

CYCLES AND EIGENVALUES OF SEQUENTIALLY GROWING RANDOM REGULAR GRAPHS

TOBIAS JOHNSON AND SOUMIK PAL

ABSTRACT. Consider the sum of d many iid random permutation matrices on n labels along with their transposes. The resulting matrix is the adjacency matrix of a random regular (multi)-graph of degree $2d$ on n vertices. It is known that the distribution of smooth linear eigenvalue statistics of this matrix is given asymptotically by sums of Poisson random variables. This is in contrast with Gaussian fluctuation of similar quantities in the case of Wigner matrices. It is also known that for Wigner matrices the joint fluctuation of linear eigenvalue statistics across minors of growing sizes can be expressed in terms of the Gaussian Free Field (GFF). In this article we explore joint asymptotic (in n) fluctuation for a coupling of all random regular graphs of various degrees obtained by growing each component permutation according to the Chinese Restaurant Process. Our primary result is that the corresponding eigenvalue statistics can be expressed in terms of a family of independent Yule processes with immigration. These processes track the evolution of short cycles in the graph. If we now take d to infinity, certain GFF-like properties emerge.

1. INTRODUCTION

We consider graphs that have labeled vertices and are regular, i.e., every vertex has the same degree. We allow our graphs to have loops and multiple edges (such graphs are sometimes called multigraphs or pseudographs). Additionally, our graphs will be sparse in the sense that the degree will be negligible compared to the order. Every such graph has an associated adjacency matrix whose (i, j) th element is the number of edges between vertices i and j , with loops counted twice. When the graph is randomly selected, the matrix is random, and we are interested in studying the eigenvalues of the resulting symmetric matrix. Note that, due to regularity, it does not matter whether we consider the eigenvalues of the adjacency or the Laplacian matrix.

The precise distribution of this random regular graph is somewhat ad hoc. We will use what is called the permutation model. Consider the permutation digraphs generated by d many iid random permutations on n labels. We remove the direction of the edge and *collapse* all these graphs on one another. This results in a $2d$ -regular graph on n vertices, denoted by $G(n, 2d)$. At the matrix level this is given by adding all the d permutation matrices and their transposes.

Date: May 29, 2013.

2000 Mathematics Subject Classification. 60B20, 05C80.

Key words and phrases. Random regular graphs, eigenvalue fluctuations, Chinese Restaurant Process, minors of random matrices.

This research is partially supported by NSF grant DMS-1007563.

Our present work is an extension of the study of eigenvalue fluctuations carried out in [DJPP12]. We are motivated by the recent work by Borodin on joint eigenvalue fluctuations of minors of Wigner matrices and the (massless or zero-boundary) Gaussian Free Field (GFF) [Bor10a, Bor10b]. Eigenvalues of minors are closely related to interacting particle systems [Fer10, FF10], and the KPZ universality class of random surfaces [BF08]. See [JN06] for more on eigenvalues of minors of GUE and [ANvM11] for those of Dyson's Brownian motion.

Let us consider a particular but important case of Borodin's result in [Bor10a] (single sequence, the entire \mathbb{N}). An $n \times n$ real symmetric Wigner matrix has iid upper triangular off-diagonal elements with four moments identical to the standard Gaussian. The diagonal elements are usually taken to be iid with mean zero variance two. Notice that every principal submatrix (called minors in this context) of a Wigner matrix is again a Wigner matrix of a smaller order. Thus, on some probability space one can construct an infinite order Wigner matrix W whose $n \times n$ minor $W(n)$ is a Wigner matrix of order n .

Let z be a complex number in the upper half plane \mathbb{H} . Define $y = |z|^2$ and $x = 2\Re(z)$. Consider the minor $W(\lfloor ny \rfloor)$, and let $N(z)$ be the number of its eigenvalues that are greater than or equal to \sqrt{nx} . Define the *height function*

$$(1) \quad H_n(z) := \sqrt{\frac{\pi}{2}} N(z).$$

Then Borodin shows that $\{H_n(z) - EH_n(z), z \in \mathbb{H}\}$, viewed as distributions, converges in law to a generalized Gaussian process on \mathbb{H} with a covariance kernel

$$(2) \quad C(z, w) = \frac{1}{2\pi} \ln \left| \frac{z - w}{z - \bar{w}} \right|.$$

The above is the covariance kernel for the GFF on the upper half plane.

An equivalent assertion is the following. Let $[n]$ denote the set of integers $\{1, 2, \dots, n\}$. Consider the Chebyshev polynomials of the first kind, $\{T_n, n = 0, 1, 2, \dots\}$, on the interval $[-1, 1]$. These polynomials are given by the identity $T_n(\cos(\theta)) \equiv \cos(n\theta)$. We specialize [Bor10a, Proposition 3] for the case of GOE ($\beta = 1$). Fix m positive real numbers $t_1 < t_2 < \dots < t_m$. In the notation of [Bor10a], we take $L = n$ and $B_i(n) = \lfloor t_i n \rfloor$. Then, for any positive integers j_1, j_2, \dots, j_m , the random vector

$$\left(\operatorname{tr} T_{j_i} \left(W(\lfloor t_i n \rfloor) / 2\sqrt{t_i n} \right) - \mathbf{E} \operatorname{tr} T_{j_i} \left(W(\lfloor t_i n \rfloor) / 2\sqrt{t_i n} \right), i \in [m] \right)$$

converges in law, as n tends to infinity, to a centered Gaussian vector. For $s \leq t$,

$$(3) \quad \lim_{n \rightarrow \infty} \mathbf{Cov} \left(\operatorname{tr} T_i \left(W(\lfloor tn \rfloor) / 2\sqrt{tn} \right), \operatorname{tr} T_k \left(W(\lfloor sn \rfloor) / 2\sqrt{sn} \right) \right) = \delta_{ik} \frac{k}{2} \left(\frac{s}{t} \right)^{k/2},$$

which gives the covariance kernel of the limiting vector. In particular, all such covariances are zero when $i \neq k$. Note that the traces can be expressed as integrals of the height function of the corresponding submatrices. Thus, by approximating continuous compactly supported functions of z by a function that is piecewise constant in y and polynomial in x , one gets the kernel (2).

1.1. Main results. By a tower of random permutations we mean a sequence of random permutations $(\pi^{(n)}, n \in \mathbb{N})$ such that

- (i) $\pi^{(n)}$ is a uniformly distributed random permutation of $[n]$ for each n , and

- (ii) for each n , if $\pi^{(n)}$ is written as a product of cycles then $\pi^{(n-1)}$ is derived from $\pi^{(n)}$ by deletion of the element n from its cycle.

The stochastic process that grows $\pi^{(n)}$ from $\pi^{(n-1)}$ by sequentially inserting an element n randomly is called the Chinese Restaurant Process (CRP). We will review the basic principles at a later section. In [KOV04] and other related work, a sequence of permutations satisfying condition (ii) is called a *virtual permutation*, and the distribution on virtual permutations satisfying condition (i) is considered as a substitute for Haar measure on $S(\infty)$, the infinite symmetric group. This is used to study the representation theory of $S(\infty)$, with connections to Random Matrix Theory. A recent extension of this idea is [BNN11].

Now suppose we construct a countable collection $\{\Pi_d, d \in \mathbb{N}\}$ of towers of random permutations. We will denote the permutations in Π_d by $\{\pi_d^{(n)}, n \in \mathbb{N}\}$. Then it is possible to model every possible $G(n, 2d)$ by adding the permutation matrices (and their transposes) corresponding to $\{\pi_j^{(n)}, 1 \leq j \leq d\}$. In what follows we will keep d fixed and consider n as a growing parameter. Thus, G_n will represent $G(n, 2d)$ for some fixed d . Here and later, G_0 will represent the empty graph. We construct a continuous-time version of this by inserting new vertices into G_n with rate $n + 1$. Formally, define independent times $T_i \sim \text{Exp}(i)$, and let

$$M_t = \max \left\{ m : \sum_{i=1}^m T_i \leq t \right\},$$

and define the continuous-time Markov chain $G(t) = G_{M_t}$. When $d = 1$, this process is essentially just a continuous-time version of the CRP itself. Though this case is unusual compared to the rest—for example, $G(t)$ is likely to be disconnected when $d = 1$ and connected when d is larger—our results do still hold.

Our first result is about the process of *short cycles* in the graph process $G(t)$. By a cycle of length k in a graph, we mean what is sometimes called a simple cycle: a walk in the graph that begins and ends at the same vertex, and that otherwise repeats no vertices. We will give a more formal definition in Section 2.2. Let $(C_k^{(s)}(t), k \in \mathbb{N})$ denote the number of cycles of various lengths k that are present in $G(s + t)$. This process is not Markov, but nonetheless it converges to a Markov process (indexed by t) as s tends to infinity.

To describe the limit, define

$$a(d, k) = \begin{cases} (2d - 1)^k - 1 + 2d, & \text{when } k \text{ is even,} \\ (2d - 1)^k + 1, & \text{when } k \text{ is odd.} \end{cases}$$

Consider the set of natural numbers $\mathbb{N} = \{1, 2, \dots\}$ with the measure

$$\mu(k) = \frac{1}{2} [a(d, k) - a(d, k - 1)], \quad k \in \mathbb{N}, \quad a(d, 0) := 0.$$

Consider a Poisson point process χ on $\mathbb{N} \times [0, \infty)$ with an intensity measure given on $\mathbb{N} \times (0, \infty)$ by the product measure $\mu \otimes \text{Leb}$, where Leb is the Lebesgue measure, and with additional masses of $a(d, k)/2k$ on $(k, 0)$ for $k \in \mathbb{N}$.

Let \tilde{P}_x denote the law of an one-dimensional pure-birth process on \mathbb{N} given by the generator:

$$Lf(k) = k(f(k + 1) - f(k)), \quad k \in \mathbb{N},$$

starting from $x \in \mathbb{N}$. This is also known as the *Yule process*.

Suppose we are given a realization of χ . For any atom (k, y) of the countably many atoms of χ , we start an independent process $(X_{k,y}(t), t \geq 0)$ with law \tilde{P}_k . Define the random sequence

$$N_k(t) := \sum_{(j,y) \in \chi \cap \{[k] \times [0,t]\}} 1 \{X_{j,y}(t-y) = k\}.$$

In other words, at time t , for every site k , we count how many of the processes that started at time $y \leq t$ at site $j \leq k$ are currently at k . Note that both $(N_k(\cdot), k \in \mathbb{N})$ and $(N_k(\cdot), k \in [K])$, for some $K \in \mathbb{N}$, are Markov processes, while $N_k(\cdot)$ for fixed k is not.

Theorem 1. *As $s \rightarrow \infty$, the process $(C_k^{(s)}(t), k \in \mathbb{N}, 0 \leq t < \infty)$ converges in law, in the product topology on $D^\infty[0, \infty)$, to the Markov process $(N_k(t), k \in \mathbb{N}, 0 \leq t < \infty)$. The limiting process is stationary.*

Remark 2. In fact, the same argument used to prove Theorem 1 shows that the process $(C_k^{(s)}(t), -\infty < t < \infty)$ converges in law to the Markov process $(N_k(t), -\infty < t < \infty)$ running in stationarity. The same conclusion holds for all the following theorems in this section.

We now explore the joint convergence across various d 's. Define $C_{d,k}^{(s)}(t)$ naturally, stressing the dependence on the parameter d .

Theorem 3. *There is a joint process convergence of $(C_{i,k}^{(s)}(t), k \in \mathbb{N}, i \in [d], t \geq 0)$ to a limiting process $(N_{i,k}(t), k \in \mathbb{N}, i \in [d], t \geq 0)$. This limit is a Markov process whose marginal law for every fixed d is described in Theorem 1. Moreover, for any $d \in \mathbb{N}$, the process $(N_{d+1,k}(\cdot) - N_{d,k}(\cdot), k \in \mathbb{N})$ is independent of the process $(N_{i,k}(\cdot), k \in \mathbb{N}, i \in [d])$ and evolves as a Markov process. Its generator (defined on functions dependent on finitely many coordinates) is given by*

$$Lf(x) = \sum_{k=1}^{\infty} kx_k [f(x + e_{k+1} - e_k) - f(x)] + \sum_{k=1}^{\infty} \nu(d, k) [f(x + e_k) - f(x)],$$

where x is a nonnegative sequence, $(e_k, k \in \mathbb{N})$ are the canonical orthonormal basis of ℓ^2 , and

$$\nu(d, k) = \frac{1}{2} [a(d+1, k) - a(d+1, k-1) - a(d, k) + a(d, k-1)].$$

Remark 4. Theorems 1 and 3 show an underlying branching process structure. We actually prove a more general decomposition where cycles are tracked by edge labels. The additive structure also imparts a natural intertwining relationship between the Markov operators. See [CPY98, Section 2] and [DF90, Bor10a].

We now focus on eigenvalues of $G(t)$. Note that there is no easy exact relationship between the eigenvalues of G_n for various n since the eigenvectors play a role in determining any such identity. In fact, the eigenvalues of G_n and G_{n+1} need not be interlaced. However, one can consider linear eigenvalue statistics for the graph $G(n, 2d)$. That is, for any d -regular graph on n vertices G and function $f: \mathbb{R} \rightarrow \mathbb{R}$, define the random variable

$$\text{tr } f(G) := \sum_{i=1}^n f(\lambda_i)$$

where $\lambda_1 \geq \dots \geq \lambda_n$ are the eigenvalues of adjacency matrix of G divided by $2(2d-1)^{1/2}$. The scaling is necessary to take a limit with respect to d .

By a polynomial basis we refer to a sequence of polynomials $\{f_0 \equiv 1, f_1, f_2, \dots\}$ such that f_k is a polynomial of degree k of a single argument over reals. In the statement below $[\infty]$ will refer to \mathbb{N} .

Theorem 5. *There exists a polynomial basis $\{f_i, i \in \mathbb{N}\}$ (depending on d) such that for any $K \in \mathbb{N} \cup \{\infty\}$, the process $(\text{tr } f_k(G(s+t)), k \in [K], t \geq 0)$ converges in law, as s tends to infinity, to the Markov process $(N_k(t), k \in [K], t \geq 0)$ of Theorem 1. (The polynomials are given explicitly in (15).) Hence, for any polynomial f , the process $(\text{tr } f(G(s+t)))$ converges to a linear combination of the coordinate processes of $(N_k(t), k \in \mathbb{N})$.*

The Markov property is especially intriguing since, to the best of our knowledge, no similar property of eigenvalues of the standard Random Matrix ensembles is known. For the special case of minors of the Gaussian Unitary/Orthogonal Ensembles, the entire distribution of eigenvalues across minors of various sizes do satisfy a Markov property. However, this is facilitated by the known symmetry properties of the eigenvectors, and do not extend to other examples of Wigner matrices.

For our final result we will take d to infinity. We will make the following notational convention: for any polynomial f , we will denote the limiting process of $(\text{tr } f(G(s+t)), t \geq 0)$ by $(\text{tr } f(G(\infty+t)), t \geq 0)$. Recall that this process is a linear combination of $(N_k(t), k \in \mathbb{N}, t \geq 0)$.

Theorem 6. *Let $\{T_k, k \in \mathbb{N}\}$ denote the Chebyshev orthogonal polynomials of the first kind on $[-1, 1]$. As d tends to infinity, the collection of processes*

$$(\text{tr } T_k(G(\infty+t)) - \mathbf{E} \text{tr } T_k(G(\infty+t))), t \geq 0, k \in \mathbb{N}$$

converges weakly in $D^\infty[0, \infty)$ to a collection of independent Ornstein-Uhlenbeck processes $(U_k(t), t \geq 0, k \in \mathbb{N})$, running in equilibrium. Here the equilibrium distribution of U_k is $N(0, k/2)$ and U_k satisfies the stochastic differential equation

$$dU_k(t) = -kU_k(t) dt + k dW_k(t), \quad t \geq 0,$$

and $(W_k, k \in \mathbb{N})$ are iid standard one-dimensional Brownian motions.

Thus, the collection of random variables $(\text{tr } T_k(G(\infty+t)) - \mathbf{E} \text{tr } T_k(G(\infty+t)))$, indexed by k and t , converges as d tends to infinity to a centered Gaussian process with covariance kernel given by

$$(4) \quad \lim_{d \rightarrow \infty} \mathbf{Cov}(\text{tr } T_i(G(\infty+t)), \text{tr } T_k(G(\infty+s))) = \delta_{ik} \frac{k}{2} e^{k(s-t)}.$$

for $s \leq t$.

A comparison of (4) with Borodin's result (3) shows that the above limit captures a key property of the GFF covariance structure. The appearance of the exponential is merely due to a deterministic time-change of the process. A somewhat more detailed discussion can be found in the following section.

