Reversibility of the non-backtracking random walk

Jonathan Hermon

University of Cambridge, Cambridge, UK. E-mail: jh2129@statslab.cam.ac.uk

Abstract. Let G be a connected graph of uniformly bounded degree. A k non-backtracking random walk (k-NBRW) $(X_n)_{n=0}^{\infty}$ on G evolves according to the following rule: Given $(X_n)_{n=0}^{s}$, at time $s+1$ the walk picks at random some edge which is incident to X_s that was not crossed in the last k steps and moves to its other end-point. If no such edge exists then it makes a simple random walk step. Assume that for some $R > 0$ every ball of radius R in G contains a cycle. We show that under some "nice" random time change the 1-NBRW becomes reversible. This is used to prove that it is recurrent iff the simple random walk is recurrent. A similar result is proved for every k under stronger assumptions in general, and with no assumptions for Cayley graphs of finitely generated Abelian groups.

Résumé. Soit G un graphe connexe dont les degrés sont uniformément bornés. La marche aléatoire k non-rebroussante (k-NBRW) sur G évolue de la manière suivante: sachant $(X_n)_{n=0}^{s}$, l'état X_{s+1} est obtenu en choisissant uniformément au hasard une arête incidente à X_s qui n'a pas été empruntée durant les k mouvements précédents, et en l'empruntant. Si toutes les arêtes incidentes à X_s ont été empruntées durant les k mouvements précédents, on choisit une arête incidente à X_s uniformément au hasard. Supposons que toute boule de rayon R dans G contient un cycle. Alors nous montrons que, sous un changement de temps judicieux, la 1-NBRW devient réversible. Nous en déduisons que cette marche est récurrente si et seulement si la marche aléatoire simple l'est aussi. Un résultat similaire est établi pour tout k, sous des hypothèses supplémentaires en général, mais sans aucune hypothèse dans le cas des graphes de Cayley de groupes abéliens finiment engendrés.

1. Introduction

In this work we study a generalization of non-backtracking random walk (**NBRW**) on a graph G, in which the walk is required (when possible) to avoid either the last k edges it crossed (viewed as undirected edges), or the last k vertices it visited. If this is not possible, we say the walk got stuck, in which case it moves to a random neighbor. We call such a walk an edge (respectively, a vertex[1]) k-*NBRW* (for more precise definitions see Section 1.1).

The special case of $k = 1$, which is simply a NBRW, received much attention (e.g. [2–5,11,14,15,17,18,20]). A recurring theme in the existing literature is that when G is regular the NBRW and the simple random walk (**SRW**) on G (see Section 1.1 for the definition of SRW) are intimately related to one another. In particular, when G is regular one may deduce some properties of one walk by establishing related properties for the other, or estimate certain parameters for one walk via related parameters of the other (e.g. [2,17,20]). Conversely, when G is non-regular the two walks can exhibit quite different behaviors (e.g. [4,18]).

---

[1]When a certain statement is true for both the edge and the vertex NBRWs or if the type of walk is clear from context, we often omit the prefix edge/vertex.

Indeed, when G is d-regular for $d \geq 3$, a well-known argument involving the universal cover of G shows that the SRW $(S_n)_{n=0}^\infty$ and the NBRW $(S_n^{BN})_{n=0}^\infty$ on G can be coupled so that $S_n^{BN} = S_{\tau_n}$ for all n, where $(\tau_{n+1} - \tau_n)_{n=1}^\infty$ are i.i.d. random variables satisfying that $P(\tau_2 - \tau_1 > s) \leq e^{-cs}$ for all $s \geq 0$ for some $c > 0$ (see Proposition 1.1). However, there does not exist such a coupling when G is non-regular or when the NBRW is replaced by a k-NBRW, provided that $k > \text{girth}(G)$ (where the *girth* of a graph G, denoted by $\text{girth}(G)$, is defined as the length of the shortest cycle in G).

One of the most fundamental questions that one can ask about any random walk is whether it is recurrent or transient. A fairly simple consequence of the aforementioned coupling is that when G is regular SRW is recurrent iff the NBRW is recurrent.

Proposition 1.1. *Let $G = (V, E)$ be a d-regular connected graph for some $d \geq 3$. Then the SRW on G is recurrent iff the NBRW on G is recurrent.*

Proposition 1.1 is proved in Section A.1. It is natural to ask whether the same assertion holds even without regularity or when the NBRW is replaced with a k-NBRW. In this work we show that under some mild conditions the answer to the last question is positive. To the best of the author's knowledge, this was unknown even for $k = 1$. Moreover, we show (Theorem 2) that for Cayley graphs of finitely generated infinite Abelian groups which are not isomorphic to $\mathbb{Z}$ we have for all k that the SRW is recurrent iff the k-NBRW is recurrent. Finally, we show (Theorem 3) that for any connected graph of bounded degree the variant of the NBRW which backtracks at each step with probability $p \in (0, 1)$ is recurrent iff the SRW is recurrent.

