Ramanujan graphings and correlation decay in local algorithms

Ágnes Backhausz   Balázs Szegedy   Bálint Virág

May 30, 2013

Abstract

Let G be a large-girth d-regular graph and μ be a random process on the vertices of G produced by a randomized local algorithm. We prove the upper bound $(k + 1 - 2k/d)\left(\frac{1}{\sqrt{d-1}}\right)^k$ for the (absolute value of the) correlation of values on pairs of vertices of distance k and show that this bound is optimal. The same results hold automatically for factor of i.i.d processes on the d-regular tree. In that case we give an explicit description for the (closure) of all possible correlation sequences. Our proof is based on the fact that the Bernoulli graphing of the infinite d-regular tree has spectral radius $2\sqrt{d-1}$. Graphings with this spectral gap are infinite analogues of finite Ramanujan graphs and they are interesting on their own right.

1   Introduction

Randomized local algorithms are special type of parallelized algorithms that can be used to produce various important structures in graphs (independent sets, dominating sets, matchings, colorings, local samples, etc.) in constant running time (see [4],[5],[6], [8], [9], [10], [13], [16]).

Let $d \in \mathbb{N}$ be a fixed number. The input of the algorithm is a graph G of maximal degree at most d. The first step of the algorithm puts labels on the vertices of G independently from a probabiliy space Ω. The second step evaluates a function f (rule of the algorithm) at each vertex v that depends on the isomorphism type of the labeled neighborhood of v of radius r. (If Ω is an infinite probability space then f is assumed to be measurable.)

In this paper we focus on the case of large-girth d-regular graphs. We give a rather explicit description for the (pointwise closure) of all possible correlation sequences in a local algorithm. Moreover, we obtain that the absolute value of the correlation of the values on two vertices of distance k is at most $(k + 1 - 2k/d)(d - 1)^{-k/2}$ provided that the girth of the graph is at least $k + 2r + 2$. Surprisingly, the bound itself does not depend on the radius r which is related to the complexity of the algorithm. Furthermore, we show that our upper bound is essentially optimal.

In our proof we make use of infinite measurable graphs called *graphings*. It turns out that many properties of local algorithms on large-girth d-regular graphs can be studied through the properties of a single object B_d called *Bernoulli graphing*. For example our result on the correlation decay relies on the fact that B_d is a *Ramanujan graphing*. Ramanujan graphs were introduced in the seminal paper [14] by Lubotzky, Phillips and Sarnak. A d-regular graph is Ramanujan if its second largest eigenvalue (in the absolute value) is at most $2\sqrt{d-1}$. The analogue of the second largest eigenvalue can easily be defined for graphings by the spectral radius on the orthogonal complement of the constant function. In the infinite case $2\sqrt{d-1}$ is the smallest possible value. (Note (see [12]) that this value is also equal to the spectral radius of the adjacency operator of the infinite d-regular tree.) Ramanujan graphings (and more generally graphings with spectral gap) are interesting on their own right. (Note that M. Abért has a different, more general definition for the notion of a Ramanujan graphing. [1])

Last but not least, our results can also be interpreted from an ergodic theoretic point of view. As we will see in Section 3, d-regular graphings can be used to produce random processes on the d-regular infinite tree T_d that are invariant under the full automorphism group. Conversely, automorphism invariant processes can be used to produce d-regular graphings. Processes that come from the Bernoulli graphing are usually called *factor of i.i.d* processes [15]. Our result on the possible correlation sequences and the correlation decay naturally generalizes to "Ramanujan processes" (i.e., processes that come from Ramanujan graphings).

To broaden the view on the topic, the last section discusses an interesting connection to the representation theory of $\mathrm{Aut}(T_d)$. There is a triple correspondence between spherical representations, correlation sequences, and invariant Gaussian processes on T_d.

2 Ramanujan and Bernoulli graphings

Definition 2.1 *Let X be a Polish topological space and let ν be a probability measure on the Borel sets in X. A* graphing *is a graph $\mathcal{G}$ on $V(\mathcal{G}) = X$ with bounded maximal degree and Borel measurable edge set $E(\mathcal{G}) \subset X \times X$ such that*

$$\int_A e(x, B)d\nu(x) = \int_B e(x, A)d\nu(x) \tag{1}$$

for all measurable sets $A, B \subseteq X$, where $e(x, S)$ is the number of edges from $x \in X$ to $S \subseteq X$.

---

[1]Personal communication.

