a consequence of the fact that we can sample almost-uniform random perfect matchings of bipartite graphs in polynomial time [JSV04].

Corollary 1.2. *There exists a randomised polynomial-time algorithm that computes, with high probability, a cycle-factor of a directed d -regular graph on n vertices with at most $\mathcal{O}((n \log d)/d)$ cycles.*

If we remove one edge from each cycle in a cycle-factor, we obtain a *path-factor*, which is a partition of the vertices of a graph into paths. Magnan and Martin [MM09] conjectured that every d -regular graph on n vertices has a path-factor with at most $n/(d+1)$ paths. This would be tight as the disjoint union of $n/(d+1)$ cliques on $d+1$ vertices requires at least $n/(d+1)$ paths in any path-factor. The conjecture of Magnan and Martin is only known to be true for $d \leq 6$ [MM09, FF22], $d = \Omega(n)$ [GL21], and for partitioning almost all vertices of a graph into paths [MMP24, LMS25]. Feige, Ravi, and Singh [FRS14] showed that every d -regular graph on n vertices has a path-factor with $\mathcal{O}(n/\sqrt{d})$ paths, which was the best previous upper bound known. Using Corollary 1.2, we get the following improvement.

Corollary 1.3. *There exists a randomised polynomial-time algorithm that computes, with high probability, a path-factor of a d -regular graph on n vertices with at most $\mathcal{O}((n \log d)/d)$ paths.*

The conjecture of Magnan and Martin is also intimately related to the celebrated Linear Arboricity Conjecture of Akiyama, Exoo, and Harary [AEH80] which states that the edges of every graph with maximum degree Δ can be partitioned into at most $\lceil (\Delta+1)/2 \rceil$ path-factors. The best bound known for the Linear Arboricity Conjecture is due to Lang and Postle [LP23] and shows that $\Delta/2 + \mathcal{O}(\sqrt{\Delta} \log^4 \Delta)$ path-factors suffice. This bound cannot be improved to $\Delta/2 + o(\sqrt{\Delta})$ without being able to show that every Δ -regular graph has a path-factor with at most $o(n/\sqrt{\Delta})$ paths. Corollary 1.3 shows that such path-factors exist.

Finally, we can use our results to find short tours³ of connected d -regular graphs. The problem of finding such tours was introduced by Vishnoi [Vis12], motivated by the study of approximation algorithms for the travelling salesman problem. Feige, Ravi, and Singh [FRS14] showed that there exist tours of length $(1 + \mathcal{O}(1/\sqrt{d})) \cdot n$ and provided an example in which every tour has length at least $(1 + \Omega(1/d)) \cdot n$. Since a cycle-factor with c cycles in a connected graph can be efficiently transformed into a tour of length $n + 2c$, we obtain the following.

Corollary 1.4. *There exists a randomised polynomial-time algorithm that computes, with high probability, a tour of a connected d -regular graph on n vertices of length at most $(1 + \mathcal{O}((\log d)/d)) \cdot n$.*

The rest of the paper is structured as follows. In Section 3, we prove our main results, Theorem 1.1 and Corollary 1.2. This section also includes an overview of our proof strategy. Since the proof of Theorem 1.1 uses the language of entropy, we first provide a short introduction to entropy in Section 2. We also use this section to introduce a small amount of notation.

Note. *In upcoming work, we use a completely different approach to prove that every d -regular graph on n vertices has a path-factor with $\mathcal{O}(n/d)$ paths as well as a tour of length $(1 + \mathcal{O}(1/d)) \cdot n$ if the graph*

²A cycle-factor of an undirected graph can contain individual edges, which we consider to be cycles of length 2.

³A tour of a graph is a connected walk that visits every vertex and starts and ends at the same vertex.

is connected. We also use this approach to make progress on the Linear Arboricity Conjecture. However, this approach does not provide efficient algorithms and does not apply to directed graphs or cycle-factors.

2 Preliminaries

Let $[n] := \{1, \dots, n\}$, and let S_n denote the family of permutations of $[n]$. If $\sigma \in S_n$ and $\mathcal{I} \subseteq [n]$, we write $\sigma(\mathcal{I}) = \{\sigma(i) : i \in \mathcal{I}\}$. A cycle of $\sigma \in S_n$ is a minimal non-empty subset $\mathcal{I} \subseteq [n]$ such that $\sigma(\mathcal{I}) = \mathcal{I}$. We denote the number of cycles of σ by $|\sigma|$.