Remark 7. A common model for random regular graphs is the *configuration model* or *pairing model* (see [Wor99] for more information). The model is defined as follows: Start with n buckets, each containing d prevertices. Then, separate these dn prevertices into pairs, choosing uniformly from every possible pairing. Finally, collapse each bucket into a single vertex, making an edge between one vertex and

another if a prevertex in one bucket is paired with a prevertex in the other bucket. This model has the advantage that choosing a graph from it conditional on it containing no loops or parallel edges is the same as choosing a graph uniformly from the set of graphs without loops and parallel edges. The model also allows for graphs of odd degrees, unlike the permutation model.

It is possible to construct a process of growing random regular graphs similar to the one in this paper using a dynamic version of this model. Given some initial pairing of prevertices labeled $\{1, \dots, dn\}$, extend it to a random pairing of $\{1, \dots, dn+2\}$ by the following procedure: Choose X uniformly from $\{1, \dots, dn+1\}$. Pair $dn+2$ with X . If $X = dn+1$, leave the other pairs unchanged; if not, pair the previous partner of X with $dn+1$. This is an analogue of the CRP in the setting of random pairings, in that if the initial pairing is uniformly chosen, then so is the extended one.

If d is odd, we repeat this procedure a total of d times to extend a random d -regular graph on n vertices to have $n+2$ vertices (when d is odd, the number of vertices in the graph must be even). When d is even, repeat $d/2$ times to add one new vertex to a random graph. In this way, we can construct a sequence of growing random regular graphs. We believe that all the results of this paper hold in this model with minor changes, with similar proofs.

1.2. Existing literature. The study of the spectral properties of sparse regular random graphs is motivated by several different problems. These matrices do not fall within the purview of the standard techniques of Random Matrix Theory (RMT) due to their sparsity and lack of independence between entries. However, extensive simulations ([JMRR99]) point to conjectures that these matrices still belong to the universality class of random matrices. For example, it is conjectured via simulations ([MNS08]) that the distribution of the second largest eigenvalue (in absolute value) is given by the Tracy-Widom distribution. In the physics literature, eigenvalues of random regular graphs have been considered as a toy model of quantum chaos ([Smi10], [OGS09], [OS10]). Simulations suggest that the eigenvalue spacing distribution has the same limit as that of the Wigner matrices. A limiting *Gaussian wave* character of eigenvectors have also been conjectured ([Elo08, Elo10, ES10]). Some fine properties of eigenvalues and eigenvectors can indeed be proved for a single permutation matrix; see [Wie00] and [BAD11].

Somewhat complicating the matter is the fact that when the degree d is kept fixed and we let n go to infinity, several classical results about random matrix ensembles fail. A bit more elaboration on this point is needed. The two parameters in the ensemble of random graphs are the degree d and the order n . In the permutation model it is possible to construct random regular graphs for every possible value of (d, n) where d is an even positive integer and n is any positive integer. Hence one can consider various kinds of limits of these parameters. We will refer as the *diagonal* limit the procedure of having a sequence of (d, n) where both these parameters simultaneously go to infinity. To maintain sparsity¹, it is usually assumed that d is at most poly-logarithmic in n . No lower bound on the growth rate of d is assumed. However, results are often easier to prove when d is kept fixed and we let n go to infinity. Suppose for each d one gets a limiting object (say a probability distribution); one can now take d to infinity and explore limits of the sequence of

¹The non-sparse can be typically absorbed within standard techniques of RMT by comparing with a corresponding Erdős-Rényi graph whose adjacency matrix has independent entries.

these objects. We will refer to this procedure ($\lim_{d \rightarrow \infty} \lim_{n \rightarrow \infty}$) as the triangular limit. The triangular limit is often identical to the diagonal limit irrespective of the sequence through which the diagonal limit is taken, while maintaining sparsity. Moreover, these limiting statistics frequently match with those of the GOE ensemble and the real symmetric Wigner matrices. This is true, for example, for the empirical spectral distribution [DP12, TVW13] and fluctuations of smooth linear eigenvalue statistics [DJPP12].

Our present result is a triangular limit result. Let us first explain the connection with the massless GFF. We follow Definition 2.12 and the first example in Section 2.5 of [She07]. Consider the space of smooth real functions compactly supported on \mathbb{H} with the Dirichlet inner product $\langle f, g \rangle = \int_{\mathbb{H}} \nabla f \cdot \nabla g \, dz$. We consider the completion of this pre-Hilbert space, and call it the Sobolev space H^1 . The GFF can be thought of as a random distribution h which associates with every $f \in H^1$ a mean zero Gaussian random variable $\langle h, f \rangle$ that is an L^2 isometry in the following sense $\text{Cov}(\langle h, f \rangle, \langle h, g \rangle) = \langle f, g \rangle$. Now, one can perform the integration of a function f with h by first integrating their traces over semicircular arcs of a fixed radius, and then a further integral over the radius. Over the semicircular arcs Fourier transforms (or Chebyshev Polynomials, for real functions) provide an orthogonal basis for this Gaussian field. As one parametrizes the radius properly, one obtains independent Ornstein-Uhlenbeck processes for each Chebyshev polynomial. Hence these OU processes completely determine the GFF covariance structure. This explains the word ‘equivalent’ on page 2, paragraph 4, and is the essence of the calculations done in [Bor10a]. See also [Spo98] for a similar formalism for Dyson’s Brownian motion on the circle.

One of the reasons why we cannot prove a full GFF convergence is that the parameters d and n behave independently of one another. The degree d determines the support of the spectral distribution $[-2\sqrt{2d-1}, 2\sqrt{2d-1}]$, asymptotically independent of n . For Wigner matrices, the dimension itself determines the length of the spectral support. This results in the parametrization of (1). It should be possible to extend our results to a GFF convergence by either letting d grow with n in the graph, or, even by letting d grow with time for the limiting Poisson structure in Theorem 1. However this has not been attempted in the present article.

2. PRELIMINARIES

2.1. A primer on the Chinese Restaurant Process. The CRP, introduced by Dubins and Pitman, is a particular example of a two parameter family of stochastic processes that constructs sequentially random exchangeable partitions of the positive integers via the cyclic decomposition of a random permutation. Our short description is taken from [Pit06, Section 3.1].

An initially empty restaurant has an unlimited number of circular tables numbered $1, 2, \dots$, each capable of seating an unlimited number of customers. Customers numbered $1, 2, \dots$ arrive one by one and are seated at the tables according to the following plan. Person 1 sits at table 1. For $n \geq 1$ suppose that n customers have already entered the restaurant, and are seated in some arrangement, with at least one customer at each of the tables j for $1 \leq j \leq k$ (say), where k is the number of tables occupied by the first n customers to arrive. Let customer $n+1$ choose with equal probability to sit at any of the following $n+1$ places: to the left of customer j for some $1 \leq j \leq n$, or alone at table $k+1$. Define $\pi^{(n)}: [n] \rightarrow [n]$ as




FIGURE 1. A cycle whose word is the equivalence class of $\pi_2\pi_1^{-1}\pi_2\pi_1\pi_2\pi_3^{-1}$ in \mathcal{W}_6/D_{12} .

the permutation whose cyclic decomposition is given by the tables; that is, if after n customers have entered the restaurant, customers i and j are seated at the same table, with i to the left of j , then $\pi^{(n)}(i) = j$, and if customer i is seated alone at some table then $\pi^{(n)}(i) = i$. The sequence $(\pi^{(n)})$ then has features (i) and (ii) mentioned in the first paragraph of Section 1.1.

2.2. Combinatorics on words. The graph G_n , formed from the independent permutations $\pi_1^{(n)}, \dots, \pi_d^{(n)}$, can be considered as a directed, edge-labeled graph in a natural way. For convenience, drop superscripts and let $\pi_l = \pi_l^{(n)}$. If $\pi_l(i) = j$, then by definition G_n contains an edge between i to j . When convenient, we consider this edge to be directed from i to j and to be labeled by π_l .

Consider a walk on G_n , viewed in this way, and imagine writing down the label of each edge as it is traversed, putting π_i or π_i^{-1} according to the direction we walk over the edge. We call a walk *closed* if it starts and ends at the same vertex, and we call a closed walk a *cycle* if it never visits a vertex twice (besides the first and last one), and it never traverses an edge more than once in either direction. Thus the word $w = w_1 \cdots w_k$ formed as a cycle is traversed is cyclically reduced, i.e., $w_i \neq w_{i+1}^{-1}$ for all i , considering i modulo k . For example, following an edge and then immediately backtracking does not form a 2-cycle, and the word formed by this walk is $\pi_i\pi_i^{-1}$ or $\pi_i^{-1}\pi_i$ for some i , which is not cyclically reduced. We consider two cycles equivalent if they are both walks on an identical set of edges; that is, we ignore the starting vertex and the direction of the walk.

Let \mathcal{W}_k denote the set of cyclically reduced words of length k . We would like to associate each k -cycle in G_n with the word in \mathcal{W}_k formed by the above procedure, but since we can start the walk at any point in the cycle and walk in either of two directions, there are actually up to $2k$ different words that could be formed by it. Thus we identify elements of \mathcal{W}_k that differ only by rotation and inversion (for example, $\pi_1\pi_2^{-1}\pi_1\pi_2$ and $\pi_1^{-1}\pi_2\pi_1^{-1}\pi_2^{-1}$) and denote the resulting set by \mathcal{W}_k/D_{2k} , where D_{2k} is the dihedral group acting on the set \mathcal{W}_k in the natural way.

Definition 8 (Properties of words). For any k -cycle in G_n , the element of \mathcal{W}_k/D_{2k} given by walking around the cycle is called the *word* of the cycle (see Figure 1). For any word w , let $|w|$ denote the length of w . Let $h(w)$ be the largest number m such that $w = u^m$ for some word u . If $h(w) = 1$, we call w *primitive*. For any $w \in \mathcal{W}_k$, the orbit of w under the action of D_{2k} contains $2k/h(w)$ elements, a fact which we will frequently use. Let $c(w)$ denote the number of pairs of double letters in w , i.e., the number of integers i modulo $|w|$ such that $w_i = w_{i+1}$. For example,

$c(\pi_1\pi_1\pi_2^{-1}\pi_2^{-1}\pi_1) = 3$. We will also consider $|\cdot|$, $h(\cdot)$, and $c(\cdot)$ as functions on \mathcal{W}_k/D_{2k} , since they are invariant under cyclic rotation and inversion.

To more easily refer to words in \mathcal{W}_k/D_{2k} , choose some canonical representative $w_1 \cdots w_k \in \mathcal{W}_k$ for every $w \in \mathcal{W}_k/D_{2k}$. Based on this, we will often think of elements of \mathcal{W}_k/D_{2k} as words instead of equivalence classes, and we will make statements about the i th letter of a word in \mathcal{W}_k/D_{2k} . For $w = w_1 \cdots w_k \in \mathcal{W}_k/D_{2k}$, let $w^{(i)}$ refer to the word in $\mathcal{W}_{k+1}/D_{2k+2}$ given by $w_1 \cdots w_i w_i w_{i+1} \cdots w_k$. We refer to this operation as *doubling* the i th letter of w . A related operation is to *halve* a pair of double letters, for example producing $\pi_1\pi_2\pi_3\pi_4$ from $\pi_1\pi_2\pi_3\pi_4\pi_1$. (Since we apply these operations to words identified with their rotations, we do not need to be specific about which letter of the pair is deleted.) The following technical lemma underpins most of our combinatorial calculations.

Lemma 9. *Let $u \in \mathcal{W}_k/D_{2k}$ and $w \in \mathcal{W}_{k+1}/D_{2k+2}$. Suppose that a letters in u can be doubled to form w , and b pairs of double letters in w can be halved to form u . Then*

$$\frac{a}{h(u)} = \frac{b}{h(w)}.$$

Remark 10. At first glance, one might expect that $a = b$. The example $u = \pi_1\pi_2\pi_1\pi_1\pi_2$ and $w = \pi_1\pi_1\pi_2\pi_1\pi_1\pi_2$ shows that this is wrong, since only one letter in u can be doubled to give w , but two different pairs in w can be halved to give u .

Proof. Let $\text{Orb}(u)$ and $\text{Orb}(w)$ denote the orbits of u and w under the action of the dihedral group in \mathcal{W}_k and \mathcal{W}_{k+1} , respectively. When we speak of halving a pair of letters in a word in $\text{Orb}(w)$, always delete the second of the two letters (for example, $\pi_1\pi_2\pi_1$ becomes $\pi_1\pi_2$, not $\pi_2\pi_1$). When we double a letter in a word in $\text{Orb}(u)$, put the new letter after the doubled letter (for example, doubling the second letter of $\pi_1\pi_2^{-1}$ gives $\pi_1\pi_2^{-1}\pi_2^{-1}$, not $\pi_2^{-1}\pi_1\pi_2^{-1}$.)

For each of the $2k/h(u)$ words in $\text{Orb}(u)$, there are a doubling operations yielding a word in $\text{Orb}(w)$. For each of the $(2k+2)/h(w)$ words in $\text{Orb}(w)$, there are b halving operations yielding a word in $\text{Orb}(u)$. For every halving operation on a word in $\text{Orb}(w)$, there is a corresponding doubling operation on a word in $\text{Orb}(u)$ and vice versa, except for halving operations that straddle the ends of the word, as in $\pi_1\pi_2\pi_1$. There are $2b/h(w)$ of these, giving us

$$\begin{aligned} \frac{2ka}{h(u)} &= \frac{(2k+2)b}{h(w)} - \frac{2b}{h(w)} \\ &= \frac{2kb}{h(w)}, \end{aligned}$$

and the lemma follows from this. \square

Let $\mathcal{W}' = \bigcup_{k=1}^{\infty} \mathcal{W}_k/D_{2k}$, and let $\mathcal{W}'_K = \bigcup_{k=1}^K \mathcal{W}_k/D_{2k}$. We will use the previous lemma to prove the following technical property of the $c(\cdot)$ statistic.

Lemma 11. *In the vector space with basis $\{q_w\}_{w \in \mathcal{W}'_K}$,*

$$\sum_{w \in \mathcal{W}'_{K-1}} \sum_{i=1}^{|w|} \frac{1}{h(w)} q_{w^{(i)}} = \sum_{w \in \mathcal{W}'_K} \frac{c(w)}{h(w)} q_w.$$




FIGURE 2. The vertex 6 is inserted between vertices 2 and 3 in π_1 , causing the above cycle to grow.

Proof. Fix some $w \in \mathcal{W}_k/D_{2k}$, and let $a(u)$ denote the number of letters of u that can be doubled to give w , for any $u \in \mathcal{W}_{k-1}/D_{2k-2}$. We need to prove that

$$\sum_{u \in \mathcal{W}_{k-1}/D_{2k-2}} \frac{a(u)}{h(u)} = \frac{c(w)}{h(w)}.$$

Let $b(u)$ be the number of pairs in w that can be halved to give u . By Lemma 9,

$$\sum_{u \in \mathcal{W}_{k-1}/D_{2k-2}} \frac{a(u)}{h(u)} = \sum_{u \in \mathcal{W}_{k-1}/D_{2k-2}} \frac{b(u)}{h(w)},$$

and $\sum_{u \in \mathcal{W}_{k-1}/D_{2k-2}} b(u) = c(w)$. \square

3. THE PROCESS LIMIT OF THE CYCLE STRUCTURE

As the graph $G(t)$ grows, new cycles form, which we can classify into two types. Suppose a new vertex numbered n is inserted at time t , and this insertion creates a new cycle. If the edges entering and leaving vertex n in the new cycle have the same edge label, then the new cycle has “grown” from a cycle with one fewer vertex, as in Figure 2. If the edges entering and leaving n in the cycle have different labels, then the cycle has formed “spontaneously” as in Figure 3, rather than growing from a smaller cycle. This classification will prove essential in understanding the evolution of cycles in $G(t)$.

Once a cycle comes into existence in $G(t)$, it remains until a new vertex is inserted into one of its edges. Typically, this results in the cycle growing to a larger cycle, as in Figure 2. If a new vertex is simultaneously inserted into multiple edges of the same cycle, the cycle is instead split into smaller cycles as in Figure 4. These new cycles are spontaneously formed, according to the classification of new cycles given in the previous paragraph. Tracking the evolution of these smaller cycles in turn, we see that as the graph evolves, a cycle grows into a cluster of overlapping cycles. However, it will follow from Proposition 19 that for short cycles, this behavior is not typical. Thus in our limiting object, cycles will grow only into larger cycles.




FIGURE 3. A cycle forms “spontaneously” when the vertex 6 is inserted into the graph.




FIGURE 4. The vertex 6 is inserted into the cycle in two different places in the same step, causing the cycle to split in two. Note that each new cycle would be classified as spontaneously formed.