Note that finite graphs are special graphings defined on finite probability spaces with uniform distribution. Let $\mathcal{G}$ be as in Definition 2.1. If $f : X \to \mathbb{C}$ is a measurable function then we define $\mathcal{G}f$ by

$$\mathcal{G}f(x) = \sum_{(x,v) \in E(\mathcal{G})} f(v). \tag{2}$$

A short calculation shows (see [10]) that $\mathcal{G}$ is a self-adjoint operator on $L^2(X)$ of norm at most d where d is the maximal degree in $\mathcal{G}$.

To keep our notation simple, in this paper we will only consider d-regular graphings (every vertex has degree d). In this case we say that $\mathcal{M} = \mathcal{G}/d$ is the Markov operator corresponding to $\mathcal{G}$. We have that for $f \in L^2(X)$ the value of $(\mathcal{M}^k f)(x)$ is equal to the expected value of f at the end of a random walk of length k started at x. Furthermore if $H \subseteq X$ is a positive measure set then

$$p_k(H) := (1_H, \mathcal{M}^k 1_H)/\nu(H)$$

is the probability that a random walk of length k started at a random point of H ends in H.

Let $L_0^2(X)$ denote the subspace in $L^2(X)$ consisting of functions with integral equal to zero. If $\mathcal{G}$ is a d-regular graphing then the constant 1 function is an eigenfunction of $\mathcal{G}$ with eigenvalue d and so its orthogonal complement $L_0^2(X)$ is invariant under the action of $\mathcal{G}$. We denote the norm of $\mathcal{G}$ on $L_0^2(X)$ by $\rho(\mathcal{G})$. If $\rho(\mathcal{G}) < d$ then we say that $\mathcal{G}$ *has spectral gap.* Note that the spectral gap is closely related to return probabilities and mixing rates of random walks on $\mathcal{G}$.

The graphing $\mathcal{G}$ is called *ergodic* if there is no measurable connected component $S \subset X$ of $\mathcal{G}$ such that $0 < \nu(S) < 1$. Graphings are typically not connected as abstract graphs so ergodicity is a good substitute for the notion of connectivity. It is easy to see that if a d-regular graphing $\mathcal{G}$ has spectral gap then it has to be ergodic. Furthermore, an ergodic graphing is either a finite connected graph or its probability space has no atoms.

The following statement on $\varrho(\mathcal{G})$ is a modification of well-known facts about finite graphs (see e.g. [7, Theorem 7.1.]) for graphings.

Lemma 2.2 *If $\mathcal{G}$ is an arbitrary d-regular graphing then*

$$\rho(\mathcal{G}) = d \sup_{H,k} \left| \frac{p_k(H) - \nu(H)}{1 - \nu(H)} \right|^{1/k}$$

where H runs through all positive measure sets in X and $k \in \mathbb{N}$.

We prove the following theorem.

Theorem 2.1 *Let $\mathcal{G}$ be a d-regular graphing on an atomless probability space (X, ν). Then $\rho(\mathcal{G}) \geq 2\sqrt{d-1}$.*

Motivated by the previous theorem we will use the following definition.

Definition 2.3 *A d-regular graphing $\mathcal{G}$ is Ramanujan if $\rho(\mathcal{G}) \leq 2\sqrt{d-1}$.*

It follows from Theorem 2.1 that a Ramanujan graphing $\mathcal{G}$ is either a finite Ramanujan graph or $\rho(\mathcal{G}) = 2\sqrt{d-1}$.

We continue with the definition of the Bernoulli graphing of the d-regular tree. Let T_d denote the d-regular infinite tree and let T_d^* denote the version of T_d in which a special vertex o called *root* is distinguished. The set $Y = [0,1]^{T_d^*}$ is a probability space with the product measure. The group of root preserving automorphisms of T_d^* is also acting on Y by the permutation of the coordinates. This action is obviously measure preserving. We denote by Ω_d the space $Y/\mathrm{Aut}(T_d^*)$ with the inherited probability measure ν_d. We connect two elements in Ω_d by an edge if one can be obtained from the other by replacing the root to a neighboring vertex. The graph B_d constructed this way is a d-regular graphing (see [10]) that is called the Bernoulli graphing of T_d. Note that with probability one the connected component of a random element in Ω_d is isomorphic to T_d.

The following theorem seems to have been known for a while (see [11] or [15] Theorem 2.1.) We include a simple proof for completeness.

Theorem 2.2 *For every $d \geq 2$ the Bernoulli graphing B_d is Ramanujan.*

To prove Theorem 2.1 and Theorem 2.2 we will need some preparation. The next lemma is an easy consequence of the spectral theorem.