For our proofs, we will need Stirling's approximation which states that as $n \rightarrow \infty$, we have $n! = (1 + o(1)) \cdot \sqrt{2\pi n} (n/e)^n$. In fact, it will be more useful for us to use the weaker bounds $(n/e)^n \leq n! \leq en(n/e)^n$ which hold for all n .

We now recall some basic definitions and properties of Shannon entropy which will be essential for our proofs. For a more thorough exposition, we refer the reader to [AZ18, Chapter 37]. This also contains a nice exposition of Radhakrishnan's proof [Rad97] of the Brégman-Minc inequality (see Theorem 3.1), which serves as the starting point of our approach. In the following, all logarithms are base 2, and we use the convention $p \log p = 0$ for $p = 0$.

Suppose that X is a random variable that takes values in a finite set \mathcal{X} . Then, the *Shannon entropy* of X is defined as

$$\mathcal{H}(X) := - \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) \log(\mathbb{P}(X = x)).$$

One basic property of entropy is that $\mathcal{H}(X) \leq \log|\mathcal{X}|$ with equality if and only if X is distributed uniformly at random in \mathcal{X} . For $p \in [0, 1]$, write $\mathcal{H}(p) := -p \log p - (1 - p) \log(1 - p)$ for the entropy of a Bernoulli random variable with probability p .

The definition of entropy easily extends to conditional probability. If E is some event, then the entropy of X conditioned on E occurring is

$$\mathcal{H}(X \mid E) := - \sum_{x \in \mathcal{X}} \mathbb{P}(X = x \mid E) \log(\mathbb{P}(X = x \mid E)).$$

If Y is another random variable that takes values in a finite set \mathcal{Y} , then the *conditional entropy* of X given Y is

$$\mathcal{H}(X \mid Y) := \sum_{y \in \mathcal{Y}} \mathbb{P}(Y = y) \cdot \mathcal{H}(X \mid Y = y).$$

Note that $\mathcal{H}(X \mid Y) = 0$ if and only if Y completely determines X .

The only fact about conditional entropy that we will need is the *chain rule* which says that if X_1, \dots, X_n are random variables, then

$$\mathcal{H}(X_1, \dots, X_n) = \mathcal{H}(X_1) + \mathcal{H}(X_2 \mid X_1) + \dots + \mathcal{H}(X_n \mid X_1, \dots, X_{n-1})$$

where the tuple X_1, \dots, X_i is viewed as a single random variable whose distribution is the joint distribution of X_1, \dots, X_i .

3 Expected number of cycles in a random cycle-factor

Suppose that G is a directed d -regular graph on n vertices. We wish to show that the expected number of cycles of a uniformly random cycle-factor of G is $\mathcal{O}((n \log d)/d)$. To achieve this, we will reveal the edges of the random cycle-factor one by one in a random order and bound the probability that each new revealed edge closes a cycle. We do this by going through the vertices

in a random order and revealing their out-neighbours in the cycle-factor. If the revealed out-neighbour of a vertex were always uniformly distributed among all remaining available out-neighbours, then we could show the desired bound. Indeed, since the edges are revealed in a random order, it is possible to show that the number t of remaining out-neighbours of a vertex is uniformly distributed in $[d]$, and so the probability that each new revealed edge closes a cycle would be at most

$$\frac{1}{d} \sum_{t=1}^d \frac{1}{t} = \mathcal{O}\left(\frac{\log d}{d}\right).$$

To handle the general case where out-neighbours are not uniformly distributed, we connect our process to Radhakrishnan's entropy proof [Rad97] of the Brègman-Minc inequality [Brè73]. Towards that end, construct a d -regular bipartite graph H with bipartition $V(H) = U \cup V$ where U and V are both copies of $V(G)$, and for each directed edge (u, v) of G we add an edge between $u \in U$ and $v \in V$ to H . It is easy to see that cycle-factors of G correspond to perfect matchings of H ⁴. The Brègman-Minc inequality gives the following upper bound on the number of perfect matchings in regular bipartite graphs.

Theorem 3.1. *Every d -regular bipartite graph on $2n$ vertices has at most $(d!)^{n/d}$ perfect matchings.*

On a high level, Radhakrishnan's proof of Theorem 3.1 picks a uniformly random perfect matching and shows that its entropy is at most $\log((d!)^{n/d})$. The idea here is to reveal the edges of the random matching one by one in a random order⁵ and to bound the amount of entropy added by each new revealed edge. Note that the amount of entropy added is maximised if each new revealed edge is uniformly distributed among its remaining options. This can be used to prove Theorem 3.1.