3.1. Heuristics for the limiting process. We give some estimates that will motivate the definition of the limiting process in Section 3.2. This section is entirely motivational, and we will not attempt to make anything rigorous.

Suppose that vertex n is inserted into $G(t)$ at some time t . First, we consider the rate that cycles form spontaneously with some word $w \in \mathcal{W}_k/D_{2k}$. There are $2k/h(w)$ words in the orbit of w under the action of D_{2k} , and out of these, $2(k - c(w))/h(w)$ have nonequal first and last letters. For each such word $u = u_1 \cdots u_k$, we can give a walk on the graph by starting at vertex n and following the edges indicated by u , going from n to $u_1(n)$ to $u_2(u_1(n))$ and so on. If this walk happens to be a cycle, the condition $u_1 \neq u_k$ implies that it would be spontaneously formed.


In a short interval Δt when $G(t)$ has $n - 1$ vertices, the probability that vertex n is inserted is about $n \Delta t$. For any word u , the walk from vertex n generated by u is a cycle with probability approximately $1/n$, since after applying the random permutations u_1, \dots, u_k in turn, we will be left at an approximately uniform random vertex. Any new spontaneous cycle formed with word w will be counted by one of these walks, with u in the orbit of w , and it will be counted again by the walk generated by $u_k^{-1} \cdots u_1^{-1}$. The expected number of spontaneous cycles formed in a

short interval Δt is then approximately

$$\frac{1}{h(w)}(k - c(w)) \frac{n \Delta t}{n} = \frac{1}{h(w)}(k - c(w)) \Delta t.$$

Thus we will model the spontaneous formation of cycles with word w by a Poisson process with rate $(k - c(w))/h(w)$.

Next, we consider how often a cycle with word $w \in \mathcal{W}_k$ grows into a larger cycle. Suppose that $G(t)$ has $n - 1$ vertices, and that it contains a cycle of the form



When vertex n is inserted into the graph, the probability that it is inserted after s_{i-1} in permutation w_i is $1/n$. Thus, after a spontaneous cycle with word w has formed, we can model the evolution of its word as a continuous-time Markov chain where each letter is doubled with rate one.

3.2. Formal definition of the limiting process. Consider the measure μ on \mathcal{W}' given by

$$\mu(w) = \frac{|w| - c(w)}{h(w)}.$$

Consider a Poisson point process χ on $\mathcal{W}' \times [0, \infty)$ with an intensity measure given by the product measure $\mu \otimes \text{Leb}$, where Leb refers to the Lebesgue measure. Each atom (w, t) of χ represents a new spontaneous cycle with word w formed at time t .

Now, we define a continuous-time Markov chain on the countable space \mathcal{W}' governed by the following rates: From state $w \in \mathcal{W}_k/D_{2k}$, jump with rate one to each of the k words in $\mathcal{W}_{k+1}/D_{2k+2}$ obtained by doubling a letter of w . If a word can be formed in more than one way by doubling a letter in w , then it receives a correspondingly higher rate. For example, from $w = \pi_1\pi_1\pi_2$, the chain jumps to $\pi_1\pi_1\pi_1\pi_2$ with rate two and to $\pi_1\pi_1\pi_2\pi_2$ with rate one. Let \tilde{P}_w denote the law of this process started from $w \in \mathcal{W}'$.

Suppose we are given a realization of χ . For any atom (w, s) of the countably many atoms of χ , we start an independent process $(X_{w,s}(t), t \geq 0)$ with law \tilde{P}_w . Define the stochastic process

$$N_w(t) := \sum_{\substack{(u,s) \in \chi \\ s \leq t}} 1 \{X_{u,s}(t-s) = w\}.$$

Interpreting these processes as in the previous section, $N_w(t)$ counts the number of cycles formed spontaneously at time s that have grown to have word w at time t .

The fact that the process exists is obvious since one can define the countably many independent Markov chains on a suitable product space. The following lemma establishes some of its key properties.

Lemma 12. *Recall that $\mathcal{W}'_L = \bigcup_{k=1}^L \mathcal{W}_k/D_{2k}$. We have the following conclusions:*

- (i) *For any $L \in \mathbb{N}$, the stochastic process $\{(N_w(t), w \in \mathcal{W}'_L), t \geq 0\}$ is a time-homogeneous Markov process with respect to its natural filtration, with RCLL paths.*

(ii) Recall that for $w \in \mathcal{W}_k/D_{2k}$, the element $w^{(i)} \in \mathcal{W}_{k+1}/D_{2k+2}$ is the word formed by doubling the i th letter of w . The generator for the Markov process $\{(N_w(t), w \in \mathcal{W}'_L), t \geq 0\}$ acts on f at $x = (x_w, w \in \mathcal{W}'_L)$ by

$$\begin{aligned} \mathcal{L}f(x) &= \sum_{w \in \mathcal{W}'_L} \sum_{i=1}^{|w|} x_w [f(x - e_w + e_{w^{(i)}}) - f(x)] \\ &\quad + \sum_{w \in \mathcal{W}'_L} \frac{|w| - c(w)}{h(w)} [f(x + e_w) - f(x)], \end{aligned}$$

where e_w is the canonical basis vector equal to one at entry w and equal to zero everywhere else. For a word u of length greater than L , take $e_u = 0$.

(iii) The product measure of $\text{Poi}(1/h(w))$ over all $w \in \mathcal{W}'_L$ is the unique invariant measure for this Markov process.

Proof. Conclusion (i) follows from construction, as does conclusion (ii). To prove conclusion (iii), we start by the fundamental identity of the Poisson distribution: if $X \sim \text{Poi}(\lambda)$, then for any function f , we have

$$(5) \quad \mathbf{E}Xg(X) = \lambda \mathbf{E}g(X + 1).$$

We need to show that if the coordinates of $X = (X_w, w \in \mathcal{W}'_L)$ are independent Poisson random variables with $\mathbf{E}X_w = 1/h(w)$, then

$$(6) \quad \mathbf{E}\mathcal{L}f(X) = 0.$$

Since the process is an irreducible Markov chain on countable state space, the existence of one invariant distribution shows that the chain is positive recurrent and that the invariant distribution is unique.

To argue (6) we will repeatedly apply identity (5) to functions g constructed from f by keeping all but one coordinate fixed. Thus, for any $w \in \mathcal{W}'_L$ and $1 \leq i \leq |w|$, we condition on all X_u with $u \neq w$ and hold those coordinates of f fixed to obtain,

$$\mathbf{E}X_w f(X - e_w + e_{w^{(i)}}) = \frac{1}{h(w)} \mathbf{E}f(X + e_{w^{(i)}})$$

taking $e_{w^{(i)}} = 0$ when $|w| = L$. In the same way,

$$\mathbf{E}X_w f(X) = \frac{1}{h(w)} \mathbf{E}f(X + e_w).$$

By these two equalities,

$$\begin{aligned} \mathbf{E} \sum_{w \in \mathcal{W}'_L} \sum_{i=1}^{|w|} X_w [f(X - e_w + e_{w^{(i)}}) - f(X)] \\ &= \sum_{w \in \mathcal{W}'_L} \sum_{i=1}^{|w|} \frac{1}{h(w)} \mathbf{E}[f(X + e_{w^{(i)}}) - f(X + e_w)] \\ &= \sum_{w \in \mathcal{W}'_{L-1}} \sum_{i=1}^{|w|} \frac{1}{h(w)} \mathbf{E}f(X + e_{w^{(i)}}) + \sum_{w \in \mathcal{W}_L/D_{2L}} \frac{|w|}{h(w)} \mathbf{E}f(X) \\ &\quad - \sum_{w \in \mathcal{W}'_L} \frac{|w|}{h(w)} \mathbf{E}f(X + e_w). \end{aligned}$$

Specializing Lemma 11 to $q_w = \mathbf{E}f(X + e_w)$, the first sum is

$$\sum_{w \in \mathcal{W}'_{L-1}} \sum_{i=1}^{|w|} \frac{1}{h(w)} \mathbf{E}f(X + e_{w^{(i)}}) = \sum_{w \in \mathcal{W}'_L} \frac{c(w)}{h(w)} \mathbf{E}f(X + e_w),$$

which gives us

$$\begin{aligned} \mathbf{E} \sum_{w \in \mathcal{W}'_L} \sum_{i=1}^{|w|} X_w [f(X - e_w + e_{w^{(i)}}) - f(X)] \\ = \sum_{w \in \mathcal{W}'_L} \frac{c(w) - |w|}{h(w)} \mathbf{E}f(X + e_w) + \sum_{w \in \mathcal{W}_L/D_{2L}} \frac{|w|}{h(w)} \mathbf{E}f(X). \end{aligned}$$

All that remains in proving (6) is to show that

$$\sum_{w \in \mathcal{W}'_L} \frac{|w| - c(w)}{h(w)} = \sum_{w \in \mathcal{W}_L/D_{2L}} \frac{|w|}{h(w)}.$$

Specializing Lemma 11 to $q_w = 1$ shows that $\sum_{w \in \mathcal{W}'_L} c(w)/h(w) = \sum_{w \in \mathcal{W}'_{L-1}} |w|/h(w)$.

Thus

$$\begin{aligned} \sum_{w \in \mathcal{W}'_L} \frac{|w| - c(w)}{h(w)} &= \sum_{w \in \mathcal{W}'_L} \frac{|w|}{h(w)} - \sum_{w \in \mathcal{W}'_{L-1}} \frac{|w|}{h(w)} \\ &= \sum_{w \in \mathcal{W}_L/D_{2L}} \frac{|w|}{h(w)}, \end{aligned}$$

establishing (6) and completing the proof. \square

From now on, we will consider the process $(N_w(t), k \in \mathbb{N}, t \geq 0)$ to be running under stationarity, i.e., with marginal distributions given by conclusion (iii) of the last lemma. This process is easily constructed as described above, but with additional point masses of weight $1/h(w)$ for each $w \in \mathcal{W}'$ at $(w, 0)$ added to the intensity measure of χ , thus giving us the correct distribution at time zero.

3.3. Time-reversed processes. Fix some time $T > 0$. We define the time-reversal $\overleftarrow{N}_w(t) := N_w(T - t)$ for $0 \leq t \leq T$.

Lemma 13. *For any fixed $L \in \mathbb{N}$, the process $\{\{\overleftarrow{N}_w(t), w \in \mathcal{W}'_L\}, 0 \leq t \leq T\}$ is a time-homogenous Markov process with respect to the natural filtration. A trivial modification at jump times renders RCLL paths. The transition rates of this chain are given as follows. Let $u \in \mathcal{W}_{k-1}/D_{2k-1}$ and $w \in \mathcal{W}_k/D_{2k}$, and suppose that u can be obtained from w by halving b different pairs. Let $x = (x_w, w \in \mathcal{W}'_L)$.*

- (i) *The chain jumps from x to $x + e_u - e_w$ with rate bx_w .*
- (ii) *The chain jumps from x to $x - e_w$ with rate $(k - c(w))x_w$.*
- (iii) *If $w \in \mathcal{W}_L/D_{2L}$, then the chain jumps from x to $x + e_w$ with rate $L/h(w)$.*

Proof. Any Markov process run backwards under stationarity is Markov. If the chain has transition rate $r(x, y)$ from states x to y , then the transition rate of the backwards chain from x to y is $r(y, x)\nu(y)/\nu(x)$, where ν is the stationary distribution. We will let ν be the stationary distribution from Lemma 12iii and calculate the transition rates of the backwards chain, using the rates given in Lemma 12ii.

Let a denote the number of letters in u that give w when doubled. The transition rate of the original chain from $x + e_u - e_w$ to x is $a(x_u + 1)$, so the transition rate of the backwards chain from x to $x + e_u - e_w$ is

$$a(x_u + 1) \frac{\nu(x + e_{k-1, c-1} - e_{k, c})}{\nu(x)} = \frac{ah(w)x_w}{h(u)},$$

and this is equal to bx_w by Lemma 9. A similar calculation shows that the transition rate from x to $x - e_w$ is

$$\frac{(k - c(w))\nu(x - e_w)}{h(w)\nu(x)} = (k - c(w))x_w,$$

proving (ii). The transition rate from x to $x + e_w$ for $w \in \mathcal{W}_L/D_{2L}$ is

$$\frac{\nu(x + e_w)}{\nu(x)}(x_w + 1)L = \frac{L}{h(w)},$$

which completes the proof. \square

By definition,

$$\overleftarrow{N}_w(t) = \sum_{\substack{(u, s) \in \mathcal{X} \\ s \leq T-t}} 1 \{X_{u, s}(T - t - s) = w\}.$$

We will modify this slightly to define the process

$$\overleftarrow{M}_w(t) := \sum_{\substack{(u, s) \in \mathcal{X} \\ s \leq T-t}} 1 \{X_{u, s}(T - t - s) = w \text{ and } |X_{u, s}(T - s)| \leq L\}.$$

The idea is that $\overleftarrow{M}_w(t)$ is the same as $\overleftarrow{N}_w(t)$, except that it does not count cycles at time t that had more than L vertices at time zero. The process $(\overleftarrow{M}_w(t), w \in \mathcal{W}'_L)$ is a Markov chain with the same transition rates as $(\overleftarrow{N}_w(t), w \in \mathcal{W}'_L)$, except that it does not jump from x to $x + e_w$ for $w \in \mathcal{W}_L/D_{2L}$. These two chains also have the same initial distribution, but $(\overleftarrow{M}_w(t), w \in \mathcal{W}'_L)$ is not stationary (in fact, it is eventually absorbed at zero).

4. PROCESS CONVERGENCE

Recall that $C_k^{(s)}(t)$ is the number of cycles of length k in the graph $G(s + t)$, defined on p. 3. For $w \in \mathcal{W}'$, let $C_w^{(s)}(t)$ be the number of cycles in $G(s + t)$ with word w . We will prove that $(C_w^{(s)}(\cdot), w \in \mathcal{W}')$ converges to a distributional limit, from which the convergence of $(C_k^{(s)}(\cdot), k \in \mathbb{N})$ will follow. The proof depends on knowing the limiting *marginal* distribution of $C_w^{(s)}(t)$. We provide this and more in the following theorem, which should be of independent interest:

Theorem 14. *Let $G_n = G(n, 2d)$, a $2d$ -regular random graph on n vertices from the permutation model. For any k , let \mathcal{I}_k be the set of all cycles of length k on the complete graph K_n with edge labels that form a cyclically reduced word; these are the possible k -cycles that might appear in G_n . Let $\mathcal{I} = \bigcup_{k=1}^r \mathcal{I}_k$ for some integer r .*

For any cycle $\alpha \in \mathcal{I}$, let $I_\alpha = 1\{G_n \text{ contains } \alpha\}$, and let $\mathbf{I} = (I_\alpha, \alpha \in \mathcal{I})$. Let $\mathbf{Z} = (Z_\alpha, \alpha \in \mathcal{I})$ be a vector whose coordinates are independent Poisson random

variables with $\mathbf{E}Z_\alpha = 1/[n]_k$ for $\alpha \in \mathcal{I}_k$. Then for all $d \geq 2$ and $n, r \geq 1$,

$$d_{TV}(\mathbf{I}, \mathbf{Z}) \leq \frac{c(2d-1)^{2r-1}}{n}$$

for some absolute constant c , where $d_{TV}(X, Y)$ denotes the total variation distance between the laws of X and Y .

Corollary 15. *Let $\{Z_w, w \in \mathcal{W}'_K\}$ be a family of independent Poisson random variables with $\mathbf{E}Z_w = 1/h(w)$. For any fixed integer K and $d \geq 1$,*

(i) *as $t \rightarrow \infty$,*

$$(C_w(t), w \in \mathcal{W}'_K) \xrightarrow{\mathcal{L}} (Z_w, w \in \mathcal{W}'_K);$$

(ii) *as $t \rightarrow \infty$, the probability that there exist two cycles of length K or less sharing a vertex in $G(t)$ approaches zero.*

We give the proofs in the appendix, along with some further discussion. Now, we turn to the convergence of the processes.