A major difficulty in studying the k-NBRW is that even for $k = 1$ it is non-reversible (see Section 2 for the definition of reversibility). It is much harder to apply techniques from potential theory to non-reversible Markov chains. One reason is that in order to apply techniques from potential theory one first has to know a stationary measure for the chain. For non-reversible Markov chains it is sometimes hard to find a stationary measure, and if the chain is transient, it is possible that no stationary measure exists (although for the k-NBRW a stationary measure does exist, see Section 2). In particular, for $k > 1$ it seems quite difficult to find a stationary measure for the vertex k-NBRW. In fact, this is also true for the edge k-NBRW if it may get stuck.

Another reason is the lack of rich variety of comparison techniques available in the reversible setup (see Section 4). However, one available comparison technique which we shall exploit is that the additive symmetrization of a recurrent Markov chain is always recurrent (see Theorem 4.12). Unfortunately, there exist only few techniques for establishing the converse implication, or more generally for arguing that the geometry of a non-reversible chain is in some sense equivalent to that of a reversible one.

To overcome this difficulty, we show that if G has a positive density of short simple cycles then viewed at "nice" random times, the edge NBRW becomes a reversible chain which can be compared with SRW on G. So is the case for the k-NBRW under slightly more complicated conditions.

1.1. *Basic definitions – SRW, NBRW and recurrence/transience*

Let $G = (V, E)$ be a locally finite connected (simple)[2] graph. A *simple random walk* (**SRW**) on G is a reversible Markov chain with state space V and transition probabilities: $P(x, y) = \mathbf{1}\{x \sim y\}/\deg(x)$, where $x \sim y$ indicates that $\{x, y\} \in E$, and $\deg(x)$ is the degree of x (which is the number of edges which are incident to x). A *non-backtracking random walk* (**NBRW**) $(X_t)_{t=0}^\infty$ on G started from vertex u makes its first step according to the same rule as SRW started from u, and then evolves according to the following rule: the conditional distribution of X_{t+2} given $(X_i)_{i=0}^{t+1}$ is chosen from the uniform distribution on $\{z \in V : z \sim X_{t+1}, z \neq X_t\}$, unless this set is empty (i.e. $\deg(X_{t+1}) = 1$)), in which case $X_{t+2} = X_t$ (i.e., the walk is forced to backtrack). While this is **not** a Markov chain, it can be transformed into a Markov chain on $\overrightarrow{E}$, the set of directed edges of G, whose transition probabilities are given by

$$B\big((x, y), (z, w)\big) := \frac{\mathbf{1}\{y = z, w \neq x\}}{\deg(y) - 1} + \mathbf{1}\{y = z, w = x, \deg(y) = 1\}. \tag{1.1}$$

---

[2]Although all of our results can be extended to the case that G is a multigraph, for the sake of simplicity of notation and clarity of presentation we shall assume that G contains no multiple edges nor loops.

Let $k \in \mathbb{N}$. Similarly, an *edge k-NBRW* (respectively, a *vertex k-NBRW*) $(X_t)_{t=0}^{\infty}$ evolves according to the following rule: the conditional distribution of X_{t+1} given $(X_i)_{i=0}^t$ is chosen from the uniform distribution on

$$\left\{ z \in V : z \sim X_t, \{z, X_t\} \notin \left\{ \{X_{i-1}, X_i\} : t - k < i \le t \right\} \right\}$$

$$\left(\text{respectively,} \left\{ z \in V : z \sim X_t, z \notin \{X_i : t - k \le i \le t\} \right\} \right),$$

unless this set is empty, in which case we say the walk is *stuck*. If the walk is stuck then X_{t+1} is chosen from the uniform distribution over the neighbors of X_t (this is done in order to ensure the k-NBRW does not have absorbing states). These walks can be transformed into Markov chains on a subset of the set of directed paths of length k in an analogous manner to the case $k = 1$. We denote the collection of all paths of length k in the graph $G = (V, E)$ (where G will be clear from context) by

$$\mathcal{P}_k := \left\{ (x_i)_{i=0}^k \in V^{k+1} : \{x_i, x_{i-1}\} \in E \text{ for all } i \in [k] \right\}.$$

Let

$$\mathcal{P}_k^{\mathrm{e}} := \{\gamma \in \mathcal{P}_k : \gamma \text{ crosses } k \text{ distinct edges}\} \quad \text{and}$$

$$\mathcal{P}_k^{\mathrm{v}} := \{\gamma \in \mathcal{P}_k : \gamma \text{ contains } k + 1 \text{ distinct vertices}\}. \tag{1.2}$$

We may take the state space of the edge (respectively, vertex) k-NBRW to be the collection of all length k paths which are accessible from $\mathcal{P}_k^{\mathrm{e}}$ (respectively, $\mathcal{P}_k^{\mathrm{v}}$). If the walk cannot get stuck then the state spaces are simply $\mathcal{P}_k^{\mathrm{e}}$ and $\mathcal{P}_k^{\mathrm{v}}$, respectively (see Section 3.2 for details).