Lemma 2.4 *Let $\mathcal{F}$ be a bounded, self-adjoint operator on the Hilbert space $\mathcal{H}$ and assume that G spans $\mathcal{H}$. Then*

$$\|\mathcal{F}\| = \sup_{v \in G}\left(\limsup_{i \to \infty}\left|\frac{(v, \mathcal{F}^i v)}{(v,v)}\right|^{1/i}\right).$$

Using this lemma we are ready to prove Lemma 2.2.

Proof of Lemma 2.2 For a positive measure set $H \subseteq X$ let $g_H = \frac{1}{\sqrt{\nu(H)}}1_H - \sqrt{\nu(H)}1_X$. We will use that $g_H \in L_0^2(X)$. Then for every $k \in \mathbb{N}$ we have that

$$\frac{\rho(\mathcal{G})}{d} = \rho(\mathcal{M}) \geq |(g_H, \mathcal{M}^k g_H)/(g_H, g_H)|^{1/k} = |(p_k(H) - \nu(H))/(1 - \nu(H))|^{1/k}.$$

To verify the above calculation note that $\mathcal{M}$ is a self-adjoint operator and thus $(1_X, \mathcal{M}^k 1_H) = (\mathcal{M}^k 1_X, 1_H) = (1_X, 1_H) = \nu(H)$. The other inequality follows from Lemma 2.4 and the fact that functions of the form g_H span the space $L_0^2(X)$. $\square$

The next lemma is well known from probability theory [17].

Lemma 2.5 *Let $S \subset T_d^*$ be a finite subset and let $r_k(S)$ denote the probability that a random walk started at the root ends in S. Then $\lim_{k \to \infty} r_{2k}(S)^{1/2k} = \frac{2\sqrt{d-1}}{d}$, if S contains vertices at even distance from the root. Moreover, $\lim_{k \to \infty} r_{2k+1}(S)^{1/(2k+1)} = \frac{2\sqrt{d-1}}{d}$, if S contains vertices at odd distance.*

Proof of Theorem 2.1: Since T_d covers every d-regular graph it is clear that for every positive measure set $H \subseteq X$ we have that $p_k(H) \geq r_k(o)$. The fact that X is atomless implies that $\nu(H)$ can be arbitrary small in Lemma 2.2 and thus we obtain that $r_k(o)^{1/k} \leq \rho(\mathcal{G})/d$. By Lemma 2.5 this completes the proof. $\qquad\square$

Proof of Theorem 2.2: It follows from Theorem 2.1 that $\rho(B_d) \geq 2\sqrt{d-1}$. Thus by Lemma 2.4 it remains to show that for some spanning set $G \subset L_0^2(\Omega_d)$ the inequality $\limsup_{i \to \infty} |(v, B_d^i v)|^{1/i} \leq 2\sqrt{d-1}$ holds whenever $v \in G$. Let G be the set of all functions with 0 integral and norm 1 on Ω_d that depend only on the labels in a bounded neighborhood of the root.

For $r \in \mathbb{N}$ let S_r denote the neighborhood of the root in T_d^* of radius r. Assume that $f \in G$ is a function which depends only on the labels in S_r. Using the notation of Lemma 2.5 we claim that $c_k = |(\mathcal{M}^k f, f)| \leq r_k(S_{2r})$ where $\mathcal{M} = B_d/d$ is the Markov operator. This fact together with Lemma 2.5 will complete the proof.

To prove the claim observe that c_k is equal to the correlation of the values of f at the two endpoints of a random walk of length k started at a random point $x \in X$. Using the construction of Ω_d we lift the situation to the probability space $[0, 1]^{T_d^*}$. By abusing the notation we assume that f is defined on $[0, 1]^{T_d^*}$ and it is invariant under $\mathrm{Aut}(T_d^*)$. For an element $\omega \in [0, 1]^{T_d^*}$ and $v \in T_d^*$ let $g(v, \omega)$ denote the value of f when the root is replaced to v. (The fact that $g(v, \omega)$ is well-defined relies on the fact that f is invariant under $\mathrm{Aut}(T_d^*)$.) The value of c_k has the following description. We choose a random labeling ω of the vertices T_d^* with $[0, 1]$, start a random walk of length k at the root of T_d^* and on this probability space we take the correlation between $g(o, \omega)$ and $g(v, \omega)$ where v is the endpoint of the walk. Conditioned on the fact that the random walk ends outside S_{2r} it is clear that $g(o, \omega)$ and $g(v, \omega)$ are independent and so the correlation is 0. It follows that the return probability to S_{2r} is an upper bound for c_k. $\qquad\square$