Theorem 16. *The process $(C_w^{(s)}(\cdot), w \in \mathcal{W}')$ converges in law as $s \rightarrow \infty$ to $(N_w(\cdot), w \in \mathcal{W}')$.*

Proof. The main difficulty in turning the intuitive ideas of Section 3.1 into an actual proof is that $(C_w^{(s)}(t), w \in \mathcal{W}')$ is not Markov. We now sketch how we evade this problem. We will run our chain backwards, defining $\overleftarrow{G}_s(t) = G(s+T-t)$ for some fixed $T > 0$. Then, we ignore all of $\overleftarrow{G}_s(0)$ except for the subgraph consisting of cycles of size L and smaller, which we will call $\overleftarrow{\Gamma}_s(0)$. The graph $\overleftarrow{\Gamma}_s(t)$ is the evolution of this subgraph as time runs backward, ignoring the rest of $\overleftarrow{G}_s(t)$. Then, we consider the number of cycles with word w in $\overleftarrow{\Gamma}_s(t)$, which we call $\phi_w(\overleftarrow{\Gamma}_s(t))$. Choose $K \ll L$. Then $\phi_w(\overleftarrow{\Gamma}_s(t))$ is likely to be the same as $C_w^{(s)}(T-t)$ for any word w with $|w| \leq K$. The remarkable fact that makes $\phi_w(\overleftarrow{\Gamma}_s(t))$ possible to analyze is that if $\overleftarrow{\Gamma}_s(0)$ consists of disjoint cycles, then $(\phi_w(\overleftarrow{\Gamma}_s(t)), w \in \mathcal{W}'_L)$ is a Markov chain governed by the same transition rates as $(\overleftarrow{M}_w(t), w \in \mathcal{W}'_L)$.

Another important idea of the proof is to ignore the vertex labels in $\overleftarrow{G}_s(t)$, so that we do not know in what order the vertices will be removed. Thus we can view $\overleftarrow{G}_s(t)$ as a Markov chain with the following description: Assign each vertex an independent $\text{Exp}(1)$ clock. When the clock of vertex v goes off, remove it from the graph, and patch together the π_i -labeled edges entering and leaving v for each $1 \leq i \leq d$.

Step 1. *Definitions of $\overleftarrow{\Gamma}_s(t)$ and ϕ_w and analysis of $(\phi_w(\overleftarrow{\Gamma}_s(t)), w \in \mathcal{W}'_L)$.*

Fix $T > 0$ and define $\overleftarrow{G}_s(t) = G(s+T-t)$. As mentioned above, we will consider $\overleftarrow{G}_s(t)$ only up to relabeling of vertices, which makes it a process on the countable state space consisting of all edge-labeled graphs on finitely many unlabeled vertices. With respect to its natural filtration, it is a Markov chain in which each vertex is removed with rate one, as described above.

To formally define $\overleftarrow{\Gamma}_s(t)$, fix integers $L > K$ and let $\overleftarrow{\Gamma}_s(0)$ be the subgraph of $\overleftarrow{G}_s(0)$ made up of all cycles of length L or less. We then evolve $\overleftarrow{\Gamma}_s(t)$ in parallel with $\overleftarrow{G}_s(t)$. When a vertex v is deleted from $\overleftarrow{G}_s(t)$, the corresponding vertex v in

$\overleftarrow{\Gamma}_s(t)$ is deleted if it is present. If v has a π_i -labeled edge entering and leaving it in $\overleftarrow{\Gamma}_s(t)$, then these two edges are patched together. Other edges in $\overleftarrow{\Gamma}_s(t)$ adjacent to v are deleted. This makes $\overleftarrow{\Gamma}_s(t)$ a subgraph of $\overleftarrow{G}_s(t)$, as well as a continuous-time Markov chain on the countable state space consisting of all edge-labeled graphs on finitely many unlabeled vertices. The transition probabilities of $\overleftarrow{\Gamma}_s(t)$ do not depend on s .

From Corollary 15, we can find the limiting distribution of $\overleftarrow{\Gamma}_s(0)$. Suppose that γ is a graph in the process's state space that is not a disjoint union of cycles. By Corollary 15ii,

$$\lim_{s \rightarrow \infty} \mathbf{P}[\overleftarrow{\Gamma}_s(0) = \gamma] = 0.$$

Suppose instead that γ is made up of disjoint cycles, with z_w cycles of word w for each $w \in \mathcal{W}'_L$. By Corollary 15i,

$$(7) \quad \lim_{s \rightarrow \infty} \mathbf{P}[\overleftarrow{\Gamma}_s(0) = \gamma] = \prod_{w \in \mathcal{W}'_L} \mathbf{P}[Z_w = z_w],$$

where $(Z_w, w \in \mathcal{W}'_L)$ are independent Poisson random variables with $\mathbf{E}Z_w = 1/h(w)$. Thus $\overleftarrow{\Gamma}_s(0)$ converges in law as $s \rightarrow \infty$ to a limiting distribution supported on the graphs made up of disjoint unions of cycles. For different values of s , the chains $\overleftarrow{\Gamma}_s(t)$ differ only in their initial distributions, and the convergence in law of $\overleftarrow{\Gamma}_s(0)$ as $s \rightarrow \infty$ induces the process convergence of $\{\overleftarrow{\Gamma}_s(t), 0 \leq t \leq T\}$ to a Markov chain $\{\overleftarrow{\Gamma}(t), 0 \leq t \leq T\}$ with the same transition rates whose initial distribution is the limit of $\overleftarrow{\Gamma}_s(0)$.

For any finite edge-labeled graph G , let $\phi_w(G)$ be the number of cycles in G with word w . By the continuous mapping theorem, the process $(\phi_w(\overleftarrow{\Gamma}_s(t)), w \in \mathcal{W}'_L)$ converges in law to $(\phi_w(\overleftarrow{\Gamma}(t)), w \in \mathcal{W}'_L)$ as $s \rightarrow \infty$.

We will now demonstrate that this process has the same law as $(\overleftarrow{M}_w(t), w \in \mathcal{W}'_L)$. The graph $\overleftarrow{\Gamma}(t)$ consists of disjoint cycles at time $t = 0$, and as it evolves, these cycles shrink or are destroyed. The process $(\phi_w(\overleftarrow{\Gamma}(t)), w \in \mathcal{W}'_L)$ jumps exactly when a vertex in a cycle in $\overleftarrow{\Gamma}(t)$ is deleted. If the deleted vertex lies in a cycle between two edges with the same label, the cycle shrinks. If the deleted vertex lies in a cycle between two edges with different labels, the cycle is destroyed. The only relevant consideration in where the process will jump at time t is the number of vertices of these two types in $\overleftarrow{\Gamma}(t)$, which can be deduced from $(\phi_w(\overleftarrow{\Gamma}(t)), w \in \mathcal{W}'_L)$. Thus this process is a Markov chain.

Consider two words $u, w \in \mathcal{W}'_K$ such that w can be obtained from u by doubling a letter. Suppose that u can be obtained from w by halving any of b pairs of letters. Suppose that the chain is at state $x = (x_v, v \in \mathcal{W}'_L)$. There are bx_w vertices that when deleted cause the chain to jump from x to $x - e_w + e_u$, each of which is removed with rate one. Thus the chain jumps from x to $x - e_w + e_u$ with rate bx_w . Similarly, it jumps to $x - e_w$ with rate $(|w| - c(w))x_w$. These are the same rates as the chain $(\overleftarrow{M}_w(t), w \in \mathcal{W}'_L)$ from Section 3.3. The initial distribution given by (7) is also the same as that of $(\overleftarrow{M}_w(t), w \in \mathcal{W}'_L)$, demonstrating that the two processes $(\phi_w(\overleftarrow{\Gamma}(t)), w \in \mathcal{W}'_L)$ and $(\overleftarrow{M}_w(t), w \in \mathcal{W}'_L)$ have the same law.

Step 2. *Approximation of $\overleftarrow{C}_w^{(s)}(t)$ by $\phi_w(\overleftarrow{\Gamma}_s(t))$.*

We will compare the two processes $\{(\overleftarrow{C}_w^{(s)}(t), w \in \mathcal{W}'_K), 0 \leq t \leq T\}$ and $\{(\phi_w(\overleftarrow{\Gamma}_s(t)), w \in \mathcal{W}'_K), 0 \leq t \leq T\}$ and show that for sufficiently large L , they are identical with probability arbitrarily close to one.

Consider some cycle in $\overleftarrow{G}_s(t)$; we can divide its vertices into those that lie between two edges of the cycle with different labels, and those that lie between two edges with the same label. We call this second class the *shrinking vertices* of the cycle, because if one is deleted from $\overleftarrow{G}_s(t)$ as it evolves, the cycle shrinks. We define $E_s(L)$ to be the event that for some cycle in $\overleftarrow{G}_s(0)$ of size $l > L$, at least $l - K$ of its shrinking vertices are deleted by time T .

We claim that outside of the event $E_s(L)$, the two processes $\{(\overleftarrow{C}_w^{(s)}(t), w \in \mathcal{W}'_K), 0 \leq t \leq T\}$ and $\{(\phi_w(\overleftarrow{\Gamma}_s(t)), w \in \mathcal{W}'_K), 0 \leq t \leq T\}$ are identical. Suppose that these two processes are not identical. Then there is some cycle α of size K or less present in $\overleftarrow{G}_s(t)$ but not in $\overleftarrow{\Gamma}_s(t)$ for $0 < t \leq T$. As explained in Section 3, as a cycle evolves (in forward time), it grows into an overlapping cluster of cycles. Thus $\overleftarrow{G}_s(0)$ contains some cluster of overlapping cycles that shrinks to α at time t . One of the cycles in this cluster has length greater than L , or the cluster would be contained in $\overleftarrow{\Gamma}_s(0)$ and α would have been contained in $\overleftarrow{\Gamma}_s(t)$.

To see that $l - K$ shrinking vertices must be deleted from this cycle, consider the evolution of α into the cluster of cycles in both forward and reverse time. If a vertex is inserted into a single edge of a cycle in forward time, we see in reverse time the deletion of a shrinking vertex. If a vertex is simultaneously inserted into two edges of a cycle, causing the cycle to split, we see in reverse time the deletion of a non-shrinking vertex of a cycle. As α grows, a cycle of size greater than L can form only by single-insertion of at least $l - K$ vertices into the eventual cycle. In reverse time, this is seen as deletion of $l - K$ shrinking vertices. This demonstrates that $E_s(L)$ holds.

We will now show that for any $\epsilon > 0$, there is an L sufficiently large that $\mathbf{P}[E_s(L)] < \epsilon$ for any s . Let $w \in \mathcal{W}_l/D_{2l}$ with $l > L$, and let $I \subseteq [l]$ such that $|I| = l - K$ and $w_i = w_{i-1}$ for all $i \in I$, considering indices modulo l . For any cycle in $\overleftarrow{G}_s(0)$ with word w , the set I corresponds to a set of $l - K$ shrinking vertices of the cycle.

We define $F(w, I)$ to be the event that $\overleftarrow{G}_s(0)$ contains one or more cycles with word w , and that the vertices corresponding to I in one of these cycles are all deleted within time T . By a union bound,

$$(8) \quad \mathbf{P}[E_s(L)] \leq \sum_{w, I} \mathbf{P}[F(w, I)].$$

We proceed by enumerating all pairs of w and I . For any pair w, I , deleting the letters in w at positions given by I results in a word $u \in \mathcal{W}_K/D_{2K}$. For any given $u = u_1 \cdots u_K \in \mathcal{W}_K/D_{2K}$, the word $w \in \mathcal{W}_l/D_{2l}$ must have the form

$$w = \underbrace{u_1 \cdots u_1}_{a_1 \text{ times}} \underbrace{u_2 \cdots u_2}_{a_2 \text{ times}} \cdots \underbrace{u_K \cdots u_K}_{a_K \text{ times}},$$

with $a_i \geq 1$ and $a_1 + \cdots + a_K = l$. The number of choices for a_1, \dots, a_K is $\binom{l-1}{K-1}$, the number of compositions of l into K parts, and each of these corresponds to a choice of w and I . There are fewer than $a(d, K)$ choices for u , giving us a bound of $a(d, K) \binom{l-1}{K-1}$ choices of pairs w and I for any fixed $l > L$.

Next, we will show that for any pair w and I with $|w| = l$,

$$(9) \quad \mathbf{P}[F(w, I)] \leq (1 - e^{-T})^{l-K}.$$

Condition on $\overleftarrow{G}_s(0)$ having n vertices. Consider any of the $[n]_l$ possible sequences of l vertices. Choose some representative $w' \in \mathcal{W}_l$ of w . For each of these sequences, the probability that it forms a cycle with word w' is at most $1/[n]_l$ (recall the original definition of our random graphs in terms of random permutations). Given that the sequence forms a cycle, the probability that the vertices of the cycle at positions I are all deleted within time T is $(1 - e^{-T})^{l-K}$. Hence

$$\begin{aligned} \mathbf{P}\left[F(w, I) \mid \overleftarrow{G}_s(0) \text{ has } n \text{ vertices}\right] &\leq [n]_l \frac{1}{[n]_l} (1 - e^{-T})^{l-K}, \\ &\leq (1 - e^{-T})^{l-K}. \end{aligned}$$

This holds for any n , establishing (9).

Applying all of this to (8),

$$\mathbf{P}[E_s(L)] \leq \sum_{l=L+1}^{\infty} a(d, K) \binom{l-1}{K-1} (1 - e^{-T})^{l-K}.$$

This sum converges, which means that for any $\epsilon > 0$, we have $\mathbf{P}[E_s(L)] < \epsilon$ for large enough L , independent of s .

Step 3. *Approximation of $\overleftarrow{N}_w(t)$ by $\overleftarrow{M}_w(t)$.*

Recall that we defined the processes $\{\overleftarrow{M}_w(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$ and $\{\overleftarrow{N}_w(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$ on the same probability space. We will show that for sufficiently large L , the two processes are identical with probability arbitrarily close to one.

By their definitions, these two processes are identical unless one of the processes $X_{u,s}(\cdot)$ started at each atom of χ grows from a word of size K or less to a word of size $L+1$ before time T ; we call this event $E(L)$. Let

$$Y = \left| \{(u, s) \in \chi : |u| \leq K, s \leq T\} \right|,$$

the number of processes starting from a word of size K or less before time T .

Suppose that $X(\cdot)$ has law \tilde{P}_w for some word $w \in \mathcal{W}_k/D_{2k}$. We can choose L large enough that $\mathbf{P}[|X(T)| > L] < \epsilon$ for all $k \leq K$. Then $\mathbf{P}[E(L) \mid Y] < \epsilon Y$ by a union bound, and so $\mathbf{P}[E(L)] < \epsilon \mathbf{E}Y$. Since $\mathbf{E}Y < \infty$, we can make $\mathbf{P}[E(L)]$ arbitrarily small by choosing sufficiently large L .

Step 4. *Weak convergence of $\{\overleftarrow{C}_w^{(s)}(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$ to $\{\overleftarrow{N}_w(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$.*

If two processes are identical with probability $1 - \epsilon$, then the total variation distance between their laws is at most ϵ . Thus, by steps 2 and 3, we can choose L large enough that the laws of the processes $\{\overleftarrow{C}_w^{(s)}(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$ and $\{(\phi_w(\overleftarrow{\Gamma}_s(t)), w \in \mathcal{W}'_K), 0 \leq t \leq T\}$ are arbitrarily close in total variation distance, uniformly in s , and so that the laws of $\{\overleftarrow{M}_w(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$ and $\{\overleftarrow{N}_w(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$ are arbitrarily close in total variation distance. Since total variation distance dominates the Prokhorov metric (or any other metric for the topology of weak convergence), we can choose L such that these two pairs are each within $\epsilon/3$ in the Prokhorov metric. Since $\{(\phi_w(\overleftarrow{\Gamma}_s(t)), w \in \mathcal{W}'_K), 0 \leq$

$t \leq T$ converges in law to $\{\overleftarrow{M}_w(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$ as $s \rightarrow \infty$, there is an s_0 such that for all $s \geq s_0$, the laws of these processes are within $\epsilon/3$ in the Prokhorov metric. We have thus shown that for every $\epsilon > 0$, the laws of $\{\overleftarrow{C}_w^{(s)}(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$ and $\{\overleftarrow{N}_w(t), w \in \mathcal{W}'_K, 0 \leq t \leq T\}$ are within ϵ for sufficiently large s , which proves that the first random vector converges in law to the second as $s \rightarrow \infty$.

Step 5. *Weak convergence of $\{(C_w^{(s)}(t), w \in \mathcal{W}'), t \geq 0\}$ to $\{(N_w(t), w \in \mathcal{W}'), t \geq 0\}$.*

It follows immediately from the previous step that the (not time-reversed) process $\{(C_w^{(s)}(t), w \in \mathcal{W}'_K), 0 \leq t \leq T\}$ converges in law to $\{(N_w(t), w \in \mathcal{W}'_K), 0 \leq t \leq T\}$ for any $T > 0$. By Theorem 16.17 in [Bil99], $\{(C_w^{(s)}(t), w \in \mathcal{W}'_K), t \geq 0\}$ converges in law to $\{(N_w(t), w \in \mathcal{W}'_K), t \geq 0\}$, which also proves that $\{(C_w^{(s)}(t), w \in \mathcal{W}'), t \geq 0\}$ converges in law to $\{(N_w(t), w \in \mathcal{W}'), t \geq 0\}$. \square

Proof of Theorem 1. We now consider the case of short cycles in the graph. We will express these as functionals of $(C_w^{(s)}(t), w \in \mathcal{W}')$. For example, consider the count of cycles of size $k \in \mathbb{N}$. Then $C_k^{(s)}(t) = \sum_{w \in \mathcal{W}_k/D_{2k}} C_w^{(s)}(t)$ is the number of k -cycles in $G(s+t)$, and let

$$N_k(t) = \sum_{w \in \mathcal{W}_k/D_{2k}} N_w(t).$$

It follows immediately from the continuous mapping theorem that $\{(C_k^{(s)}(t), k \in \mathbb{N}), t \geq 0\}$ converges in law to $\{(N_k(t), k \in \mathbb{N}), t \geq 0\}$ as $s \rightarrow \infty$.