Remark 1.2. Note that (when G is simple) the NBRW is both a vertex and an edge 1-NBRW. However for $k > 1$ the restrictions imposed on the vertex k-NBRW are stronger than those imposed on the edge k-NBRW. Moreover, the former is less symmetric (the latter satisfies a symmetry relation (7.4) similar to the detailed balance equation, which the former does not) and thus harder to analyze.

Remark 1.3. In what comes we shall often assume that the k-NBRW cannot get stuck. Note that the vertex (respectively, edge) k-NBRW cannot get stuck if the minimal degree of the graph is at least $k + 1$ (respectively, $\lceil 2k/3 \rceil + 1$).

Another condition which ensures that the edge k-NBRW cannot get stuck is that all vertices are either of even degree or of degree at least $\lceil 2k/3 \rceil + 1$. To be precise (under the last condition) the only vertex at which the edge k-NBRW might potentially get stuck is the starting point, but this is not the case in the Markov chain representation of the edge k-NBRW, provided that the starting state is a path that does not get stuck. Consequently, the collection of graphs for which the edge k-NBRW cannot get stuck is large.

We say that a vertex v is *recurrent* for a certain type of random walk on a graph G (e.g. SRW and edge/vertex k-NBRW) if started from v the walk returns to v infinitely often a.s.. Otherwise, vertex v is said to be *transient* for that walk type. We say that G is recurrent (respectively, transient) for a certain type of random walk if all of its vertices are recurrent (respectively, transient) for that walk.

Remark 1.4. When the k-NBRW is viewed as a Markov chain, its state space is not the vertex set, but a certain subset of the collection of all paths of length k. In all of the cases considered in our main results, Theorems 1–3, the above notion of recurrence for the k-NBRW is in fact equivalent to the usual notion of recurrence for a Markov chain. Thus there is no ambiguity in the definition of recurrence/transience.

1.2. *Our results*

Before stating our results we first list some conditions for later reference.

(1) There exists some $L > 0$ such that every *2-path* is of length at most L, where a 2-path is a path consisting of degree 2 vertices.

(2) There exists some $R > 0$ such that every ball of radius R contains a cycle.

(2_k^e) There exists some $R > 0$ such that for every $\gamma, \gamma' \in \mathcal{P}_k^e$ which contain at least one common vertex, the edge k-NBRW can reach γ' from γ in at most R steps.

(2_k^v) There exists some $R > 0$ such that for every $\gamma, \gamma' \in \mathcal{P}_k^v$ which contain at least one common vertex, the vertex k-NBRW can reach γ' from γ in at most R steps.

(3_k^e) The edge k-NBRW on G cannot get stuck.

(3_k^v) The vertex k-NBRW on G cannot get stuck.

Observe that conditions (2_k^e) and (2_k^v) both imply condition (1). Condition (1) will not be assumed in any of our main results.

Let $v_0, v_1, \ldots, v_\ell = v_0 \in V$ be such that $\{v_{i-1}, v_i\} \in E$ for all $i \in [\ell] := \{1, 2, \ldots, \ell\}$. We say that $(v_0, v_1, \ldots, v_\ell = v_0)$ is an *edge* (respectively, *vertex*) *simple cycle* (of length ℓ) if $\{v_{i-1}, v_i\} \neq \{v_{j-1}, v_j\}$ for all $i \neq j \in [\ell]$ (respectively, if $|\{v_i : i \in [\ell]\}| = \ell$).

(4_k^e) There exists some $R > 0$ such that every ball of radius R contains an edge simple cycle of length at least k.

(4_k^v) There exists some $R > 0$ such that every ball of radius R contains a vertex simple cycle of length at least $k + 1$.

We believe that conditions (2_k^e) and (2_k^v) are in fact equivalent to conditions (4_k^e) and (4_k^v), respectively. Conditions (4_k^e) and (4_k^v) will not be used in the paper.

The next theorem concerns with the case that G is of (uniformly) *bounded degree* (i.e., $\sup_{v \in V} \deg(v) < \infty$).

Theorem 1. *Let G be a connected graph of bounded degree.*

  (i) *If the SRW on G is transient then the NBRW on G is also transient.*

 (ii) *If condition* (2) *holds then the SRW on G is transient iff the NBRW on G is transient.*

(iii) *If conditions* (2_k^e) *and* (3_k^e) *hold, then the SRW on G is transient iff the edge k-NBRW is transient.*

(iv) *If G is vertex-transitive, conditions* (2_k^v) *and* (3_k^v) *hold and the SRW on G is transient, then the vertex k-NBRW is also transient.*

Let H be a countable group. Let S be a finite symmetric (i.e., $S = S^{-1} := \{s^{-1} : s \in S\}$) set of generators of H (that is, $H = \{s_1 \cdots s_n : n \geq 0, s_1, \ldots, s_n \in S\}$). The (right) **Cayley graph** of H w.r.t. S is $H(S) := (H, E(S))$, where $E(S) := \{\{h, hs\} : h \in H, s \in S\}$. As the following theorem asserts, when G is a Cayley graph of an Abelian group, no additional assumptions are necessary for the assertion of Theorem 1 to hold.