We finish this section with an observation on the spectral properties of B_d. We will use a general fact about operators. Assume $\mathcal{F}$ is a bounded self-adjoint operator on a Hilbert space $\mathcal{H}$, $f \in \mathcal{H}$ and $p \in \mathbb{R}[x]$. Furthermore let P denote the projection valued measure corresponding to $\mathcal{F}$ used in the spectral theorem. Then

$$(f, p(\mathcal{F})f) = \int p \, d\sigma_f \qquad (3)$$

where $\sigma_f(M) = \|P(M)f\|_2^2$ for $M \subset \mathbb{R}$ measurable. Note that if $\|f\|_2 = 1$ then σ_f is a probability measure and we say that σ_f is the spectral measure of f corresponding to $\mathcal{F}$.

Lemma 2.6 *Let $h : \Omega_d \to \{-1, 1\}$ be the function such that $h(\omega) = -1$ if the label on the root is in $[0, 1/2]$ and $h(\omega) = 1$ otherwise. Then the spectral measure of h corresponding to B_d is the Plancherel measure of T_d, i.e. it is concentrated on $[-2\sqrt{d-1}, 2\sqrt{d-1}]$, and its density is the following: $\frac{d}{2\pi} \frac{\sqrt{4(d-1)-t^2}}{d^2-t^2}$ (it is also called the Kesten–McKay measure, see e.g. [17]).*

Proof. Let σ'_h be the spectral measure of h corresponding to B_d/d. We use the proof of Theorem 2.2 for the specific function h. The argument yields that $(h, (B_d/d)^k h)$ is equal to the return probability of a random walk of length k started at the root. On the other hand by (3) $(h, (B_d/d)^k h)$ is equal to the k-moment of σ'_h. This completes the proof. $\qquad\square$

Corollary 2.7 *The set $\{\sigma_f | f \in L_0^2(\Omega) , \|f\|_2 = 1\}$ corresponding to B_d is dense in the set of all probability measures on $[-2\sqrt{d-1}, 2\sqrt{d-1}]$ with respect to the weak topology.*

Proof. For the specific function h defined in Lemma 2.6 we have that the support σ_h is the full interval $[-2\sqrt{d-1}, 2\sqrt{d-1}]$. The existence of one such function (using the spectral theorem) implies the statement. Indeed, for small ε, the uniform measure on the interval $[x, x+\varepsilon]$ can be approximated by $\sigma_{Ph/h(x)}$, where P is the spectral projection to the interval $[x, x+\varepsilon]$. $\qquad\square$

3   Random processes on the tree

In this section we describe how to produce random processes on the tree T_d from graphings. Furthermore, the correlation decay of the process can be bounded by a function of the largest eigenvalue of the graphing.

Definition 3.1 *Let $\mathcal{G}$ be a d-regular graphing on the space (X, ν). We denote by $\operatorname{Hom}_0(T_d, \mathcal{G})$ the set of maps $\phi : V(T_d) \to X$ with the following two properties:*

- *the image of an edge in T_d is an edge in $\mathcal{G}$,*

- *ϕ is injective on the neighborhood of any vertex in T_d.*

The set $\operatorname{Hom}_0(T_d, \mathcal{G})$ is basically the set of covering maps from T_d to $\mathcal{G}$. Note that if P is a path of length k in T_d then the image of P under any $\phi \in \operatorname{Hom}_0(T_d, \mathcal{G})$ is a non-backtracking walk on $\mathcal{G}$.

Lemma 3.2 *There is a unique probability measure κ on $\operatorname{Hom}_0(T_d, \mathcal{G})$ satisfying the following properties:*

- *for every vertex $v \in V(T_d)$ the distribution of the image of v under a κ-random function is ν,*

- *κ is invariant under the action of the automorphism group of T_d.*

Proof. We can uniquely bulid up this probability measure in the following way. Let us start with an arbitrary fixed vertex v of T_d. By the first requirement the image of v has distribution ν. Once the image of v is determined, say $\phi(v) = x$, the remaining vertices of T_d have to be mapped to the connected component of x. The second requirement guarantees that ϕ has to be a randomly

chosen covering of this connected component. Note that the graphing axioms imply that the first requirement holds for every vertex of T_d. □