It is not hard to see that this limit is Markov and admits the following representation: Cycles of size k appear spontaneously with rate $\sum_{w \in \mathcal{W}_k/D_{2k}} \mu(w)$. The size of each cycle then grows as a pure birth process with generator $Lf(i) = i(f(i+1) - f(i))$. The only thing we need to verify is that

$$(10) \quad \sum_{w \in \mathcal{W}_k/D_{2k}} \mu(w) = \sum_{w \in \mathcal{W}_k/D_{2k}} \frac{k - c(w)}{h(w)} = (a(d, k) - a(d, k-1))/2.$$

However, this follows from Lemma 11 in the following way. From that lemma we get

$$\sum_{w \in \mathcal{W}_k/D_{2k}} \frac{c(w)}{h(w)} = (k-1) \sum_{w \in \mathcal{W}_{k-1}/D_{2(k-1)}} \frac{1}{h(w)}.$$

Thus

$$\sum_{w \in \mathcal{W}_k/D_{2k}} \mu(w) = \sum_{w \in \mathcal{W}_k/D_{2k}} \frac{k}{h(w)} - \sum_{w \in \mathcal{W}_{k-1}/D_{2(k-1)}} \frac{k-1}{h(w)}.$$

However, the two terms on the right side of the above equation are simply half the total number of cyclically reduced words possible, of size k and $k-1$ respectively. The total number of cyclically reduced words of size k on an alphabet of size d is $a(d, k)$ (see Appendix of [DJPP12]). This shows (10) and completes the proof. \square

We end with the following corollary.

Corollary 17. *For any $s < t$ and $j, k \in \mathbb{N}$, one has:*

$$\mathbf{Cov}(N_k(t), N_j(s)) = \begin{cases} \frac{\alpha(d, j)}{2j} \binom{k-1}{k-j} p^j (1-p)^{k-j}, & p = e^{s-t}, \quad \text{if } k \geq j. \\ 0, & \text{otherwise.} \end{cases}$$

Proof. We will refer to the Yule processes counted by $N_k(t)$ as cycles of length k present at time t , even though these ‘‘cycles’’ in the limiting process have no connection to graphs. If $k < j$, every cycle that is of length j at time s cannot grow to a cycle of length k at time t . Thus, $N_k(t)$ depends on cycles that are independent of those that make up $N_j(s)$. Hence $N_k(t)$ is independent of $N_j(s)$.

If $k \geq j$, notice that one has the following decomposition:

$$(11) \quad N_k(t) = \sum_{j=1}^k \alpha(j, k) N_j(s) + Z,$$

where $\alpha(j, k)$ is the proportion of one-dimensional pure-birth Yule processes that were at state j at time s and grew to state k at time t , and Z is a random variable that counts the number of new births in the time interval (s, t) that grew to state k at time t . Note that, under our invariant distribution all random variables $\{Z, N_j(s), 1 \leq j \leq k\}$ are independent of one another. Thus, our conclusion follows once we show

$$(12) \quad \mathbf{E}\alpha(j, k) = \binom{k-1}{k-j} p^j (1-p)^{k-j}, \quad p = e^{s-t}.$$

The expected proportion $\mathbf{E}\alpha(j, k)$ is the probability that a one-dimensional process $X_{j,k}$, with law of an Yule process starting at j , is at state k at time $(t-s)$. If ξ_j, \dots, ξ_k are independent exponential random variables with rates j, \dots, k , then

$$\mathbf{E}\alpha(j, k) = \mathbf{P}[\{\xi_j + \dots + \xi_{k-1} \leq t-s\} \cap \{\xi_j + \dots + \xi_k > t-s\}]$$

We now use the Rényi representation: suppose Y_1, Y_2, \dots, Y_k are iid $\text{Exp}(1)$ random variables. Define the order statistics $Y_{(1)} \geq Y_{(2)} \geq \dots \geq Y_{(k)}$. Then, the following equality holds in distribution

$$(Y_{(i)} - Y_{(i+1)}), \quad j \leq i \leq k = (\xi_i, \quad j \leq i \leq k).$$

Here we have defined $Y_{(k+1)} \equiv 0$. Thus, in distribution,

$$\xi_j + \dots + \xi_{k-1} = Y_{(j)} - Y_{(k)}, \quad \xi_j + \dots + \xi_k = Y_{(j)}.$$

Thus

$$\mathbf{E}\alpha(j, k) = P(t-s < Y_{(j)} \leq Y_{(k)} + t-s).$$

Note that, by an elementary symmetry argument, for any $u > (t-s)$, we have

$$\begin{aligned} & \mathbf{P}[Y_{(j)} \in (u, u+du), Y_{(j)} - Y_{(k)} < t-s] \\ &= \mathbf{P}[Y_i = u \text{ for some } i, \text{ exactly } j-1 \text{ of } Y_1, \dots, Y_k \text{ are greater than } u, \\ & \quad \text{and the rest of } Y_1, \dots, Y_k \text{ are in } [u-t+s, u]] du \\ &= k e^{-u} \binom{k-1}{j-1} e^{-(j-1)u} [e^{-u+t-s} - e^{-u}]^{k-j} du \\ &= k \binom{k-1}{j-1} e^{-ku} (e^{t-s} - 1)^{k-j} du. \end{aligned}$$

Integrating out u in the interval $(t-s, \infty)$ we get

$$\begin{aligned} \mathbf{P} [t-s < Y_{(j)} < Y_{(k)} + t-s] &= \binom{k-1}{j-1} (e^{t-s} - 1)^{k-j} \int_{t-s}^{\infty} k e^{-ku} du \\ &= \binom{k-1}{j-1} (e^{t-s} - 1)^{k-j} e^{-k(t-s)} = \binom{k-1}{j-1} e^{j(s-t)} (1 - e^{s-t})^{k-j}. \end{aligned}$$

This shows (12) and completes the proof of the corollary. \square

4.1. Two-dimensional convergence. So far, we have considered d as a constant. We now view it as a parameter of the graph and allow it to vary. Recall that $(\Pi_d, d \in \mathbb{N})$ are independent towers of random permutations, with $\Pi_d = (\pi_d^{(n)}, n \in \mathbb{N})$, and that $G(n, 2d)$ is defined from $\pi_1^{(n)}, \dots, \pi_d^{(n)}$. For each d , we follow the construction used to define $G(t)$ and construct $G_{2d}(t)$, a continuous-time version of $(G(n, 2d), n \in \mathbb{N})$. Let $\mathcal{W}'(d)$ be the set of equivalence classes of cyclically reduced words as before, with the parameter d made explicit. Define $C_{d,k}^{(s)}(t)$ as the number of k -cycles in $G_{2d}(s+t)$ and consider the convergence of the two-dimensional field $\{(C_{d,k}^{(s)}(t), d, k \in \mathbb{N}), t \geq 0\}$ as $s \rightarrow \infty$.

Again, we will consider this process as a functional of another one. Define $\mathcal{W}'(\infty) = \bigcup_{d=1}^{\infty} \mathcal{W}'(d)$, noting that $\mathcal{W}'(1) \subseteq \mathcal{W}'(2) \subseteq \dots$. For any $w \in \mathcal{W}'(d)$, the number of cycles in $G_{2d'}(s+t)$ with word w is the same for all $d' \geq d$. We define $C_w^{(s)}(t)$ by this, so that

$$C_{d,k}^{(s)}(t) = \sum_{\substack{w \in \mathcal{W}'(d) \\ |w|=k}} C_w^{(s)}(t).$$

Then we will prove convergence of $\{(C_w^{(s)}(t), w \in \mathcal{W}'(\infty)), t \geq 0\}$ as $s \rightarrow \infty$.

To define a limit for this process, we extend μ to a measure on all of $\mathcal{W}'(\infty)$ and define the Poisson point process χ on $\mathcal{W}'(\infty) \times [0, \infty)$. The rest of the construction is identical to the one in Section 3.2, giving us random variables $(N_w(t), w \in \mathcal{W}'(\infty))$.

Theorem 18. *The process $(C_w^{(s)}(\cdot), w \in \mathcal{W}'(\infty))$ converges in law as $s \rightarrow \infty$ to $(N_w(\cdot), w \in \mathcal{W}'(\infty))$.*

Proof. It suffices to prove that $(C_w^{(s)}(\cdot), w \in \mathcal{W}'(d))$ converges in law as $s \rightarrow \infty$ to $(N_w(\cdot), w \in \mathcal{W}'(d))$ for each d , which we did in Theorem 16. \square

Proof of Theorem 3. Let

$$N_{d,k}(t) = \sum_{\substack{w \in \mathcal{W}'(d) \\ |w|=k}} N_w(t).$$

By the continuous mapping theorem, $(N_{d,k}(\cdot), d, k \in \mathbb{N})$ is the limit of $(C_{d,k}^{(s)}(\cdot), d, k \in \mathbb{N})$ as $s \rightarrow \infty$.

Let us now describe what the limiting process is. It is obvious that $(N_{d,k}(\cdot), k \in \mathbb{N}, d \in \mathbb{N})$ is jointly Markov. For every fixed d , the law of the corresponding marginal is given by Theorem 1. To understand the relationship across d , notice that cycles of size k for $(d+1)$ consist of cycles of size k for d and the extra ones

that contain an edge labeled by π_{d+1} of π_{d+1}^{-1} . Thus

$$N_{d+1,k}(t) - N_{d,k}(t) = \sum_{\substack{w \in \mathcal{W}'(d+1) \setminus \mathcal{W}'(d) \\ |w|=k}} N_w(t)$$

This process is independent of $(N_{i,\cdot}, i \in [d])$, since the set of words involved are disjoint. Moreover, the rates for this process are clearly the following: cycles of size k grow at rate k and new cycles of size k appear at rate $[a(d+1, k) - a(d+1, k-1) - a(d, k) + a(d, k-1)]/2$. This completes the proof of the result. \square

5. PROCESS LIMIT FOR LINEAR EIGENVALUE STATISTICS

Let us recall some of the basic facts established in [DJPP12, Sections 3, 5] that connect linear eigenvalue statistics with cycle counts. A closed non-backtracking walk is a walk that begins and ends at the same vertex, and that never follows an edge and immediately follows that same edge backwards. If the last step of a closed non-backtracking walk is anything other than the reverse of the first step, we say that the walk is *cyclically non-backtracking*. Cyclically non-backtracking walks on G_n are exactly the closed non-backtracking walks whose words are cyclically reduced. Let $\text{CNBW}_k^{(n)}$ denote the number of closed cyclically non-backtracking walks of length k on G_n .

Cyclically non-backtracking walks are useful because they can be enumerated by linear functionals of a graph's eigenvalues. Let $\{T_n(x)\}_{n \in \mathbb{N}}$ be the Chebyshev polynomials of the first kind on the interval $[-1, 1]$. We define a set of polynomials

$$\begin{aligned} \Gamma_0(x) &= 1, \\ \Gamma_{2k}(x) &= 2T_{2k}(x) + \frac{2d-2}{(2d-1)^k}, \quad \forall k \geq 1, \\ \Gamma_{2k+1}(x) &= 2T_{2k+1}(x), \quad \forall k \geq 0. \end{aligned}$$

Let A_n be the adjacency matrix of G_n , and let $\lambda_1 \geq \dots \geq \lambda_n$ be the eigenvalues of $(2d-1)^{-1/2}A_n/2$. Then

$$(13) \quad \sum_{i=1}^n \Gamma_k(\lambda_i) = (2d-1)^{-k/2} \text{CNBW}_k^{(n)}.$$

Now, for any cycle in G_n of length $j \mid k$, we obtain $2j$ non-backtracking walks of length k by choosing a starting point and direction and then walking around the cycle repeatedly. In [DJPP12, Corollary 18], it is shown that with certain conditions on the growth of d and r , all cyclically non-backtracking walks of length r or less have this form with high probability. Thus the random vectors $(\text{CNBW}_k^{(n)}, 1 \leq k \leq r)$ and $(\sum_{j \mid k} 2j C_k^{(n)}, 1 \leq k \leq r)$ have the same limiting distribution, and the problem of finding the limiting distributions of polynomial linear eigenvalue statistics is reduced to finding limiting distributions of cycle counts. We will prove Theorem 5 by arguing that this holds for the entire process $(G(t), t \geq 0)$.

Call a cyclically non-backtracking walk *bad* if it is anything other than a repeated walk around a cycle.

Proposition 19. *Fix an integer K . There is a random time T , almost surely finite, such that there are no bad cyclically non-backtracking walks of length K or less in $G(t)$ for all $t \geq T$.*

Proof. We will work with the discrete-time version of our process $(G_n, n \in \mathbb{N})$. We first define some machinery introduced in [LP10]. Consider some cyclically non-backtracking walk of length k on the edge-labeled complete graph K_n of the form

$$s_0 \xrightarrow{w_1} s_1 \xrightarrow{w_2} s_2 \xrightarrow{w_3} \cdots \xrightarrow{w_k} s_k = s_0.$$

Here, $s_i \in [n]$ and $w = w_1 \cdots w_k$ is the word of the walk (that is, each w_i is π_j or π_j^{-1} for some j , indicating which permutation provided the edge for the walk). We say that G_n contains the walk if the random permutations π_1, \dots, π_d satisfy $w_i(s_{i-1}) = s_i$. In other words, G_n contains a walk if considering both as edge-labeled directed graphs, the walk is a subgraph of G_n .

If $(s'_i, 0 \leq i \leq k)$ is another walk with the same word, we say that the two walks are of the same *category* if $s_i = s_j \iff s'_i = s'_j$. In other words, two walks are of the same category if they are identical up to relabeling vertices. The probability that G_n contains a walk depends only on its category. If a walk contains e distinct edges, then G_n contains the walk with probability at most $1/[n]_e$.

Let $X_k^{(n)}$ be the number of bad walks of length k in G_n that start at vertex n . We will first prove that with probability one, $X_k^{(n)} > 0$ for only finitely many n . Call a category bad if the walks in the category are bad. Let $\mathcal{T}_{k,d}$ be the number of bad categories of walks of length k . For any particular bad category whose walks contain v distinct vertices, there are $[n-1]_{v-1}$ walks of that category whose first vertex is n . Any bad walk contains more edges than vertices, so

$$\mathbf{E}X_k^{(n)} \leq \frac{\mathcal{T}_{k,d}[n-1]_{v-1}}{[n]_{v+1}} \leq \frac{\mathcal{T}_{k,d}}{n(n-k)}.$$

Since $X_k^{(n)}$ takes values in the nonnegative integers, $\mathbf{P}[X_k^{(n)} > 0] \leq \mathbf{E}X_k^{(n)}$. By the Borel-Cantelli lemma, $X_k^{(n)} > 0$ for only finitely many values of n .