We pursue a very similar strategy to analyse random cycle-factors of G . Conveniently, revealing an edge of a random perfect matching of the auxiliary bipartite graph H corresponds precisely to revealing an edge of the random cycle-factor of G . As mentioned before, we wish to understand the case where a revealed edge of the cycle-factor is not uniformly distributed among its remaining options. If this happens, the amount of entropy added for the corresponding edge in H is lower than the entropy of the uniform distribution. This, in turn, translates into a better upper bound on the number of perfect matchings of H . However, this upper bound cannot drop too low. This is because the Van der Waerden inequality, proved by Egor'chev [Ego81] and Falikman [Fal81], gives the following lower bound on the number of perfect matchings in regular bipartite graphs.

Theorem 3.2. *Every d -regular bipartite graph on $2n$ vertices has at least $n! \cdot d^n / n^n$ perfect matchings.*

This lower bound allows us to quantify how much entropy can be lost while revealing the edges of the random cycle-factor of G . We will see that the total loss in entropy is at most $\mathcal{O}((n \log d)/d)$ compared to what the entropy would have been had the edges been uniformly distributed throughout the process. The key fact that ties everything together is that any loss in entropy increases the probability of closing a cycle at most by an additive linear term. Together, this implies Theorem 1.1.

We start by proving the following simple lemma that formalises the relationship between a loss in entropy and an increase in the probability of a random variable attaining a specific value. In our case, this will correspond to the probability of closing a cycle.

⁴This observation was also used by Vishnoi in [Vis12] to show that a typical cycle-factor in a d -regular graph has $\mathcal{O}(n/\sqrt{\log d})$ cycles.

⁵To the best of our knowledge, the random order was first used by Alon and Spencer [AS16] to prove the Brègman-Minc inequality.

Lemma 3.3. Let X be a random variable taking values in a set \mathcal{X} of size s , and let $\ell := \log s - \mathcal{H}(X)$. Then, for all $x \in \mathcal{X}$ it holds that

$$\mathbb{P}(X = x) \leq \frac{2}{s} + \ell.$$

Proof. Let $x \in \mathcal{X}$ and let $p := \mathbb{P}(X = x)$. Consider the indicator random variable $\mathbb{1}_{X=x}$ of the event that $X = x$. By the chain rule for entropy,

$$\mathcal{H}(X) = \mathcal{H}(\mathbb{1}_{X=x}) + \mathcal{H}(X \mid \mathbb{1}_{X=x}) = \mathcal{H}(p) + p \cdot \mathcal{H}(X \mid X = x) + (1 - p) \cdot \mathcal{H}(X \mid X \neq x).$$

Clearly, $\mathcal{H}(X \mid X = x) = 0$. Also, if $X \neq x$, then X takes at most $s - 1$ distinct values, and so $\mathcal{H}(X \mid X \neq x) \leq \log(s - 1)$. This implies that

$$\mathcal{H}(X) \leq \mathcal{H}(p) + (1 - p) \log(s - 1).$$

Let $h(q) := \mathcal{H}(q) + (1 - q) \log(s - 1)$. Note that $h(q) = \mathcal{H}(Y)$ for the random variable Y with $\mathbb{P}(Y = x) = q$ and $\mathbb{P}(Y = y) = (1 - q)/(s - 1)$ for all $y \in \mathcal{X} \setminus \{x\}$, and so $h(q) \leq \log s$. Moreover,

$$h'(q) = -\log q + \log(1 - q) - \log(s - 1) = \log\left(\frac{1 - q}{q(s - 1)}\right).$$

Since this is decreasing in q , we get for $q \geq 2/s$ that

$$h'(q) \leq h'\left(\frac{2}{s}\right) = \log\left(\frac{s - 2}{2(s - 1)}\right) \leq \log\left(\frac{1}{2}\right) = -1.$$

So, for $q \geq 2/s$, we have $h(q) \leq h(2/s) - (q - 2/s) \leq \log s - (q - 2/s)$. Since we know that $h(p) \geq \mathcal{H}(X) = \log s - \ell$, it follows that $p \leq 2/s + \ell$. \square

We can now prove our main result using the strategy outlined earlier in this section. The first part of the proof closely follows Radhakrishnan's proof of the Brègman-Minc inequality, with the twist being that we keep track of the total loss in entropy. This is used in the second part to prove an upper bound on the expected number of cycles, by passing through [Lemma 3.3](#).