Using the previous lemma we can create probability distributions of $\mathbb{R}$-valued functions on the vertices of T_d in the following way. A similar construction was used in [1], see Lemmas 35 and 36. Let $f : X \to \mathbb{R}$ be a function in $L^2(X)$. Then the composition $f \circ \phi$ is an $\mathbb{R}$ valued process on the vertices of T_d where ϕ is chosen randomly from $\mathrm{Hom}_0(T_d, \mathcal{G})$. We denote this process by μ_f. An important special case is when $\mathcal{G} = B_d$. In this case processes of the form μ_f are called *factor of i.i.d processes*. These processes can be equivalently defined without graphings in the following way. Let $f : [0,1]^{T_d^*} \to \mathbb{R}$ be a measurable function that is invariant under the natural action of $\mathrm{Aut}(T_d^*)$ on $[0,1]^{T_d^*}$. Then we construct a probability measure on $\mathbb{R}^{T_d}$ such that first we put uniform independent labels on the vertices of T_d and then at each vertex v we evaluate f by putting the root on v. (Note that this is well-defined since f is automorphism invariant.)

Theorem 3.1 *Let $\mathcal{G}$ be a Ramanujan graphing on (X, ν) and let $f : X \to \mathbb{R}$ be a function in $L^2(X)$. Then, in the process μ_f, the absolute value of the correlation of the values on two vertices $v, w \in T_d$ is at most*

$$(k + 1 - 2k/d)\left(\frac{1}{\sqrt{d-1}}\right)^k$$

where k is the distance of v and w.

The rest of the section is the proof of the above theorem. We imitate the proof from the paper [2] in the infinite setting. The main idea is that the correlation decay in μ_f can be expressed in terms of non-backtracking random walks on $\mathcal{G}$.

We define a graphing $\mathcal{G}^{(k)}$ on (X, ν); two vertices are connected in $\mathcal{G}^{(k)}$ if and only if their distance is exactly k in $\mathcal{G}$. More precisely, we need weighted graphings. That is, instead of subsets of $X \times X$, we label the edges with nonnegative integers in a Borel measurable way. These will be the multiplicities of the edges in the graphing. Otherwise the definition is the same as the original one. If $v, w \in T_d$ with distance k, and $f \in L_0^2(X)$ with $||f||_2 = 1$, then by the definition of μ_f the reader can easily check that

$$\mathrm{corr}(\mu_f(v), \mu_f(w)) = \left(f, \frac{\mathcal{G}^{(k)} f}{d(d-1)^{k-1}}\right). \tag{4}$$

The arguments in the proof of Theorem 1.1. of [2] are valid for d-regular graphings as well. Therefore for $k \geq 1$ we have

$$\frac{\mathcal{G}^{(k)}}{d(d-1)^{k-1}} = \frac{1}{\sqrt{d(d-1)^{k-1}}} q_k\left(\frac{\mathcal{G}}{2\sqrt{d-1}}\right), \tag{5}$$

where

$$q_k(x) = \sqrt{\frac{d-1}{d}} U_k(x) - \frac{1}{\sqrt{d(d-1)}} U_{k-2}(x); \quad U_k(\cos \vartheta) = \frac{\sin((k+1)\vartheta)}{\sin \vartheta},$$

7

i.e., U_k is the kth Chebyshev polynomial of the second kind for $k \geq 0$, and $U_{-1} \equiv 0$.

The spectral mapping theorem implies that if $\mathcal{F} : \mathcal{H} \to \mathcal{H}$ is a bounded self-adjoint operator on a Hilbert space $\mathcal{H}$, and p is a polynomial, then $\|p(\mathcal{F})\| \leq \max_{x \in [-\|\mathcal{F}\|, \|\mathcal{F}\|]} |p(x)|$. Since $\mathcal{G}$ is a Ramanujan graphing, its norm on L_0^2 is $2\sqrt{d-1}$. Hence in our case this yields that

$$\varrho\left(\frac{\mathcal{G}^{(k)}}{d(d-1)^{k-1}}\right) \leq \max_{x \in [-1,1]} \frac{|q_k(x)|}{\sqrt{d(d-1)^{k-1}}}.$$

We claim that

$$\max_{x \in [-1,1]} q_k(x) = \sqrt{d/(d-1)}(k+1-2k/d).$$

To see this let $T_k(\cos(\theta)) = \cos(k\theta)$ be the defining equation for the k-th Chebyshev polynomial of the first kind. It is easy to see that

$$\sqrt{d(d-1)} q_k(x) = (d-2)U_k(x) + 2T_k(x).$$

Both $|U_k|$ and $|T_k|$ have their maximal values at 1 and $U_k(1) = k+1$, $T_k(1) = 1$. By substituting 1 into q_k we get the claim.