Thus, for any fixed $r+1$, there exists a random time N such that there are no bad walks on G_n of length $r+1$ or less starting with vertex n , for $n \geq N$. We claim that for $n \geq N$, there are no bad walks at all on G_n with length r or less. Suppose that G_m contains some bad walk of length $k \leq r$, for some $m \geq N$. As the graph evolves, it is easy to compute that with probability one, a new vertex is eventually inserted into an edge of this walk. But at the time $n > m \geq N$ when this occurs, G_n will contain a bad walk of length $r+1$ or less starting with vertex n , a contradiction. Thus we have proven that G_n eventually contains no bad walks of length r or less. The equivalent statement for the continuous-time version of the graph process follows easily from this. \square

Proof of Theorem 5. Let $\text{CNBW}_k^{(s)}(t)$ denote the number of cyclically non-backtracking walks of length k in $G(s+t)$. We decompose these into those that are repeated walks around cycles of length j for some j dividing k , and the remaining bad walks, which we denote $B_k^{(s)}(t)$, giving us

$$\text{CNBW}_k^{(s)}(t) = \sum_{j|k} 2jC_j^{(s)}(t) + B_k^{(s)}(t).$$

Proposition 19 implies that

$$\lim_{s \rightarrow \infty} \mathbf{P}[B_k^{(s)}(t) = 0 \text{ for all } k \leq K, t \geq 0] = 1.$$

By Theorem 1 together with the continuous mapping theorem and Slutsky's theorem and, as s tends to infinity,

$$(14) \quad (\text{CNBW}_k^{(s)}(\cdot), 1 \leq k \leq K) \xrightarrow{\mathcal{L}} \left(\sum_{j|k} 2jN_j(\cdot), 1 \leq k \leq K \right).$$

Now, we modify the polynomials Γ_k to form a new basis $\{f_k, k \in \mathbb{N}\}$ with the right properties, which amounts to expressing each $N_k(t)$ as a linear combination of terms $\sum_{j|l} 2jN_j(t)$. We do this with the Möbius inversion formula. Define the polynomial

$$(15) \quad f_k(x) = \frac{1}{2k} \sum_{j|k} \mu\left(\frac{k}{j}\right) (2d-1)^{j/2} \Gamma_j(x),$$

where μ is the Möbius function, given by

$$\mu(n) = \begin{cases} (-1)^a & \text{if } n \text{ is the product of } a \text{ distinct primes,} \\ 0 & \text{otherwise.} \end{cases}$$

The theorem then follows from (13), (14), and the continuous mapping theorem. \square

Proof of Theorem 6. We start by recalling that, for any fixed d ,

$$2 \operatorname{tr} T_i(G(\infty + t)) + n \frac{2d-2}{(2d-1)^{i/2}} 1\{i \text{ is even}\} = (2d-1)^{-i/2} \sum_{k|i} 2kN_k(t).$$

Now, we will prove finite-dimensional convergence to the stated Ornstein-Uhlenbeck process. Consider two time points $s \leq t$ and two positive integers i, k . We will first show that, for any $i, k \in \mathbb{N}$, the pair $((2d-1)^{-i/2}(N_i(s) - \mathbf{E}N_i(s)), (2d-1)^{-k/2}(N_k(t) - \mathbf{E}N_k(t)))$ converges to a Gaussian limit as d tends to infinity. When $s = t$, this trivially follows via the Central Limit Theorem and their independent Poisson joint distribution.

When $s < t$, observe from (11) that

$$N_k(t) = \sum_{j=1}^k \alpha(j, k) N_j(s) + Z.$$

Here $\alpha(j, k)N_j(s)$, $j \in [k]$, and Z are independent Poisson random variables of various means. Moreover Z is independent of the history of the process till time s . Under the stationary law, the vector $(N_j(s), j \in \mathbb{N})$ are independent Poisson random variables. Thus, if $i > k$, then $N_i(s)$ is independent of $N_k(t)$. Otherwise, by the thinning property of Poisson, $\alpha(i, k)N_i(s)$ is independent of $(1 - \alpha(i, k))N_i(s)$. Therefore, $N_k(t) - \alpha(i, k)N_i(s)$, $\alpha(i, k)N_i(s)$, and $(1 - \alpha(i, k))N_i(s)$ are three independent Poisson random variable.

By the normal approximation to Poisson, we get the appropriate distributional convergence to corresponding independent Gaussian random variables. This shows the joint convergence of $(N_i(s), N_k(t))$ to Gaussian after centering and scaling.

A similar Gram-Schmidt orthogonalization can be carried out for the case of time points $t_1 \leq t_2 \leq \dots \leq t_m$ and corresponding positive integers j_1, j_2, \dots, j_m . This proves the joint Gaussian convergence of any finite collection of $(N_{j_i}(t_i), i \in [m])$ under centering and suitable scaling. Since the traces of Chebyshev polynomials are

linear combinations of coordinates of N , the joint Gaussian convergence extends to them by an argument invoking the Continuous Mapping and Slutsky's theorems.

For a fixed d , the covariance computation follows from Corollary 17 and (13). Hence, if $s < t$, then

$$(16) \quad \mathbf{Cov}(\mathrm{tr} T_i(G(\infty + t)), \mathrm{tr} T_j(G(\infty + s))) = \frac{1}{4} (2d - 1)^{-(i+j)/2} \sum_{k|i, l|j} 4lk \mathbf{Cov}(N_k(t), N_l(s))$$

Here

$$\mathbf{Cov}(N_k(t), N_l(s)) = \begin{cases} \frac{a(d,l)}{2l} \binom{k-1}{k-l} p^l (1-p)^{k-l}, & p = e^{s-t}, \quad \text{if } k \geq l. \\ 0, & \text{otherwise.} \end{cases}$$

We now fix any i, j, t, s and take d to infinity. Any term $a(d, r)$ is asymptotically the same as $(2d - 1)^r$. Thus the highest order term (in d) on the right side of (16) is $(2d - 1)^{\min(i,j)}$. Unless $i = j$, this term is negligible compared to $(2d - 1)^{(i+j)/2}$. This shows that the limiting covariance is zero unless $i = j$.

On the other hand, when $i = j$, every term on the right side of (16) vanishes, except when $k = i = l = j$. Hence,

$$\lim_{d \rightarrow \infty} \mathbf{Cov}(\mathrm{tr} T_i(G(\infty + t)), \mathrm{tr} T_i(G(\infty + s))) = \frac{1}{4} 2ip^i = \frac{i}{2} e^{i(s-t)}.$$

Finally we prove the process convergence. One simply needs to argue tightness. Fix a $K \in \mathbb{N}$ and, for every d , consider the process

$$\left(X_k(t) := (2d - 1)^{-k/2} (2kN_k(t) - a(d, k)), \quad k \in [K], t \geq 0 \right).$$

We claim that it suffices to show tightness for this process. This follows, since then, due to unequal scaling, the difference between this process and the centered and scaled traces go to zero in probability as d tends to infinity.

To show tightness of X , note that, by [EK86, Chapter 11, Problem 22 (c)], it suffices to show tightness for each of the individual one-dimensional processes

$$\{X_k, k \in \mathbb{N}\} \quad \text{and} \quad \{X_k + X_l, k, l \in \mathbb{N}\}.$$

For each of these one-dimensional processes, it is enough to show tightness in $D[0, T]$ (take $T = 1$, without loss of generality) which we will show using [Bil99, Theorem 13.5] for $\beta = 1, \alpha = 1$.

Consider X_k . We already know that it has finite-dimensional convergence to the Ornstein-Uhlenbeck process. Now, for any $r \leq s \leq t$, we want to estimate

$$\mathbf{E} \left[(X_k(s) - X_k(r))^2 (X_k(t) - X_k(s))^2 \right] = \frac{16k^4}{(2d - 1)^{2k}} \mathbf{E} \left[(N_k(s) - N_k(r))^2 (N_k(t) - N_k(s))^2 \right].$$

We claim that the above is $O(t - r)^2$ as d tends to infinity.

To see this, first use the elementary inequality $2ab \leq a^2 + b^2$ to obtain

$$\mathbf{E} \left[(X_k(s) - X_k(r))^2 (X_k(t) - X_k(s))^2 \right] \leq \frac{1}{2} \left[\mathbf{E} (X_k(s) - X_k(r))^4 + \mathbf{E} (X_k(t) - X_k(s))^4 \right].$$

Consider $\mathbf{E} (X_k(t) - X_k(s))^4$, which can be written as

$$(2d - 1)^{-2k} (2k)^4 \mathbf{E} (N_k(t) - N_k(s))^4.$$

From the orthogonal decomposition described in the beginning of this proof, one can write $N_k(t) - N_k(s)$ as the difference of two independent Poisson random variables

$Z_1 - Z_2$ with the same mean $\nu = (1 - r(k, k))a(d, k)/2k$, where $r(k, k) = \mathbf{E}\alpha(k, k)$. Thus

$$\begin{aligned} \mathbf{E}(N_k(t) - N_k(s))^4 &= \mathbf{E}[(Z_1 - \nu) - (Z_2 - \nu)]^4 \\ &= \mathbf{E}(Z_1 - \nu)^4 - 4\mathbf{E}(Z_1 - \nu)^3\mathbf{E}(Z_2 - \nu) + 6\mathbf{E}(Z_1 - \nu)^2\mathbf{E}(Z_2 - \nu)^2 \\ &\quad - 4\mathbf{E}(Z_1 - \nu)\mathbf{E}(Z_2 - \nu)^3 + \mathbf{E}(Z_2 - \nu)^4 \\ &= 2\nu(1 + 3\nu) + 6\nu^2 = 2\nu(1 + 6\nu). \end{aligned}$$

Since $\nu = (1 - e^{k(s-t)})a(d, k)/2k$, we get

$$(2d - 1)^{-2k}(2k)^4\mathbf{E}(N_k(t) - N_k(s))^4 \leq C'_k(1 - e^{k(s-t)})^2 \leq C_k(t - s)^2.$$

Here C_k, C'_k are constants depending on k .

Similarly $\mathbf{E}(X_k(s) - X_k(r))^4 \leq C_k(s - r)^2$. Combining the two parts we get $\mathbf{E}\left[(X_k(s) - X_k(r))^2(X_k(t) - X_k(s))^2\right] \leq C_k(t - r)^2$. This allows us to use [Bil99, Theorem 13.5, eqn. (13.14)] to show tightness of X_k .

A similar argument for the pairs $X_k + X_l$ proves our claim and completes the proof of the theorem. \square

APPENDIX: A BROAD POISSON APPROXIMATION RESULT

This appendix provides the proofs of Theorem 14 and Corollary 15. A less general version of Theorem 14 can be found in [DJPP12, Theorem 11]; we show in Corollary 24i how it follows from Theorem 14. Our theorem here also improves the total variation bound from $O((2d - 1)^{2r}/n)$ to $O((2d - 1)^{2r-1}/n)$. We conjecture that Theorem 14 is sharp.

As in the proof of Theorem 11 in [DJPP12], the main tool is the Stein-Chen method for Poisson approximation by size-biased couplings as described in [BHJ92], which uses the following idea: Recall the definition of $(I_\beta, \beta \in \mathcal{I})$ from Theorem 14. For each $\alpha \in \mathcal{I}$, let $(J_{\beta\alpha}, \beta \in \mathcal{I})$ be distributed as $(I_\beta, \beta \in \mathcal{I})$ conditioned on $I_\alpha = 1$. The goal is to construct a coupling of $(I_\beta, \beta \in \mathcal{I})$ and $(J_{\beta\alpha}, \beta \in \mathcal{I})$ so that the two random vectors are “close together”. We hope that for each $\alpha \in \mathcal{I}$, the cycles in $\mathcal{I} \setminus \{\alpha\}$ can be partitioned into two sets \mathcal{I}_α^- and \mathcal{I}_α^+ such that

$$(17) \quad J_{\beta\alpha} \leq I_\beta \quad \text{if } \beta \in \mathcal{I}_\alpha^-,$$

$$(18) \quad J_{\beta\alpha} \geq I_\beta \quad \text{if } \beta \in \mathcal{I}_\alpha^+.$$

If this is the case, then one can approximate $(I_\beta, \beta \in \mathcal{I})$ by a Poisson process by calculating $\mathbf{Cov}(I_\alpha, I_\beta)$ for every $\alpha, \beta \in \mathcal{I}$, according to the following proposition.

Proposition 20 (Corollary 10.B.1 in [BHJ92]). *Suppose that $\mathbf{I} = (I_\alpha, \alpha \in \mathcal{I})$ is a vector of 0-1 random variables with $\mathbf{E}I_\alpha = p_\alpha$. Suppose that $(J_{\beta\alpha}, \beta \in \mathcal{I})$ is distributed as described above, and that for each α there exists a partition and a coupling of $(J_{\beta\alpha}, \beta \in \mathcal{I})$ with $(I_\beta, \beta \in \mathcal{I})$ such that (17) and (18) are satisfied.*

Let $\mathbf{Y} = (Y_\alpha, \alpha \in \mathcal{I})$ be a vector of independent Poisson random variables with $\mathbf{E}Y_\alpha = p_\alpha$. Then

$$(19) \quad d_{TV}(\mathbf{I}, \mathbf{Y}) \leq \sum_{\alpha \in \mathcal{I}} p_\alpha^2 + \sum_{\alpha \in \mathcal{I}} \sum_{\beta \in \mathcal{I}_\alpha^-} |\mathbf{Cov}(I_\alpha, I_\beta)| + \sum_{\alpha \in \mathcal{I}} \sum_{\beta \in \mathcal{I}_\alpha^+} \mathbf{Cov}(I_\alpha, I_\beta).$$

We introduce two lemmas, whose proofs we will defer to the end of the appendix. The first will let us approximate \mathbf{I} by \mathbf{Z} rather than by \mathbf{Y} , and the second provides a technical bound that we need.

Lemma 21. *Let $\mathbf{Y} = (Y_\alpha, \alpha \in \mathcal{I})$ and $\mathbf{Z} = (Z_\alpha, \alpha \in \mathcal{I})$ be vectors of independent Poisson random variables. Then*

$$d_{TV}(\mathbf{Y}, \mathbf{Z}) \leq \sum_{\alpha \in \mathcal{I}} |\mathbf{E}Y_\alpha - \mathbf{E}Z_\alpha|.$$

Lemma 22. *Let a and b be d -dimensional vectors with nonnegative integer components, and let $\langle a, b \rangle$ denote the standard Euclidean inner product.*

$$\prod_{i=1}^d \frac{1}{[n]_{a_i+b_i}} - \prod_{i=1}^d \frac{1}{[n]_{a_i}[n]_{b_i}} \leq \frac{\langle a, b \rangle}{n} \prod_{i=1}^d \frac{1}{[n]_{a_i+b_i}}$$

Proof of Theorem 14. We will give the proof in three sections: First, we make the coupling and show that it satisfies (17) and (18). Next, we apply Proposition 20 to approximate \mathbf{I} by \mathbf{Y} , a vector of independent Poissons with $\mathbf{E}Y_\alpha = \mathbf{E}I_\alpha$. Last, we approximate \mathbf{Y} by \mathbf{Z} to prove the theorem.

If $d > n^{1/2}$ or $r > n^{1/10}$, then $c(2d-1)^{2r-1}/n > 1$ for a sufficiently large choice of c , and the theorem holds trivially. Thus we will assume throughout that $d \leq n^{1/2}$ and $r \leq n^{1/10}$ (the choice of $1/10$ here is completely arbitrary). The expression $O(f(d, r, n))$ should be interpreted as a function of d, r , and n whose absolute value is bounded by $Cf(d, r, n)$ for some absolute constant C , for all d, r , and n satisfying $2 \leq d \leq n^{1/2}$ and $r \leq n^{1/10}$.

Step 1. Constructing the coupling.

Fix some $\alpha \in \mathcal{I}$. We will construct a random vector $(J_{\beta\alpha}, \beta \in \mathcal{I})$ distributed as $(I_\beta, \beta \in \mathcal{I})$ conditioned on $I_\alpha = 1$. We do this by constructing a random graph G'_n distributed as G_n conditioned to contain the cycle α . Once this is done, we will define $J_{\beta\alpha} = 1\{G'_n \text{ contains cycle } \beta\}$.

Let π_1, \dots, π_d be the random permutations that give rise to G_n . We will alter them to form permutations π'_1, \dots, π'_d , and we will construct G'_n from these. Let us first consider what distributions π'_1, \dots, π'_d should have. For example, suppose that α is the cycle

$$1 \xrightarrow{\pi_3} 2 \xleftarrow{\pi_1} 3 \xrightarrow{\pi_3} 4 \xrightarrow{\pi_1} 1.$$

Then π'_1 should be distributed as a uniform random n -permutation conditioned to make $\pi'_1(3) = 2$ and $\pi'_1(4) = 1$, and π'_3 should be distributed as a uniform random n -permutation conditioned to make $\pi'_3(1) = 2$ and $\pi'_3(3) = 4$, while π'_2 should just be a uniform random n -permutation. A random graph constructed from π'_1, π'_2 , and π'_3 will be distributed as G_n conditioned to contain α .

We now describe the construction of π'_1, \dots, π'_d . Suppose α is the cycle

$$(20) \quad s_0 \xrightarrow{w_1} s_1 \xrightarrow{w_2} s_2 \xrightarrow{w_3} \dots \xrightarrow{w_k} s_k = s_0,$$

with each edge directed according to whether $w_i(s_{i-1}) = s_i$ or $w_i(s_i) = s_{i-1}$. Fix some $1 \leq l \leq d$, and suppose that the edge-label π_l appears M times in the cycle α . Let (a_m, b_m) for $1 \leq m \leq M$ be these directed edges. We must construct π'_l to have the uniform distribution conditioned on $\pi'_l(a_m) = b_m$ for $1 \leq m \leq M$.