Theorem 2. *Let H be a finitely generated infinite Abelian group which is not isomorphic to $\mathbb{Z}$. Let $G = H(S)$ be the Cayley graph of H w.r.t. some finite symmetric set of generators S. Let $k \in \mathbb{N}$. Then the following are equivalent*:

  (i) *The SRW on G is recurrent.*

 (ii) *The edge k-NBRW on G is recurrent.*

(iii) *The vertex k-NBRW on G is recurrent.*

Remark 1.5. While the assertion of Theorem 2 clearly fails when $H = \mathbb{Z}$ and $S = \{\pm 1\}$, it is not hard to modify the proof of Theorem 2 to show that if $H = \mathbb{Z}$ and $|S| > 2$ then for all $k \in \mathbb{N}$ the edge k-NBRW is recurrent.

Remark 1.6. It is classical that whether SRW on $H(S)$ is recurrent or not does not depend on S (that is, if S_1, S_2 are two finite symmetric sets of generators of H then the SRW on $H(S_1)$ is recurrent iff the SRW on $H(S_2)$ is recurrent). It follows from Theorem 2 that if $H \neq \mathbb{Z}$ is Abelian, the same is true for the edge and vertex k-NBRWs for all $k \in \mathbb{N}$.

We now define a variant of the NBRW for which condition (2) is not required for the assertion of Part (ii) of Theorem 1 to hold.

Definition 1.7. Let $G = (V, E)$ be a locally-finite connected graph. Fix some $p \in (0, 1)$. Consider a walk on the set of directed edges $\overrightarrow{E}$ whose transition kernel is given by

$$B_p\big((x, y), (z, w)\big) := (1 - p)B\big((x, y), (z, w)\big) + p \cdot \mathbf{1}\{z = y, w = x\},$$

where B is the transition kernel of the NBRW on G (i.e., at each step with probability p the walk backtracks and otherwise it evolves like a NBRW). We call this walk a *p-BRW*.

While the following theorem is of self-interest, its proof will be used to highlight the key ideas behind the proofs of Theorems 1 and 2 in a simpler setup.

Theorem 3. *Let $G = (V, E)$ be a connected graph of bounded degree. Then for every $p \in (0, 1)$ the p-BRW on G is transient iff the SRW on G is transient.*

1.3. *Open problems and remarks*

Recall that two graphs G_1, G_2 are roughly-isometric if one can be embedded in the other in an "almost surjective" manner while distorting distances by at most a constant additive and multiplicative factor (see Definition 4.6 for a precise formulation). Recall further that the relation of "being roughly-isometric" is an equivalence relation which preserves recurrence/transience (see Fact 4.9). Moreover, it is not hard to verify that condition (2) is invariant under rough-isometries. Hence the following corollary is an immediate consequence of Theorems 1 and 3.

Corollary 1.8. *Let G and H be two connected roughly-isometric graphs of bounded degree. Then the following hold*:

  (i) *If the SRW on G is transient then the NBRW on G and on H are both transient.*
 (ii) *If condition (2) holds for G then the NBRW on G is transient iff the NBRW on H is transient.*
(iii) *For every $p \in (0, 1)$ the p-BRW on G is transient iff the p-BRW on H is transient.*
(iv) *If G and H both satisfy conditions (2_k^e) and (3_k^e), then the edge k-NBRW on G is transient iff the edge k-NBRW on H is transient.*

We leave the following as open problems.

Question 1.9. Let G be a connected graph of bounded degree which satisfies conditions (2_k^v) and (3_k^v). Is it the case that the SRW on G is transient iff the vertex k-NBRW is transient?

Question 1.10. Does the assertion of Theorem 2 remain valid when G is only assumed to be vertex-transitive?

Question 1.11. Let G be a connected graph of bounded degree satisfying condition (1). Is it the case that the SRW on G is transient iff the NBRW on G is transient?

The following example demonstrates that if condition (1) fails then it is possible that the SRW is recurrent while the NBRW is transient.

Example 1.12. Let $H = (V, E)$ be some transient graph of bounded degree. Fix some $o \in V$. Let Π_n be the collection of all edges with one end-point of distance $n - 1$ from o and the other of distance n. Let G be the graph obtained from H by replacing for all $n \in \mathbb{N}$ each edge in Π_n by a path of length $|\Pi_n|$.

Observe that when the NBRW on G enters one of its 2-paths it must reach the other end of that 2-path. Thus if we observe the NBRW on G only when it visits sites belonging to V we obtain a realization of the NBRW on H. Since the SRW on H is transient, by Theorem 1 so is the NBRW on H and hence the NBRW on G must be transient as well.

To see that the SRW on G is recurrent, first apply a standard network reduction, replacing each 2-path added to G by a single edge of edge-weight which equals the inverse of the length of that path (obtaining a network on H in which each edge in Π_n has edge-weight $1/|\Pi_n|$). Then apply the Nash-Williams criterion (e.g. [19, (2.13)]) to the cut sets $\Pi_1, \Pi_2, \ldots$.