We obtain that

$$\varrho\left(\frac{\mathcal{G}^{(k)}}{d(d-1)^{k-1}}\right) \leq \left(\frac{1}{\sqrt{d-1}}\right)^k \left(k+1-\frac{2k}{d}\right).$$

This finishes the proof. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad\square$

4 Randomized local algorithms

As it was described in the introduction, a randomized local algorithm produces a random labeling of the vertices of a bounded degree graph using an initial i.i.d labeling and a local rule denoted by f. To give a precise definition we will need the following notation.

Let S be an arbitrary set. An S-labeled graph G is a graph together with a function $h : V(G) \to S$. A rooted graph is a graph in which a special vertex denoted by o and called root is distinguished. Assume that $r, d \in \mathbb{N}$. We denote by $\mathcal{N}(r, d, S)$ the set of isomorphism types of rooted, S-labeled graphs of maximum degree d such that each vertex is of distance at most r from the root. (Isomorphisms are assumed to be root and label preserving.)

A rule of radius r and degree d is a function $f : \mathcal{N}(r, d, S) \to S_2$ where S_2 is some set. Assume that G is a graph of maximal degree at most d and that $h : V(G) \to S$ is some labeling. Then we can use f to produce a new labeling $h_2 : V(G) \to S_2$ such that $h_2(v)$ is equal to the value of f on the S-labeled rooted neighborhood of radius r of v where the root is placed on v. We denote the labeling h_2 by h^f.

8

Definition 4.1 *A randomized local algorithm of radius r and degree d is given by a measurable function $f : \mathcal{N}(r, d, \Omega) \rightarrow L$ (called rule of the algorithm) where Ω is a probability space and L is a measure space. The input of the algorithm is a graph of maximal degree at most d and the output is the random labeling h^f where h is a labeling of G with independent, random elements from Ω.*

Note that local algorithms can also be computed on infinite graphs if they have bounded maximum degree. The next example produces independent sets in graphs [3].

Let $f : \mathcal{N}(1, d, [0, 1]) \rightarrow \{0, 1\}$ be the rule such that the value of f is 1 if and only if the label on the root is the smallest among all labels. It is clear that if $h : V(G) \rightarrow [0, 1]$ is an arbitrary injective function then the support of h^f is an independent set in G. Since random labelings $h : V(G) \rightarrow [0, 1]$ are injective with probability 1 we have that the local algorithm with rule f produces a random independent set with probability 1.

In the rest of this section we focus on the case when G is a d-regular graph with girth more than twice the radius of f. In this case it is enough to define f on Ω-labeled versions of the neighborhood S_r of the root in T_d^* of radius r. In other words we can assume that f is a function of the form $f : \Omega^{S_r} \rightarrow L$ that is invariant under the automorphisms of S_r.

We can also represent f as a function $g : \Omega_d \rightarrow L$ on the vertex set of the Bernoulli graphing B_d. Let $\phi : [0, 1] \rightarrow \Omega$ be an arbitrary measure preserving map and $\psi : \Omega_d \rightarrow \Omega^{S_r}$ be the map defined by deleting the vertices outside S_r and taking the ϕ images of the original labels. Let $g = f \circ \psi$. It is clear that the process μ_g on T_d is the same as the process produced by the local algorithm on T_d with rule f. On the other hand if G is any d-regular graph of girth at least $2(k + r + 1)$ then the distribution of the local algorithm in any ball of radius k is the same as its distribution on T_d in a similar ball. It follows for example that to analyze local properties (such as correlation decay) of local algorithms in the large-girth setting, it is enough to consider the algorithm on the tree T_d. This creates the connection between Bernoulli graphings and local algorithms. As a corollary of Theorem 3.1 and Theorem 2.2 we obtain the following.

Theorem 4.1 *Assume that $r, k \in \mathbb{N}$ and that $f \in L^2(\Omega^{S_r})$ is invariant under $\mathrm{Aut}(S_r)$. Let G be a d-regular graph of girth at least $k + 2r + 2$. If v, w are two vertices of distance k in G then the absolute value of the correlation of the values on v and w in the local algorithm given by f is at most $(k + 1 - 2k/d)(d - 1)^{-k/2}$.*

5   Characterization of correlation sequences

Our goal is to give an algebraic characterization (up to closure with respect to pointwise convergence) for possible correlation sequences in factor of i.i.d processes. We return to the proof of

Theorem 3.1. Let us apply (3) to f and $\mathcal{G}/(2\sqrt{d-1})$ in the calculation. We obtain that the value of (4) for two vertices of distance k is equal to

$$\int_{[-1,1]} \frac{1}{\sqrt{d(d-1)^{k-1}}} q_k \, d\sigma_f$$

where σ_f is the spectral measure of f with respect to $\mathcal{G}/(2\sqrt{d-1})$. The next theorem follows immediately from Corollary 2.7.