We define a sequence of random transpositions by the following algorithm: Let τ_1 swap $\pi_l(a_1)$ with b_1 . Let τ_2 swap $\tau_1\pi_l(a_2)$ with b_2 , and so on. We then define $\pi'_l = \tau_M \cdots \tau_1\pi_l$. This permutation satisfies $\pi'_l(a_m) = b_m$ for $1 \leq m \leq M$, and it is distributed uniformly, subject to the given constraints, which can be proven by

induction on each swap. We now define G'_n from the permutations π'_1, \dots, π'_d in the usual way. It is defined on the same probability space as G_n , and it is distributed as G_n conditioned to contain α , giving us a random vector $(J_{\beta\alpha}, \beta \in \mathcal{I})$ coupled with $(I_\beta, \beta \in \mathcal{I})$.

Now, we will give a partition $\mathcal{I}^- \cup \mathcal{I}^+ = \mathcal{I} \setminus \{\alpha\}$ satisfying (17) and (18). Suppose that G_n contains an edge $s_i \xrightarrow{w_{i+1}} v$ with $v \neq s_{i+1}$, or an edge $v \xrightarrow{w_{i+1}} s_{i+1}$ with $v \neq s_i$. The graph G'_n cannot contain this edge, since it contains α . In fact, edges of this form are the *only* ones found in G_n but not G'_n :

Lemma 23. *Suppose there is an edge $i \xrightarrow{\pi_l} j$ contained in G_n but not in G'_n . Then α contains either an edge $i \xrightarrow{\pi_l} v$ with $v \neq j$, or α contains an edge $v \xrightarrow{\pi_l} j$ with $v \neq i$.*

Proof. Suppose $\pi_l(i) = j$, but $\pi'_l(i) \neq j$. Then j must have been swapped when making π'_l , which can happen only if $\pi_l(a_m) = j$ or $b_m = j$ for some m . In the first case, $a_m = i$ and α contains the edge $i \xrightarrow{\pi_l} b_m$ with $b_m \neq j$, and in the second α contains the edge $a_m \xrightarrow{\pi_l} j$ with $a_m \neq i$. \square

Define \mathcal{I}_α^- as all cycles in \mathcal{I} that contain an edge $s_i \xrightarrow{w_{i+1}} v$ with $v \neq s_{i+1}$ or an edge $v \xrightarrow{w_{i+1}} s_{i+1}$ with $v \neq s_i$, and define \mathcal{I}_α^+ to be the rest of $\mathcal{I} \setminus \{\alpha\}$. Since G'_n cannot contain any cycle in \mathcal{I}_α^- , we have $J_{\beta\alpha} = 0$ for all $\beta \in \mathcal{I}_\alpha^-$, satisfying (17). For any $\beta \in \mathcal{I}_\alpha^+$, Lemma 23 shows that if β appears in G_n , it must also appear in G'_n . Hence $J_{\beta\alpha} \geq I_\beta$, and (18) is satisfied.

Step 2. Approximation of \mathbf{I} by \mathbf{Y} .

The conditions of Proposition 20 are satisfied, and we need only bound the sums in (19). Let $p_\alpha = \mathbf{E}I_\alpha$, the probability that cycle α appears in G_n . Recall that this equals $\prod_{i=1}^d 1/[n]_{e_i}$, where e_i is the number of times π_i and π_i^{-1} appear in the word of α . This means that

$$(21) \quad \frac{1}{n^k} \leq p_\alpha \leq \frac{1}{[n]_k},$$

where $k = |\alpha|$, the length of cycle α .

We bound the first sum in (19) by

$$(22) \quad \begin{aligned} \sum_{\alpha \in \mathcal{I}} p_\alpha^2 &= \sum_{k=1}^r \sum_{\alpha \in \mathcal{I}_k} p_\alpha^2 \leq \sum_{k=1}^r \sum_{\alpha \in \mathcal{I}_k} \frac{1}{[n]_k^2} \\ &= \sum_{k=1}^r \left(\frac{[n]_k a(d, k)}{2k} \right) \left(\frac{1}{[n]_k^2} \right) \\ &\leq \sum_{k=1}^r \frac{2d(2d-1)^{k-1}}{2k[n]_k} = O\left(\frac{d}{n}\right). \end{aligned}$$

To bound the second sum in (19), we investigate the size of \mathcal{I}_α^- . Suppose that $\alpha \in \mathcal{I}_k$, and α has the form given in (20). Any $\beta \in \mathcal{I}_\alpha^-$ must contain an edge $s_i \xrightarrow{w_{i+1}} v$ with $v \neq s_{i+1}$, or an edge $v \xrightarrow{w_{i+1}} s_{i+1}$ with $v \neq s_i$, and there are at most $2k(n-1)$ edges of this form. For any given edge, there are at most $[n-2]_{j-2}(2d-1)^{j-1}$ cycles in \mathcal{I}_j that contain that edge, for any $j \geq 2$. Thus for any $\alpha \in \mathcal{I}_k$, the number of cycles of length $j \geq 2$ in \mathcal{I}_α^- is at most $2k[n-1]_{j-1}(2d-1)^{j-1}$, and this bound also holds for $j = 1$.

For any $\beta \in \mathcal{I}_\alpha^-$, it holds that $\mathbf{E}[I_\alpha I_\beta] = 0$, so that $\mathbf{Cov}(I_\alpha, I_\beta) = -p_\alpha p_\beta$. Putting this all together and applying (21), we have

$$\begin{aligned}
\sum_{\alpha \in \mathcal{I}} \sum_{\beta \in \mathcal{I}_\alpha^-} |\mathbf{Cov}(I_\alpha, I_\beta)| &= \sum_{k=1}^r \sum_{\alpha \in \mathcal{I}_k} \sum_{j=1}^r \sum_{\beta \in \mathcal{I}_\alpha^- \cap \mathcal{I}_j} p_\alpha p_\beta \\
&\leq \sum_{k=1}^r |\mathcal{I}_k| \frac{1}{[n]_k} \sum_{j=1}^r |\mathcal{I}_\alpha^- \cap \mathcal{I}_j| \frac{1}{[n]_j} \\
&\leq \sum_{k=1}^r \frac{a(d, k)}{2k} \sum_{j=1}^r \frac{2k(2d-1)^{j-1}}{n} \\
(23) \qquad \qquad \qquad &= \sum_{k=1}^r a(d, k) O\left(\frac{(2d-1)^{r-1}}{n}\right) = O\left(\frac{(2d-1)^{2r-1}}{n}\right).
\end{aligned}$$

The final sum in (19) is the most difficult to bound. We partition \mathcal{I}_α^+ into sets $\mathcal{I}_\alpha^+ = \mathcal{I}_\alpha^0 \cup \dots \cup \mathcal{I}_\alpha^{|\alpha|-1}$, where \mathcal{I}_α^l is all cycles in \mathcal{I}_α^+ that share exactly l labeled edges with α . For any $\beta \in \mathcal{I}_\alpha^+$,

$$\mathbf{E}[I_\alpha I_\beta] = \mathbf{P}[G \text{ contains } \alpha \text{ and } \beta] = \prod_{i=1}^d \frac{1}{[n]_{e_i}},$$

where e_i is the number of π_i -labeled edges in $\alpha \cup \beta$. Thus for $\beta \in \mathcal{I}_\alpha^l$,

$$(24) \qquad \qquad \qquad \frac{1}{n^{|\alpha|+|\beta|-l}} \leq \mathbf{E}[I_\alpha I_\beta] \leq \frac{1}{[n]_{|\alpha|+|\beta|-l}}.$$


We start by seeking estimates on the size of \mathcal{I}_α^l for $l \geq 1$. Fix some choice of l edges of α . We start by counting the cycles in \mathcal{I}_α^l that share exactly these edges with α . We illustrate this in Figure 5. Call the graph consisting of these edges H , and suppose that H has p components. Since it is a forest, H has $l + p$ vertices.

Let A_1, \dots, A_p be the components of H . We can assemble any element $\beta \in \mathcal{I}_\alpha^l$ that overlaps with α in H by stringing together these components in some order, with other edges in between. Each component can appear in β in one of two orientations. Since the vertices in β have no fixed ordering, we can assume without loss of generality that β begins with component A_1 with a fixed orientation. This leaves $(p-1)!2^{p-1}$ choices for the order and orientation of A_2, \dots, A_p in β .


Imagine now the components laid out in a line, with gaps between them, and count the number of ways to fill the gaps. Suppose that β is to have length j . Each of the p gaps must contain at least one edge, and the total number of edges in all the gaps is $j - l$. Thus the total number of possible gap sizes is the number of compositions of $j - l$ into p parts, or $\binom{j-l-1}{p-1}$.

Now that we have chosen the number of edges to appear in each gap, we choose the edges themselves. We can do this by giving an ordered list $j - p - l$ vertices to go in the gaps, along with a label and an orientation for each of the $j - l$ edges this gives. There are $[n - p - l]_{j-p-l}$ ways to choose the vertices. We can give each new edge any orientation and label subject to the constraint that the word of the cycle we construct must be reduced. This means we have at most $2d - 1$ choices for the orientation and label of each new edge, for a total of at most $(2d - 1)^{j-l}$.


All together, there are at most $(p-1)!2^{p-1} \binom{j-l-1}{p-1} [n - p - l]_{j-p-l} (2d - 1)^{j-l}$ elements of \mathcal{I}_j that overlap with the cycle α at the subgraph H . We now calculate




The cycle α , with H dashed.
 The subgraph H has components A_1, \dots, A_p . In this example, the number of components of H is $p = 3$, the size of α is $k = 11$, and the number of edges in H is $l = 4$.
 In this example, we will construct a cycle β of length $j = 10$ that overlaps with α at H .



Step 1. We lay out the components A_1, \dots, A_p . We can order and orient A_2, \dots, A_p however we would like, for a total of $(p - 1)!2^{p-1}$ choices. Here, we have ordered the components A_1, A_3, A_2 , and we have reversed the orientation of A_3 .



Step 2. Next, we choose how many edges will go in each gap between components. Each gap must contain at least one edge, and we must add a total of $j - l$ edges, giving us $\binom{j-l-1}{p-1}$ choices. In this example, we have added one edge after A_1 , three after A_3 , and two after A_2 .



Step 3. We can choose the new vertices in $[n - p - l]_{j-p-l}$ ways, and we can direct and give labels to the new edges in at most $(2d - 1)^{j-l}$ ways.

FIGURE 5. Assembling an element $\beta \in \mathcal{T}_\alpha^l$ that overlaps with α at a given subgraph H .

the number of different ways to choose a subgraph H of α with l edges and p components. Suppose α is given as in (20). We first choose a vertex s_{i_0} . Then, we can specify which edges to include in H by giving a sequence $a_1, b_1, \dots, a_p, b_p$ instructing us to include in H the first a_1 edges after s_{i_0} , then to exclude the next b_1 , then to include the next a_2 , and so on. Any sequence for which a_i and b_i are positive integers, $a_1 + \dots + a_p = l$, and $b_1 + \dots + b_p = k - l$ gives us a valid choice of l edges of α making up p components. This counts each subgraph H a total of p times, since we could begin with any component of H . Hence the number of subgraphs H with l edges and p components is $(k/p) \binom{l-1}{p-1} \binom{k-l-1}{p-1}$. This gives us

the bound

$$|\mathcal{I}_\alpha^l \cap \mathcal{I}_j| \leq \sum_{p=1}^{l \wedge (j-l)} \binom{k/p}{p-1} \binom{l-1}{p-1} \binom{k-l-1}{p-1} (p-1)! \times \\ 2^{p-1} \binom{j-l-1}{p-1} [n-p-l]_{j-p-l} (2d-1)^{j-l}.$$

We apply the bounds

$$\binom{l-1}{p-1} \leq \frac{r^{p-1}}{(p-1)!}, \\ \binom{k-l-1}{p-1}, \binom{j-l-1}{p-1} \leq (er/(p-1))^{p-1},$$

to get

$$|\mathcal{I}_\alpha^l \cap \mathcal{I}_j| \leq k(2d-1)^{j-l} [n-1-l]_{j-1-l} \left(1 + \sum_{p=2}^{i \wedge (k-i)} \frac{1}{p} \left(\frac{2e^2 r^3}{(p-1)^2} \right)^{p-1} \frac{1}{[n-1-l]_{p-1}} \right).$$

Since $r \leq n^{1/10}$, the sum in the above equation is bounded by an absolute constant. Applying this bound and (24), for any $\alpha \in \mathcal{I}_k$ and $l \geq 1$,

$$\sum_{\beta \in \mathcal{I}_\alpha^l} \mathbf{Cov}(I_\alpha, I_\beta) \leq \sum_{j=l+1}^r \sum_{\beta \in \mathcal{I}_\alpha^l \cap \mathcal{I}_j} \frac{1}{[n]_{k+j-l}} \\ \leq \sum_{j=l+1}^r O\left(\frac{k(2d-1)^{j-l}}{n^{k+1}}\right) \\ = O\left(\frac{k(2d-1)^{r-l}}{n^{k+1}}\right).$$

Therefore

$$\sum_{\alpha \in \mathcal{I}} \sum_{l \geq 1} \sum_{\beta \in \mathcal{I}_\alpha^l} \mathbf{Cov}(I_\alpha, I_\beta) = \sum_{k=1}^r \sum_{\alpha \in \mathcal{I}_k} \sum_{l=1}^{k-1} \sum_{\beta \in \mathcal{I}_\alpha^l} \mathbf{Cov}(I_\alpha, I_\beta) \\ \leq \sum_{k=1}^r \sum_{\alpha \in \mathcal{I}_k} \sum_{l=1}^{k-1} O\left(\frac{k(2d-1)^{r-l}}{n^{k+1}}\right) \\ = \sum_{k=1}^r \frac{[n]_k a(d, k)}{2k} O\left(\frac{k(2d-1)^{r-1}}{n^{k+1}}\right) \\ = \sum_{k=1}^r O\left(\frac{(2d-1)^{r+k-1}}{n}\right) \\ (25) \quad = O\left(\frac{(2d-1)^{2r-1}}{n}\right).$$

Last, we must bound $\sum_{\alpha \in \mathcal{I}} \sum_{\beta \in \mathcal{I}_\alpha^0} \mathbf{Cov}(I_\alpha, I_\beta)$. For any word w , let e_i^w be the number of appearances of π_i and π_i^{-1} in w . Let α and β be cycles with words w

and u respectively, and let $k = |\alpha|$ and $j = |\beta|$. Suppose that $\beta \in \mathcal{I}_\alpha^0$. Then

$$\begin{aligned} \mathbf{Cov}(I_\alpha, I_\beta) &= \prod_{i=1}^d \frac{1}{[n]_{e_i^w + e_i^u}} - \prod_{i=1}^d \frac{1}{[n]_{e_i^w} [n]_{e_i^u}} \\ &\leq \frac{\langle e^w, e^u \rangle}{n} \prod_{i=1}^d \frac{1}{[n]_{e_i^w + e_i^u}} \leq \frac{\langle e^w, e^u \rangle}{n[n]_{k+j}} \end{aligned}$$

by Lemma 22. For any pair of words $w \in \mathcal{W}_k$ and $u \in \mathcal{W}_j$, there are at most $[n]_k [n]_j$ pairs of cycles $\alpha, \beta \in \mathcal{I}$ with words w and u , respectively. Enumerating over all $w \in \mathcal{W}_k$ and $u \in \mathcal{W}_j$, we count each pair of cycles α, β exactly $4kj$ times. Thus

$$\begin{aligned} \sum_{\alpha \in \mathcal{I}_k} \sum_{\beta \in \mathcal{I}_\alpha^0 \cap \mathcal{I}_j} \mathbf{Cov}(I_\alpha, I_\beta) &\leq \frac{[n]_k [n]_j}{4kj n [n]_{k+j}} \sum_{w \in \mathcal{W}_k} \sum_{u \in \mathcal{W}_j} \langle e^w, e^u \rangle \\ &\leq \frac{1 + O(r^2/n)}{4kj n} \left\langle \sum_{w \in \mathcal{W}_k} e^w, \sum_{u \in \mathcal{W}_j} e^u \right\rangle. \end{aligned}$$

The vector $\sum_{w \in \mathcal{W}_k} e^w$ has every entry equal by symmetry, as does $\sum_{u \in \mathcal{W}_j} e^u$. Thus each entry of $\sum_{w \in \mathcal{W}_k} e^w$ is $ka(d, k)/d$, and each entry of $\sum_{u \in \mathcal{W}_j} e^u$ is $ja(d, j)/d$. The inner product in the above equation comes to $kja(d, k)a(d, j)/d$, giving us

$$\begin{aligned} \sum_{\alpha \in \mathcal{I}_k} \sum_{\beta \in \mathcal{I}_\alpha^0 \cap \mathcal{I}_j} \mathbf{Cov}(I_\alpha, I_\beta) &\leq \frac{a(d, k)a(d, j)(1 + O(r^2/n))}{4dn} \\ &= O\left(\frac{(2d-1)^{j+k-1}}{n}\right). \end{aligned}$$

Summing over all $1 \leq k, j \leq r$,

$$(26) \quad \sum_{\alpha \in \mathcal{I}} \sum_{\beta \in \mathcal{I}_\alpha^0} \mathbf{Cov}(I_\alpha, I_\beta) = \left(\frac{(2d-1)^{2r-1}}{n}\right).$$

We can now combine equations (22), (23), (25), and (26) with Proposition 20 to show that

$$(27) \quad d_{TV}(\mathbf{I}, \mathbf{Y}) = O\left(\frac{(2d-1)^{2r-1}}{n}\right).$$

Step 3. Approximation of \mathbf{Y} by \mathbf{Z} .