Proof of Theorem 1.1. We identify the vertices of G with $[n]$. Let \mathcal{C} be the set of all permutations σ of $[n]$ such that $(i, \sigma(i)) \in E(G)$ for all $i \in [n]$. Note that \mathcal{C} corresponds to the collection of all cycle-factors of G . [Theorem 3.2](#) in combination with Stirling's approximation implies that $|\mathcal{C}| \geq n! \cdot d^n / n^n \geq (d/e)^n$.

Pick $\sigma \in \mathcal{C}$ uniformly at random, so $\mathcal{H}(\sigma) = \mathcal{H}(\sigma(1), \dots, \sigma(n)) = \log|\mathcal{C}|$. For every permutation $\tau \in S_n$, the chain rule for entropy implies that

$$\mathcal{H}(\sigma) = \sum_{j=1}^n \mathcal{H}(\sigma(\tau(j)) \mid \sigma(\tau([j - 1]))).$$

For $i \in [n]$, let $\tau_{<i} := \tau([j - 1])$ where $j = \tau^{-1}(i)$. If we reorder the above sum, we get

$$\mathcal{H}(\sigma) = \sum_{i=1}^n \mathcal{H}(\sigma(i) \mid \sigma(\tau_{<i})).$$

Expanding the definition of conditional entropy, we get

$$\mathcal{H}(\sigma) = \sum_{i=1}^n \sum_{\pi} \mathbb{P}(\sigma(\tau_{<i}) = \pi) \cdot \mathcal{H}(\sigma(i) \mid \sigma(\tau_{<i}) = \pi),$$

where π ranges over all possible values of $\sigma(\tau_{<i})$. Since σ is chosen uniformly at random, observe that $\mathbb{P}(\sigma(\tau_{<i}) = \pi) = |\{\sigma' \in \mathcal{C} : \sigma'(\tau_{<i}) = \pi\}|/|\mathcal{C}|$. Therefore,

$$\mathcal{H}(\sigma) = \sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} \mathcal{H}(\sigma(i) \mid \sigma(\tau_{<i}) = \sigma'(\tau_{<i})).$$

For all $i \in [n]$, $\sigma' \in \mathcal{C}$, and $\tau \in S_n$, let $s(i, \sigma', \tau) := |N^+(i) \setminus \sigma'(\tau_{<i})|$ be the number of out-neighbours of i in G that are not contained in $\sigma'(\tau_{<i})$. If $\sigma(\tau_{<i}) = \sigma'(\tau_{<i})$, then $\sigma(i)$ must be one of these out-neighbours, and so $\mathcal{H}(\sigma(i) \mid \sigma(\tau_{<i}) = \sigma'(\tau_{<i})) \leq \log s(i, \sigma', \tau)$. Define $\ell(i, \sigma', \tau) := \log s(i, \sigma', \tau) - \mathcal{H}(\sigma(i) \mid \sigma(\tau_{<i}) = \sigma'(\tau_{<i}))$. Then,

$$\mathcal{H}(\sigma) = \sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} (\log s(i, \sigma', \tau) - \ell(i, \sigma', \tau)).$$

Since this holds for all $\tau \in S_n$, we may pick $\tau \in S_n$ uniformly at random. This yields

$$\mathcal{H}(\sigma) = \mathbb{E}_\tau \left(\sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} (\log s(i, \sigma', \tau) - \ell(i, \sigma', \tau)) \right) = \sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} \mathbb{E}_\tau (\log s(i, \sigma', \tau)) - \mathcal{L}$$

where

$$\mathcal{L} := \sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} \mathbb{E}_\tau (\ell(i, \sigma', \tau))$$

denotes the loss in entropy compared to a uniform distribution.

For a fixed i and σ' , note that $s(i, \sigma', \tau)$ is the number of out-neighbours $j \in N^+(i)$ of i in G such that $(\sigma')^{-1}(j)$ does not appear before i in τ . Since τ is a uniformly random permutation, this number is distributed uniformly at random in $[d]$, and so it follows that

$$\mathcal{H}(\sigma) = \sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} \frac{1}{d} \sum_{t=1}^d \log t - \mathcal{L} = \frac{n}{d} \log(d!) - \mathcal{L}.$$

Using Stirling's approximation and recalling that $|\mathcal{C}| \geq (d/e)^n$, we get

$$n \log \left(\frac{d}{e} \right) \leq \log |\mathcal{C}| = \mathcal{H}(\sigma) = \frac{n}{d} \log(d!) - \mathcal{L} \leq n \log \left(\frac{d}{e} \right) + \frac{n}{d} \log(ed) - \mathcal{L},$$

and therefore $\mathcal{L} \leq (n \log(ed))/d$.