To end the introduction, the author would like to thank Itai Benjamini for suggesting most of the problems considered in this paper (private communication).

1.4. *Organization of the paper*

In Section 2 we present some background on Markov chains. In Section 3 we introduce additional notation and describe a representation of the k-NBRW as a Markov chain. In Section 4 we review a variety of comparison techniques. In Section 5 we apply some of these techniques in order to compare SRW with the additive symmetrization of the k-NBRW. In Section 6 we give an overview of the main ideas behind the proofs of our main results and prove Theorem 3. In Section 7 we find a stationary measure for the edge k-NBRW (under condition (3_k^e)). In Section 8–9 we prove Theorems 1–2, respectively.

2. Some definitions related to Markov chains

We now recall some basic definitions related to Markov chains. For further background on Markov chains see e.g. [1, 16,19]. We call $(V, E, (c_e)_{e \in E})$ a *network* if (V, E) is a graph (we allow it to contain loops, i.e. $E \subseteq \{\{u, v\} : u, v \in V\}$) and $(c_e)_{e \in E}$ are symmetric edge weights (i.e., $c_{u,v} = c_{v,u} > 0$ for every $\{u, v\} \in E$ and $c_{u,v} = 0$ if $\{u, v\} \notin E$) such that $c_v := \sum_w c_{v,w} < \infty$ for all $v \in V$. We say the network is connected if the graph (V, E) is connected. The (weighted) nearest neighbor random walk corresponding to $(V, E, (c_e)_{e \in E})$ repeatedly does the following: when the current state is $v \in V$, the walk will move to vertex u with probability $c_{u,v}/c_v$. The choice $c_e = \mathbf{1}\{e \in E\}$ corresponds to SRW on (V, E). We denote this network by $(V, E, (1)_{e \in E})$.

Consider an ergodic Markov chain $(X_k)_{k=0}^{\infty}$ on a countable state space V with a stationary measure π and transition probabilities given by P. We say that P is *reversible* w.r.t. π if $\pi(x)P(x, y) = \pi(y)P(y, x)$ for all $x, y \in V$. Observe that a weighted nearest-neighbor random walk is reversible w.r.t. the measure π given by $\pi(v) := c_v$ and hence π is stationary for the walk. Conversely, every reversible chain can be presented as a network with weights $c_{x,y} = \pi(x)P(x, y)$. The *time-reversal* of P (w.r.t. π), denoted by P^*, is given by $\pi(x)P^*(x, y) := \pi(y)P(y, x)$ and its *additive symmetrization* (w.r.t. π) is a reversible Markov chain (w.r.t. π) whose transition probabilities are given by $S := \frac{1}{2}(P + P^*)$.

Note that when the chain is recurrent, π is unique up to a constant factor, and so P_* and S are uniquely defined. This may fail when the chain is transient! A transient Markov chain may fail to have a stationary measure and may have two different stationary measures which are not constant multiples of one another. Consider a random walk on $\mathbb{Z}$ with $P(i, i+1) = p = 1 - P(i, i-1)$ for all $i \in \mathbb{Z}$, with $p \in (1/2, 1)$. Note that P is reversible w.r.t. $\pi(i) = (p/(1-p))^i$ and is stationary also w.r.t. the counting measure μ on $\mathbb{Z}$. The additive symmetrization w.r.t. π is again P while w.r.t. μ is SRW on $\mathbb{Z}$.

Harris [13] gave a general mild sufficient condition for the existence of a stationary measure for an irreducible transient Markov chain on a countable state space. In particular, it follows from his result that (in this setup) if the transition matrix P satisfies that $\{y : P(y, x) > 0\}$ is finite for all x, then P has a stationary measure. Clearly this implies that the k-NBRW has a stationary measure.

The *hitting time* of a set $A \subset V$ is $T_A := \inf\{t \geq 0 : X_t \in A\}$. Similarly, let $T_A^+ := \inf\{t \geq 1 : X_t \in A\}$. When $A = \{x\}$ is a singleton, we write T_x and T_x^+ instead of $T_{\{x\}}$ and $T_{\{x\}}^+$. For a set $A \subset V$ such that $P_v[T_A^+ < \infty] = 1$ for all $v \in V$ (under irreducibility we may replace 'for all' by 'for some') the *induced chain on* A is a Markov chain with state space A and transition probabilities given by $Q_A(a, b) := P_a[X_{T_A^+} = b]$, where P_a denotes the law of the entire chain, started from vertex u. It is not hard to verify that if the original chain is reversible w.r.t. π, then the induced chain on A is reversible w.r.t. to the restriction of π to A (e.g. [16, p. 186]). In fact, if the original chain is recurrent then this holds even without reversibility. Indeed in this case the stationary measure π is unique up to a multiplication by a constant factor and is given by $\pi(y) = \mathbb{E}_x[|\{t \in [0, X_{T_x^+}) : X_t = y\}|]$ (where x is arbitrary). Taking $x, y \in A$ we see that the r.h.s. is the same as the corresponding expectation for the induced chain on A, and so the stationary measure of the induced chain on A is simply the restriction of π to A as claimed.