By Lemma 21 and (21),

$$\begin{aligned} d_{TV}(\mathbf{Y}, \mathbf{Z}) &\leq \sum_{\alpha \in \mathcal{I}} |\mathbf{E}Y_\alpha - \mathbf{E}Z_\alpha| \leq \sum_{k=1}^r \sum_{\alpha \in \mathcal{I}_k} \left(\frac{1}{[n]_k} - \frac{1}{n^k}\right) \\ &= \sum_{k=1}^r \frac{a(d, k)}{2k} \left(1 - \frac{[n]_k}{n^k}\right). \end{aligned}$$

Since $[n]_k \geq n^k(1 - k^2/2n)$,

$$d_{TV}(\mathbf{Y}, \mathbf{Z}) \leq \sum_{k=1}^r \frac{a(d, k)k}{4n} = O\left(\frac{r(2d-1)^r}{n}\right).$$

Together with (27), this bounds the total variation distance between the laws of \mathbf{I} and \mathbf{Z} and proves the theorem. \square

The distributions of any functionals of \mathbf{I} and \mathbf{Z} satisfy the same bound in total variation distance. This gives us several results as easy corollaries, including an improvement on [DJPP12, Theorem 11].

Corollary 24.

- i) Let $(Z_k, 1 \leq k \leq r)$ be a vector of independent Poisson random variables with $\mathbf{E}Z_k = a(d, k)/2k$. Let C_k denote the number of k -cycles in G_n , a $2d$ -regular permutation random graph on n vertices. Then for some absolute constant c ,

$$d_{TV}((C_k, 1 \leq k \leq r), (Z_k, 1 \leq k \leq r)) \leq \frac{c(2d-1)^{2r-1}}{n}.$$

- ii) Let $(Z_w, w \in \mathcal{W}'_K)$ be a vector of independent Poisson random variables with $\mathbf{E}Z_w = 1/h(w)$. Let C_w denote the number of cycles with word w in G_n , a $2d$ -regular permutation random graph on n vertices. Then for some absolute constant c ,

$$d_{TV}((C_w, w \in \mathcal{W}'_K), (Z_w, w \in \mathcal{W}'_K)) \leq \frac{c(2d-1)^{2K-1}}{n}.$$

Proof. Observe that $C_k = \sum_{\alpha \in \mathcal{I}_k} I_\alpha$, and that if we define $Z_k = \sum_{\alpha \in \mathcal{I}_k} Z_\alpha$, then $(Z_k, 1 \leq k \leq r)$ is distributed as described. Thus (i) follows from Theorem 14.

To prove (ii), note that $C_w = \sum_{\alpha} I_\alpha$, where the sum is over all cycles in \mathcal{I} with word w . We then define Z_w as the analogous sum over Z_α . Since the number of cycles in \mathcal{I} with word w is $[n]_k/h(w)$, we have $\mathbf{E}Z_w = 1/h(w)$, and the total variation bound follows from Theorem 14. \square

We can also use Theorem 14 to bound the likelihood that G_n contains two overlapping cycles of size r or less.

Corollary 25. Let G_n be a $2d$ -regular permutation random graph on n vertices. Let E be the event that G_n contains two cycles of length r or less with a vertex in common. Then for some absolute constant c' , for all $d \geq 2$ and $n, r \geq 1$,

$$\mathbf{P}[E] \leq \frac{c'(2d-1)^{2r}}{n}.$$

Proof. Let E' be the event that $Z_\alpha = Z_\beta = 1$ for two cycles $\alpha, \beta \in \mathcal{I}$ that have a vertex in common. By Theorem 14,

$$\mathbf{P}[E] \leq \mathbf{P}[E'] + \frac{c(2d-1)^{2r-1}}{n}.$$

For any cycle $\alpha \in \mathcal{I}_k$, there are at most $k[n-1]_{j-1}a(d, j)$ cycles in \mathcal{I}_j that share a vertex with α . For any such cycle β , the chance that $Z_\alpha = 1$ and $Z_\beta = 1$ is less than $1/[n]_k[n]_j$. By a union bound,

$$\begin{aligned} \mathbf{P}[E'] &\leq \sum_{k=1}^r \frac{a(d, k)[n]_k}{2k} \sum_{j=1}^r \frac{k[n-1]_{j-1}a(d, j)}{[n]_k[n]_j} \\ &\leq \sum_{k=1}^r \sum_{j=1}^r \frac{a(d, k)a(d, j)}{2n} = O\left(\frac{(2d-1)^{2r}}{n}\right). \end{aligned} \quad \square$$

Proof of Corollary 15. When $d = 1$, there is only one word of each length in \mathcal{W}'_K , and statement (i) reduces to the well-known fact that the cycle counts of a random permutation converge to independent Poisson random variables (see [AT92] for much more on this subject). In this case, $G(t)$ is made up of disjoint cycles for all times t , so that statement (ii) is trivially satisfied.

When $d \geq 2$, let $C_w^{(n)}$ be the number of cycles with word w in G_n , as in Corollary 24ii. The random vector $(C_w(t), w \in \mathcal{W}'_K)$ is a mixture of the random vectors $(C_w^{(n)}, w \in \mathcal{W}'_K)$ over different values of n . That is,

$$\mathbf{P}[(C_w(t), w \in \mathcal{W}'_K) \in A] = \sum_{n=1}^{\infty} \mathbf{P}[M_t = n] \mathbf{P}[(C_w^{(n)}, w \in \mathcal{W}'_K) \in A]$$

for any set A , recalling that $G(t) = G_{M_t}$. Corollary 24ii together with the fact that $\mathbf{P}[M_t > N] \rightarrow 1$ as $t \rightarrow \infty$ for any N imply that $(C_w(t), w \in \mathcal{W}'_K)$ converges in law to $(Z_w, w \in \mathcal{W}'_K)$, establishing statement (i). Statement (ii) follows in the same way from Corollary 25. \square

Proof of Lemma 21. We will apply the Stein-Chen method directly. Define the operator \mathcal{A} by

$$\mathcal{A}h(x) = \sum_{\alpha \in \mathcal{I}} \mathbf{E}[Z_\alpha] (h(x + e_\alpha) - h(x)) + \sum_{\alpha \in \mathcal{I}} x_\alpha (h(x - e_\alpha) - h(x))$$

for any $h: \mathbb{Z}_+^{|\mathcal{I}|} \rightarrow \mathbb{R}$ and $x \in \mathbb{Z}_+^{|\mathcal{I}|}$. This is the Stein operator for the law of \mathbf{Z} , and $\mathbf{E}\mathcal{A}h(\mathbf{Z}) = 0$ for any bounded function h . By Proposition 10.1.2 and Lemma 10.1.3 in [BHJ92], for any set $A \subseteq \mathbb{Z}_+^{|\mathcal{I}|}$, there is a function h such that

$$\mathcal{A}h(x) = 1\{x \in A\} - \mathbf{P}[\mathbf{Z} \in A],$$

and this function has the property that

$$(28) \quad \sup_{\substack{x \in \mathbb{Z}_+^{|\mathcal{I}|} \\ \alpha \in \mathcal{I}}} |h(x + e_\alpha) - h(x)| \leq 1.$$

Thus we can bound the total variation distance between the laws of \mathbf{Y} and \mathbf{Z} by bounding $|\mathbf{E}\mathcal{A}h(\mathbf{Y})|$ over all such functions h .

We write $\mathcal{A}h(\mathbf{Y})$ as

$$\begin{aligned} \mathcal{A}h(\mathbf{Y}) &= \sum_{\alpha \in \mathcal{I}} \mathbf{E}[Y_\alpha] (h(\mathbf{Y} + e_\alpha) - h(\mathbf{Y})) + \sum_{\alpha \in \mathcal{I}} Y_\alpha (h(\mathbf{Y} - e_\alpha) - h(\mathbf{Y})) \\ &\quad + \sum_{\alpha \in \mathcal{I}} (\mathbf{E}Z_\alpha - \mathbf{E}Y_\alpha) (h(\mathbf{Y} + e_\alpha) - h(\mathbf{Y})). \end{aligned}$$

The first two of these sums have expectation zero, so

$$|\mathbf{E}\mathcal{A}h(\mathbf{Y})| \leq \sum_{\alpha \in \mathcal{I}} |\mathbf{E}Z_\alpha - \mathbf{E}Y_\alpha| \mathbf{E}|h(\mathbf{Y} + e_\alpha) - h(\mathbf{Y})|.$$

By (28), $|h(\mathbf{Y} + e_\alpha) - h(\mathbf{Y})| \leq 1$, which proves the lemma. \square

Proof of Lemma 22. We define a family of independent random maps σ_i and τ_i for $1 \leq i \leq d$. Choose σ_i uniformly from all injective maps from $[a_i]$ to $[n]$, and choose

τ_i uniformly from all injective maps from $[b_i]$ to $[n]$. Effectively, σ_i and τ_i are random ordered subsets of $[n]$. We say that σ_i and τ_i *clash* if their images overlap.

$$\mathbf{P}[\sigma_i \text{ and } \tau_i \text{ clash for some } i] = 1 - \prod_{i=1}^d \frac{[n]_{a_i+b_i}}{[n]_{a_i}[n]_{b_i}}.$$

For any $1 \leq i \leq d$, $1 \leq j \leq a_i$, and $1 \leq k \leq b_i$, the probability that $\sigma_i(j) = \tau_i(k)$ is $1/n$. By a union bound,

$$\mathbf{P}[\sigma_i \text{ and } \tau_i \text{ clash for some } i] \leq \sum_{i=1}^d \frac{a_i b_i}{n} = \frac{\langle a, b \rangle}{n}.$$

We finish the proof by dividing both sides of this inequality by $\prod_{i=1}^d [n]_{a_i+b_i}$. \square

REFERENCES

- [ANvM11] Mark Adler, Eric Nordenstam, and Pierre van Moerbeke. The Dyson Brownian minor process. Preprint. Available at arXiv:1006.2956, 2011.
- [AT92] Richard Arratia and Simon Tavaré. The cycle structure of random permutations. *Ann. Probab.*, 20(3):1567–1591, 1992.
- [BAD11] Gérard Ben Arous and Kim Dang. On fluctuations of eigenvalues of random permutation matrices. 2011. Preprint. Available at arXiv:1106.2108.
- [BF08] Alexei Borodin and Patrik Ferrari. Anisotropic growth of random surfaces in 2+1 dimensions. Preprint. Available at arXiv:0804.3035, 2008.
- [BHJ92] A. D. Barbour, Lars Holst, and Svante Janson. *Poisson approximation*, volume 2 of *Oxford Studies in Probability*. The Clarendon Press Oxford University Press, New York, 1992. Oxford Science Publications.
- [Bil99] Patrick Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1999. A Wiley-Interscience Publication.
- [BNN11] Paul Bourgade, Joseph Najnudel, and Ashkan Nikeghbali. A unitary extension of virtual permutations. Preprint. Available at arXiv:1102.2633, 2011.
- [Bor10a] Alexei Borodin. CLT for spectra of submatrices of Wigner random matrices. Preprint. Available at arXiv:1010.0898, 2010.
- [Bor10b] Alexei Borodin. CLT for spectra of submatrices of Wigner random matrices II. Stochastic evolution. Preprint. Available at arXiv:1011.3544, 2010.
- [CPY98] Philippe Carmona, Frédérique Petit, and Marc Yor. Beta-gamma random variables and intertwining relations between certain Markov processes. *Rev. Mat. Iberoamericana*, 14(2):311–367, 1998.
- [DF90] Persi Diaconis and James Allen Fill. Strong stationary times via a new form of duality. *Ann. Probab.*, 18(4):1483–1522, 1990.
- [DJPP12] Ioana Dumitriu, Tobias Johnson, Soumik Pal, and Elliot Paquette. Functional limit theorems for random regular graphs. *Probab. Theory Related Fields*, pages 1–55, 2012. Published online, 25 August 2012.
- [DP12] Ioana Dumitriu and Soumik Pal. Sparse regular random graphs: Spectral density and eigenvectors. *Ann. Probab.*, 40(5):2197–2235, 2012.
- [EK86] Stewart N. Ethier and Thomas G. Kurtz. *Markov processes*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1986. Characterization and convergence.
- [Elo08] Yehonatan Elon. Eigenvectors of the discrete Laplacian on regular graphs—a statistical approach. *J. Phys. A*, 41(43):435203, 17, 2008.
- [Elo10] Yehonatan Elon. Gaussian waves on the regular tree. Preprint. Available at arXiv:0907.5065, 2010.
- [ES10] Yehonatan Elon and Uzy Smilansky. Percolating level sets of the adjacency eigenvectors of d -regular graphs. *J. Phys. A*, 43(45):455209, 13, 2010.
- [Fer10] Patrik L. Ferrari. From interacting particle systems to random matrices. *J. Stat. Mech. Theory Exp.*, (10):P10016, 15, 2010.

- [FF10] Patrik L. Ferrari and René Frings. On the partial connection between random matrices and interacting particle systems. *J. Stat. Phys.*, 141(4):613–637, 2010.
- [JMRR99] Dmitry Jakobson, Stephen D. Miller, Igor Rivin, and Zeév Rudnick. Eigenvalue spacings for regular graphs. In *Emerging applications of number theory (Minneapolis, MN, 1996)*, volume 109 of *IMA Vol. Math. Appl.*, pages 317–327. Springer, New York, 1999.
- [JN06] Kurt Johansson and Eric Nordenstam. Eigenvalues of GUE minors. *Electron. J. Probab.*, 11:no. 50, 1342–1371, 2006.
- [KOV04] Sergei Kerov, Grigori Olshanski, and Anatoly Vershik. Harmonic analysis on the infinite symmetric group. *Invent. Math.*, 158(3):551–642, 2004.
- [LP10] Nati Linial and Doron Puder. Word maps and spectra of random graph lifts. *Random Structures Algorithms*, 37(1):100–135, 2010.
- [MNS08] Steven J. Miller, Tim Novikoff, and Anthony Sabelli. The distribution of the largest nontrivial eigenvalues in families of random regular graphs. *Experiment. Math.*, 17(2):231–244, 2008.
- [OGS09] Idan Oren, Amit Godel, and Uzy Smilansky. Trace formulae and spectral statistics for discrete Laplacians on regular graphs. I. *J. Phys. A*, 42(41):415101, 20, 2009.
- [OS10] Idan Oren and Uzy Smilansky. Trace formulas and spectral statistics for discrete Laplacians on regular graphs (II). *J. Phys. A*, 43(22):225205, 13, 2010.
- [Pit06] Jim Pitman. *Combinatorial stochastic processes*, volume 1875 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2006. Lectures from the 32nd Summer School on Probability Theory held in Saint-Flour, July 7–24, 2002, With a foreword by Jean Picard.
- [She07] Scott Sheffield. Gaussian free fields for mathematicians. *Probability Theory and Related Fields*, 139(3-4):521–541, 2007.
- [Smi10] Uzy Smilansky. Discrete graphs - a paradigm model for quantum chaos. In *Séminaire Poincaré*, volume XIV, pages 1–26. 2010.
- [Spo98] Herbert Spohn. Dyson’s model of interacting brownian motions at arbitrary coupling strength. In *MARKOV PROC. REL. FIELDS*, pages 649–661, 1998.
- [TVW13] Linh V. Tran, Van H. Vu, and Ke Wang. Sparse random graphs: Eigenvalues and eigenvectors. *Random Structures Algorithms*, 42(1):110–134, 2013.
- [Wie00] Kelly Wieand. Eigenvalue distributions of random permutation matrices. *Ann. Probab.*, 28(4):1563–1587, 2000.
- [Wor99] Nicholas C. Wormald. Models of random regular graphs. In *Surveys in combinatorics, 1999 (Canterbury)*, volume 267 of *London Math. Soc. Lecture Note Ser.*, pages 239–298. Cambridge Univ. Press, Cambridge, 1999.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
E-mail address: toby@math.washington.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
E-mail address: soumik@u.washington.edu