We now bound the number of cycles $|\sigma|$ of σ using a very similar calculation. For every permutation $\tau \in S_n$, we say that $\sigma(i)$ closes a cycle if $\tau_{<i}$ contains all other vertices that are on the same cycle as i in σ . That is, adding $\sigma(i)$ to $\sigma(\tau_{<i})$ creates one additional cycle. Note that $|\sigma|$ is equal to the number of vertices $i \in [n]$ for which $\sigma(i)$ closes a cycle, and so

$$\begin{aligned} \mathbb{E}_\sigma(|\sigma|) &= \sum_{i=1}^n \mathbb{P}(\sigma(i) \text{ closes a cycle}) \\ &= \sum_{i=1}^n \sum_{\pi} \mathbb{P}(\sigma(\tau_{<i}) = \pi) \cdot \mathbb{P}(\sigma(i) \text{ closes a cycle} \mid \sigma(\tau_{<i}) = \pi) \\ &= \sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} \mathbb{P}(\sigma(i) \text{ closes a cycle} \mid \sigma(\tau_{<i}) = \sigma'(\tau_{<i})). \end{aligned}$$

If $\sigma(\tau_{<i}) = \sigma'(\tau_{<i})$, then either $\sigma(i)$ cannot close a cycle, or $\sigma(i)$ needs to be one specific out-neighbour $j \in N^+(i) \setminus \sigma'(\tau_{<i})$ in order to close a cycle. So, by [Lemma 3.3](#) we know that

$$\begin{aligned} \mathbb{P}(\sigma(i) \text{ closes a cycle} \mid \sigma(\tau_{<i}) = \sigma'(\tau_{<i})) &\leq \max_{j \in N^+(i)} \mathbb{P}(\sigma(i) = j \mid \sigma(\tau_{<i}) = \sigma'(\tau_{<i})) \\ &\leq \frac{2}{s(i, \sigma', \tau)} + \ell(i, \sigma', \tau). \end{aligned}$$

This implies that

$$\mathbb{E}_\sigma(|\sigma|) \leq \sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} \left(\frac{2}{s(i, \sigma', \tau)} + \ell(i, \sigma', \tau) \right).$$

Since this holds for all $\tau \in S_n$, we may again pick $\tau \in S_n$ uniformly at random. As noted before, $s(i, \sigma', \tau)$ is then distributed uniformly at random in $[d]$, and so

$$\begin{aligned} \mathbb{E}_\sigma(|\sigma|) &\leq \mathbb{E}_\tau \left(\sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} \left(\frac{2}{s(i, \sigma', \tau)} + \ell(i, \sigma', \tau) \right) \right) \\ &= \sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} \mathbb{E}_\tau \left(\frac{2}{s(i, \sigma', \tau)} \right) + \mathcal{L} \\ &= \sum_{i=1}^n \frac{1}{|\mathcal{C}|} \sum_{\sigma' \in \mathcal{C}} \frac{1}{d} \sum_{t=1}^d \frac{2}{t} + \mathcal{L} \\ &\leq 2 \cdot \frac{n}{d} (\log d + 1) + \mathcal{L} \leq 4 \cdot \frac{n}{d} (\log d + 1). \end{aligned} \quad \square$$

We finish by showing that there is a randomised polynomial-time algorithm that finds a cycle-factor of G with at most $\mathcal{O}((n \log d)/d)$ cycles with high probability. For this, we will use the fact that there exists an algorithm that samples an almost-uniform⁶ random perfect matching in bipartite graphs, which was proved by Jerrum, Sinclair, and Vigoda [[JSV04](#)].

Theorem 3.4. *There exists a randomised polynomial-time algorithm that samples an almost-uniform random perfect matching of a bipartite graph.*

Using the fact that cycle-factors correspond to perfect matchings in the bipartite graph H constructed before, [Corollary 1.2](#) is now an easy consequence.

Proof of [Corollary 1.2](#). Recall that cycle-factors of G correspond to perfect matchings in the bipartite graph H constructed before. By [Theorem 3.4](#) we can sample an almost-uniform random perfect matching of H in polynomial time, and so we can sample an almost-uniform random cycle-factor σ of G in polynomial time.