We now assume that $\sum_{v \in V} \pi(v) = \infty$. The ℓ_2 and ℓ_∞ norms of $f \in \mathbb{R}^V$ are given by $\|f\|_2 := \sqrt{\langle f, f \rangle_\pi}$ and $\|f\|_\infty := \sup_{v \in V} |f(v)|$, where $\langle f, g \rangle_\pi := \sum_{v \in V} \pi(v) f(v) g(v)$. The space of ℓ_2 functions is given by $\mathcal{H} := L_2(V, \pi) = \{f : \mathbb{R}^V : \|f\|_2 < \infty\}$. Then P (and similarly, also S and P^*) defines a linear operator on $\mathcal{H}$ via $Pf(x) := \sum_{y \in V} P(x, y) f(y) = \mathbb{E}_x[f(X_1)]$. Note that P^* is the dual of P and that S is self-adjoint. The *Dirich-*

let form $\mathcal{E}_P(\cdot, \cdot) : \mathcal{H}^2 \to \mathbb{R}$ corresponding to P is $\mathcal{E}_P(f, g) := \langle (I - P)f, g \rangle_\pi$. Note that

$$\mathcal{E}_P(f, f) = \mathcal{E}_{P^*}(f, f) = \mathcal{E}_S(f, f) = \frac{1}{2} \sum_{x,y} \pi(x) S(x, y) \big(f(x) - f(y) \big)^2. \tag{2.1}$$

3. Notation and a construction of the k-NBRW as a Markov chain

3.1. *Notation*

Fix some graph $G = (V, E)$. Here we allow G to contain loops (however we do not allow it to have multiple edges between a pair of vertices). Recall that we write $u \sim v$ iff $\{u, v\} \in E$. We denote a directed edge from x to y by (x, y) and an undirected edge between x and y by $\{x, y\}$. We denote the set of directed edges by $\overrightarrow{E} := \{(x, y) : \{x, y\} \in E\}$. For $\overrightarrow{e} = (x, y) \in \overrightarrow{E}$ we denote its *reversal*, tail and head (respectively) by $\overrightarrow{e}^{\mathrm{r}} := (y, x)$, $\overrightarrow{e}_- := x$ and $\overrightarrow{e}_+ := y$. Let $k \in \mathbb{N}$. We denote the set $\{1, 2, \dots, k\}$ by $[k]$. Recall that we denote the collection of all paths of length k in a graph G by

$$\mathcal{P}_k := \big\{ (x_i)_{i=0}^k \in V^{k+1} : \{x_i, x_{i-1}\} \in E \text{ for all } i \in [k] \big\}.$$

Let $\gamma = (\gamma_0, \dots, \gamma_k) \in \mathcal{P}_k$. We denote the ith coordinate of γ by γ_i. The *reversal* of γ is given by $\gamma^{\mathrm{r}} := (\gamma_k, \dots, \gamma_0)$ (that is, $\gamma_i^{\mathrm{r}} := \gamma_{k-i}$ for all i). Its length is denoted by $|\gamma|$ (where $|\gamma| + 1$ is the number of coordinates in γ). For $e \in E$ (respectively, $v \in V$) we write $e \in \gamma$ (respectively, $v \in \gamma$) if $e \in \gamma^{\mathrm{e}} := \{\{\gamma_{i-1}, \gamma_i\} : i \in [k]\}$ (respectively, $v \in \gamma^{\mathrm{v}} := \{\gamma_i : 0 \le i \le k\}$). For $\gamma = (\gamma_0, \dots, \gamma_k) \in \mathcal{P}_k$ and $v \sim \gamma_k$ we write $v \overset{\mathrm{e}}{\sim} \gamma$ (respectively, $v \overset{\mathrm{v}}{\sim} \gamma$) if $\{\gamma_k, v\} \notin \gamma$ (respectively, $v \notin \gamma$). We also write $\gamma(v) := (\gamma_1, \dots, \gamma_k, v)$ (note that $\gamma(v) \in \mathcal{P}_k$ and $\gamma(v)_i := \gamma_{i+1}$ for $0 \le i < k$). As in (1.2) let

$$\mathcal{P}_k^{\mathrm{e}} := \big\{ \gamma \in \mathcal{P}_k : |\gamma^{\mathrm{e}}| = k \big\} \quad \text{and} \quad \mathcal{P}_k^{\mathrm{v}} := \big\{ \gamma \in \mathcal{P}_k : |\gamma^{\mathrm{v}}| = k + 1 \big\}.$$

Observe that $\gamma \in \mathcal{P}_k^{\mathrm{e}}$ (respectively, $\gamma \in \mathcal{P}_k^{\mathrm{v}}$) iff $\gamma^{\mathrm{r}} \in \mathcal{P}_k^{\mathrm{e}}$ (respectively, $\gamma^{\mathrm{r}} \in \mathcal{P}_k^{\mathrm{v}}$). For $\gamma \in \mathcal{P}_k$ let $A(\gamma) := \{\gamma(v) : v \sim \gamma_k\}$,

$$N^{\mathrm{e}}(\gamma) := \big\{ \gamma(v) : v \overset{\mathrm{e}}{\sim} \gamma \big\} \quad \text{and} \quad N^{\mathrm{v}}(\gamma) := \big\{ \gamma(v) : v \overset{\mathrm{v}}{\sim} \gamma \big\}.$$

We denote the transition kernel of the edge (respectively, vertex) k-NBRW by $P_{k,\mathrm{e}}$ (respectively, $P_{k,\mathrm{v}}$). For every set A we denote the counting measure on A (i.e., the measure which assigns to every $a \in A$ mass 1) by π_A.