Theorem 5.1 *Let X_d denote the set of all sequences with $x_k = \int d^{-1/2}(d-1)^{(1-k)/2} q_k \, d\eta$ where η is a probability measure on $[-1,1]$. Then the closure of possible correlation sequences in factor of i.i.d processes is equal to X_d.*

We finish with an example for a local algorithm on the tree T_d. We start with the intitial i.i.d labeling $\{X_u\}_{u \in T_d}$ where $X_u = 1$ with probability $1/2$ and $X_u = -1$ otherwise. For $r \geq 0$ denote by $S_r(v)$ the neighborhood of radius r around $v \in T_d$. Note that

$$|S_r(v)| = 1 + d\frac{(d-1)^r - 1}{d-2}$$

for $r \geq 1$.

For every $w \in T_d$ we define the random variable $Y_w = \frac{1}{\sqrt{|S_r(w)|}} \sum_{u \in S_r(w)} X_u$. It is clear that $\{Y_w\}_{w \in T_d}$ is the output of a local algorithm. Furthermore we have that

$$\mathrm{corr}(\mu_f(v), \mu_f(w)) = \frac{1}{|S_r(v)|} \mathrm{cov}\left(\sum_{u \in S_r(v)} X_u, \sum_{u \in S_r(w)} X_u \right)$$

$$= \frac{1}{|S_r(v)|} \mathrm{var}\left(\sum_{u \in S_r(v) \cap S_r(w)} X_u \right) = \frac{|S_r(v) \cap S_r(w)|}{|S_r(v)|} = \frac{d(d-1)^{r-k/2} - 2}{d(d-1)^r - 2}$$

if k is even and $k < 2r$. For the last equation we use the fact that $S_r(v) \cap S_r(w)$ is equal to $S_{r-k/2}(z)$ where z is the middle point of of the path connecting v and w. Now, as r goes to infinity, the lower bound for the correlation converges to $(d-1)^{-k/2}$.

For odd k with $k \leq 2r + 1$ we have two points in the middle and so

$$\frac{|S_r(v) \cap S_r(w)|}{|S_r(v)|} = \frac{2[1 + (d-1) + \ldots + (d-1)^{r-(k+1)/2}]}{|S_r(v)|} = 2\frac{(d-1)^{r-(k-1)/2} - 1}{d(d-1)^r - 2}.$$

This converges to $\frac{2}{d}(d-1)^{-\frac{k-1}{2}}$ as $r \to \infty$.

This shows that the correlation decay is close to be optimal in this simple example.

6 Semi-definite functions, spherical representations and Gaussian processes

The goal of this section is to show how correlation sequences of invariant processes on T_d can be viewed from a representation theoretic perspective. Moreover, every such correlation sequence

produces a unique invariant Gaussian process on T_d which is interesting on its own right.

Let $\mathcal{P}_d$ denote the set of all positive semi-definite functions $p : T_d \times T_d \to [-1, 1]$ such that $p(v, v) = 1$ and the value of $p(v, w)$ depends only on the distance of v and w for every pair $v, w \in T_d$. It is clear that the correlation structure of an arbitrary (real valued) invariant process on T_d is an element in $\mathcal{P}_d$. On the other hand any element of $\mathcal{P}_d$ defines a symmetric representation of T_d in some real Hilbert space. To be more precise, there is a function ϕ from T_d to some separable Hilbert space $\mathcal{H}$ such that $p(v, w) = (\phi(v), \phi(w))$ and that $\{\phi(v) | v \in T_d\}$ generates $\mathcal{H}$. It is clear that ϕ is unique up to orthogonal transformations and that there is an orthogonal representation $\psi : \mathrm{Aut}(T_d) \to O(\mathcal{H})$ with the property that $\phi(\alpha(v)) = \psi(\alpha)(\phi(v))$ for every $v \in T_d$ and $\alpha \in \mathrm{Aut}(T_d)$. In particular $\phi(o)$ is fixed under $\mathrm{Aut}(T_d^*)$ where o is any distinguished root in T_d. Such representations of $\mathrm{Aut}(T_d)$ are called *spherical* in the literature. In other words, a representation of $\mathrm{Aut}(T_d)$ is spherical if the subgroup $\mathrm{Aut}(T_d^*)$ has a fixed vector x of length 1 such that the images of x under $\mathrm{Aut}(T_d)$ generate the underlying Hilbert space. It is clear that each spherical representation ψ of $\mathrm{Aut}(T_d)$ gives rise to an element in $\mathcal{P}_d$ by $p(v, w) = (\psi(\alpha_v)(x), \psi(\alpha_w)(x))$ where $\alpha_v(o) = v, \alpha_w(o) = w$. (The spherical property guarantees that p is well-defined.) This construction yields a one to one correspondence between spherical representations and elements in $\mathcal{P}_d$.