By [Theorem 1.1](#) and Markov's inequality, we know that a uniformly random cycle-factor of G has at most $\mathcal{O}((n \log d)/d)$ cycles with probability at least $1/2$. Since σ is almost-uniform with some error δ , this implies that σ has at most $\mathcal{O}((n \log d)/d)$ cycles with probability at least $1/2 - \delta$. If we now sample σ independently at least $\Omega(\log n)$ times and return the cycle-factor with the least number of cycles, then the returned cycle-factor has at most $\mathcal{O}((n \log d)/d)$ cycles with high probability. \square

Acknowledgements. We would like to thank Thomas Karam for an informative lecture on Radhakrishnan's entropy proof of the Brègman-Minc inequality.

⁶If X is a random variable that takes values in a finite set \mathcal{X} , then X is *almost-uniform* with error δ if the total variation distance between X and the uniform distribution is at most δ , that is $\sum_{x \in \mathcal{X}} |1/|\mathcal{X}| - \mathbb{P}(X = x)| \leq \delta$.

References

- [AEH80] JIN AKIYAMA, GEOFFREY EXOO, and FRANK HARARY (1980). [Covering and packing in graphs. III. Cyclic and acyclic invariants](#). *Mathematica Slovaca* **30**(4), 405–417. [↑2](#)
- [AS16] NOGA ALON and JOEL H. SPENCER (2016). The probabilistic method. Fourth edn. (John Wiley & Sons). [↑4](#)
- [AZ18] MARTIN AIGNER and GÜNTER M. ZIEGLER (2018). [Proofs from The Book](#). Sixth edn. (Springer, Berlin). [↑3](#)
- [Brè73] LEV MEEROVICH BRÈGMAN (1973). [Some properties of nonnegative matrices and their permanents](#). *Soviet Mathematics-Doklady* **14**, 945–949. [↑4](#)
- [Cau40] AUGUSTIN LOUIS CAUCHY (1840). Exercices d’analyse et de physique mathématique: 1, vol. 1 (Bachelier, imprimeur-libraire). [↑1](#)
- [Ego81] GREGORY P. EGORYCHEV (1981). [The solution of van der Waerden’s problem for permanents](#). *Advances in Mathematics* **42**(3), 299–305. [↑4](#)
- [Fal81] DMITRY I. FALIKMAN (1981). [Proof of the van der Waerden conjecture regarding the permanent of a doubly stochastic matrix](#). *Matematicheskie Zametki* **29**(6), 931–938. [↑4](#)
- [FF22] URIEL FEIGE and ELLA FUCHS (2022). [On the path partition number of 6-regular graphs](#). *Journal of Graph Theory* **101**(3), 345–378. [↑2](#)
- [FRS14] URIEL FEIGE, R. RAVI, and MOHIT SINGH (2014). [Short tours through large linear forests](#). *Integer Programming and Combinatorial Optimization, IPCO 2014, Lecture Notes in Computer Science*, vol. 8494, 273–284. [↑2](#)
- [GL21] VYTAUTAS GRUSLYS and SHOHAM LETZTER (2021). [Cycle partitions of regular graphs](#). *Combinatorics, Probability and Computing* **30**(4), 526–549. [↑2](#)
- [JSV04] MARK JERRUM, ALISTAIR SINCLAIR, and ERIC VIGODA (2004). [A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries](#). *Journal of the ACM (JACM)* **51**(4), 671–697. [↑2, 7](#)
- [LMS25] SHOHAM LETZTER, ABHISHEK METHUKU, and BENNY SUDAKOV (2025). [Nearly Hamilton cycles in sublinear expanders, and applications](#). arXiv:2503.07147. [↑2](#)
- [LP23] RICHARD LANG and LUKE POSTLE (2023). [An improved bound for the linear arboricity conjecture](#). *Combinatorica* **43**(3), 547–569. [↑2](#)
- [MM09] COLTON MAGNANT and DANIEL M. MARTIN (2009). [A note on the path cover number of regular graphs](#). *The Australasian Journal of Combinatorics* **43**, 211–217. [↑2](#)
- [MMPS24] RICHARD MONTGOMERY, ALP MÜYESSER, ALEXEY POKROVSKIY, and BENNY SUDAKOV (2024). [Approximate path decompositions of regular graphs](#). *Journal of the London Mathematical Society*. To appear. [↑2](#)
- [Rad97] JAIKUMAR RADHAKRISHNAN (1997). [An entropy proof of Bregman’s theorem](#). *Journal of Combinatorial Theory, Series A* **77**(1), 161–164. [↑3, 4](#)
- [Vis12] NISHEETH K. VISHNOI (2012). [A permanent approach to the traveling salesman problem](#). *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, 76–80. [↑2, 4](#)