3.2. *A representation of the k-NBRW as a Markov chain*

Let $k \ge 1$. In this section we construct the edge and vertex k-NBRWs as Markov chains. We first construct the edge and vertex k-NBRWs as Markov chains on $\mathcal{P}_k$ and then later we restrict the state space to include only the "relevant" paths. Given that the current location of the walk is γ the edge (respectively, vertex) k-NBRW evolves according to the following rule: If $N^{\mathrm{e}}(\gamma) \ne \varnothing$ (respectively, $N^{\mathrm{v}}(\gamma) \ne \varnothing$), the walk moves to some $\gamma' \in N^{\mathrm{e}}(\gamma)$ (respectively, $\gamma' \in N^{\mathrm{v}}(\gamma)$) chosen from the uniform distribution. Otherwise, it moves to some $\gamma' \in A(\gamma)$ chosen from the uniform distribution.

We take the state space of the edge (respectively, vertex) k-NBRW to be

$$\Omega_k^{\mathrm{e}} := \big\{ \gamma' \in \mathcal{P}_k : \exists \gamma \in \mathcal{P}_k^{\mathrm{e}} \text{ and } n \ge 0 \text{ such that } P_{k,\mathrm{e}}^n(\gamma, \gamma') > 0 \big\}$$

(respectively, $\Omega_k^{\mathrm{v}} := \{ \gamma' \in \mathcal{P}_k : \exists \gamma \in \mathcal{P}_k^{\mathrm{v}} \text{ and } n \ge 0 \text{ such that } P_{k,\mathrm{v}}^n(\gamma, \gamma') > 0 \}$), the collection of all states which are accessible starting from some state in $\mathcal{P}_k^{\mathrm{e}}$ (respectively, $\mathcal{P}_k^{\mathrm{v}}$).

4. Comparison techniques

In this section we present various techniques which allow one to deduce that a certain Markov chain is recurrent (or transient) if some other chain is recurrent (or transient).

The *capacity* of a finite set A (w.r.t. a certain Markov chain and some stationary measure π) is defined as

$$\mathrm{Cap}(A) := \sum_{a \in A} \pi(a) \mathrm{P}_a \big[T_A^+ = \infty \big].$$

4.1. *Comparison of two reversible Markov chains*

Under reversibility, by the Dirichlet principle we have that for a finite set A

$$\mathrm{Cap}(A) = \inf_{f \in \mathbb{R}^V : f \upharpoonright A \equiv 1, 0 \le f \le 1, \mathrm{supp}(f) \text{ is finite}} \mathcal{E}_P(f, f), \tag{4.1}$$

where $f \upharpoonright A$ is the restriction of f to A, $\mathrm{supp}(f) := \{v \in V : f(v) \ne 0\}$ is the support of f and $\mathcal{E}_P(f, f)$ is as in (2.1). The following standard lemma is an immediate consequence of (4.1).

Lemma 4.1. *Let P, Q be two transition kernels on the same (countable) state space Ω. Assume that P and Q are reversible w.r.t. π_P and π_Q, respectively. Assume further that for all $f : \Omega \to \mathbb{R}$ we have that $\mathcal{E}_P(f, f) \le M \mathcal{E}_Q(f, f)$. Then for all finite $A \subset \Omega$ we have that*

$$\mathrm{Cap}_P(A) \le M \, \mathrm{Cap}_Q(A),$$

where $\mathrm{Cap}_P(A)$ and $\mathrm{Cap}_Q(A)$ are the capacities of A w.r.t. P and Q, respectively. In particular, if Q is recurrent then so is P.

Definition 4.2. Let $G := (V, E, (c_e)_{e \in E})$ and $G' := (V, E', (c'_e)_{e \in E'})$ be two connected networks. For every $(x, y) \in \overrightarrow{E}'$ let $\mathcal{P}_{x,y}(G)$ be the collection of all non-empty (oriented) paths from x to y in G. Denote the number of times the (directed) edge (u, v) appears in a path γ by

$$r\big((u, v), \gamma\big) := \big| \{ i : \gamma_i = u, \gamma_{i+1} = v \} \big|.$$

We say that $\Phi := (\Phi_{x,y})_{(x,y) \in \overrightarrow{E}'}$ is a *G-G' flow* (or *P-P' flow*, where P and P' are the transition kernels of G and G', respectively) if for every $\{x, y\} \in E'$ the following hold:

   (i) $\Phi_{x,y}$ is a map from $\mathcal{P}_{x,y}(G)$ to $[0, c'_{x,y}]$,
   (ii) $\sum_\gamma \Phi_{x,y}(\gamma) = c'_{x,y}$ and
   (iii) $\Phi_{x,y}(\gamma) = \Phi_{y,x}(\gamma^{\mathrm{r}})$ for every γ.