To complete the picture, for every $p \in \mathcal{P}_d$ we construct an invariant process with correlation structure p. The most natural choice is an infinite dimensional Gaussian distribution γ_p on $\mathbb{R}^{T_d}$ with correlation structure p. The uniqueness of γ_p guarantees that it is an invariant under $\mathrm{Aut}(T_d)$. If μ is any other invariant process with the same correlation structure p then we can also obtain γ_p by the central limit theorem in the following way. Assume that μ is already normalized in a way that it has zero expectation and variance (at each vertex) equal to 1. Let $[\mu]_n$ denote the distibution of $q_1 + q_2 + \cdots + q_n$ where each q_i is an independent element of $\mathbb{R}^{T_d}$ chosen with distribution μ. It is clear that the weak limit of $n^{-1/2}[\mu]_n$ is a Gaussian process with correlation structure p and thus it is equal to γ_p. If in particular μ is a factor of i.i.d process then so is $[\mu]_n$ for every n. It follows that in this case γ_p is a weak limit of factor of i.i.d processes. Therefore we obtain the following.

Corollary 6.1 *A Gaussian process is a limit of factor of i.i.d. processes if and only if its correlation decay is as in Theorem 5.1.*

References

[1] M. Abért, Y. Glasner, B. Virág. The measurable Kesten theorem. *http://arxiv.org/abs/1111.2080.*

[2] N. Alon, I. Benjamini, E. Lubetzky, S. Sodin: Non-backtracking random walks mix faster, *J. CCM Vol. 09, no. 4, 2007.*

[3] N. Alon, J. Spencer: *The Probabilistic Method*, J. Wiley and. Sons, New York, 2nd edition, 2000.

[4] E. Csóka: Random local algorithms, *http://arxiv.org/abs/1202.1565*

[5] E. Csóka: Maximum flow is approximable by deterministic constant-time algorithm is sparse networks, *http://arxiv.org/abs/1005.0513*

[6] G. Elek, G. Lippner: Borel oracles. An analytic approach to constant time algorithms, *Proc. AMS,* **138***, no. 8, Aug 2010.*

[7] J. Friedman, J-P. Tillich: Generalized Alon–Boppana theorems and error-correcting codes. *SIAM J. Discrete Math.* **19** *no.3. 700-718, 2005.*

[8] D. Gaboriau, R. Lyons: A measurable group theoretic solution to von Neumann's problem, *Invent. Math.,* **177** 533–540, 2009.

[9] D. Gamarnik, M. Sudan: Limits of local algorithms over sparse random graphs, arXiv: 1304.1831 [math.PR].

[10] H. Hatami, L. Lovász, B. Szegedy: Limits of local-global convergent graph sequences, *http://arxiv.org/abs/1205.4356*

[11] A. S. Kechris, T. Tsankov: Amenable actions and almost invariant sets, *Proc. Amer. Math. Soc.,* **136** *687-697*, 2008 (electronic).

[12] H. Kesten: Symmetric random walks on groups, *Trans. Amer. Math. Soc.,* **92** *336–354, 1959.*

[13] G. Kun: Expanders have a spanning Lipschitz subgraph with large girth, *http://arxiv.org/abs/1303.4982.*

[14] A. Lubotzky, R. Phillips, P. Sarnak: Ramanujan graphs, *Combinatorica,* **8** *(3): 261–277, 1988.*

[15] R. Lyons, F. Nazarov: Perfect matchings as IID factors on non-amenable groups, *Europ. J. Combin.,* **32** *1115–1125, 2011.*

[16] M. Naor, L. Stockmeyer: What can be computed locally?, *SIAM Journal of computing* **24***, 1995.*

[17] W. Woess: *Random walks on infinite graphs and groups*, Cambridge University Press, 2000.

ÁGNES BACKHAUSZ. Eötvös Loránd University, Budapest, Hungary. `agnes@cs.elte.hu`

BALÁZS SZEGEDY. University of Toronto, Canada. `szegedyb@gmail.com`

BÁLINT VIRÁG. University of Toronto, Canada. `balint@math.toronto.edu`