The *congestion* of $(u, v) \in \overrightarrow{E}$ w.r.t. Φ is defined as

$$A_{u,v}(\Phi) := \frac{1}{c_{u,v}} \sum_{(x,y) \in \overrightarrow{E}'} \sum_{\gamma \in \mathcal{P}_{x,y}(G)} r\big((u, v), \gamma\big) |\gamma| \Phi_{x,y}(\gamma).$$

The *congestion* of Φ is defined as $A(\Phi) := \sup \{ A_{u,v}(\Phi) : (u, v) \in \overrightarrow{E} \}$.

The following lemma provides a standard technique for utilizing Lemma 4.1. The lemma is essentially Theorem 2.3 from [7]. While the formulation in [7] is slightly different, the difference is non-essential (cf. Theorem 1 in [10] and the remark thereafter).

Lemma 4.3. *In the setup of Definition 4.2 Let P and P' be two reversible transition kernels on the same state space. Let Φ be a P–P' flow. Then $\mathcal{E}_{P'}(f) \le A(\Phi) \mathcal{E}_P(f)$ for all $f \in \mathbb{R}^V$.*

We now present a variant of Lemma 4.3, useful for comparing two networks with different state spaces.

Definition 4.4. Let $G := (V, E, (c_e)_{e \in E})$ and $G' := (V', E', (c'_e)_{e \in E'})$ be two connected networks. Using notation from Definition 4.2, we say that $\phi : V' \to V$ is a *rough-embedding* from G' to G if there exist a map $\Phi : \overrightarrow{E}' \to \bigcup_{x,y:(x,y)\in\overrightarrow{E}'} \mathcal{P}_{\phi(x),\phi(y)}(G)$ and some constants $\alpha, \beta > 0$ such that the following hold:

(i) $\Phi((x, y)) \in \mathcal{P}_{\phi(x),\phi(y)}(G)$ and $\Phi((y, x))$ is its reversal for every $(x, y) \in \overrightarrow{E}'$,

(ii) for every $(x, y) \in \overrightarrow{E}'$ we have that $\frac{1}{c_{x,y}} \le \alpha \sum_{e \in \Phi((x,y))} \frac{1}{c_e}$, and

(iii) every $(u, v) \in \overrightarrow{E}$ belongs to the image (under Φ) of at most β edges $(x, y) \in \overrightarrow{E}'$.

Fact 4.5 (e.g. [19] Theorem 2.17). If there is a rough-embedding from G' to G and G' is transient, then G is also transient.

We now recall the notion of a rough-isometry.

Definition 4.6. Let $G_i := (V_i, E_i)$ (where $i = 1, 2$) be two bounded degree graphs. For $u, v \in V_i$, let $d_i(u, v)$ be graph distance (w.r.t. G_i) between u and v (i.e., the number of edges along the shortest path in G_i between u and v). We say that $f : V_1 \to V_2$ is a K-*rough-embedding* from G_1 to G_2 if the map f can stretch or shrink distances only by a multiplicative and additive factor of at most K in the following sense:

$$\forall u, v \in V_1, \quad \frac{1}{K}d_1(u, v) - 1 \le d_2\big(f(u), f(v)\big) \le K\big(d_1(u, v) + 1\big). \tag{4.2}$$

We say that it is a K-*rough-isometry* from G_1 to G_2 if (4.2) holds and in addition f is also "almost surjective" in the following sense: For every $w \in V_2$, there exists some $v \in V_1$ such that $d_2(f(v), w) \le K$. We say that f is a *rough-embedding* (respectively, *rough-isometry*) if for some $K \in \mathbb{N}$ it is a K-rough-embedding (respectively, K-rough-isometry). We say that G_1 and G_2 are *roughly-isometric* if there exists a rough-isometry between them.

Remark 4.7. Let $G_i := (V_i, E_i)$ (where $i = 1, 2$) be two bounded degree graphs. It is easy to see that if $f : V_1 \to V_2$ is a rough-embedding from G_1 to G_2 (in the sense of Definition 4.6), then f is also a rough-embedding from $(V_1, E_1, (1)_{e \in E_1})$ to $(V_2, E_2, (1)_{e \in E_2})$ (i.e., from SRW on G_1 to SRW on G_2) in the sense of Definition 4.4.

Remark 4.8. Let $G_i := (V_i, E_i)$ ($i = 1, 2$) be two bounded degree graphs. It is easy to verify that if f is a rough-isometry from G_1 to G_2, then any $g : V_2 \to V_1$ such that $g(u) \in \{v \in V_1 : d_2(f(v), u) = \min_{v' \in V_1} d_2(f(v'), u)\}$ for all $u \in V_2$, is also a rough-isometry.

In the bounded degree setup the relation of being roughly-isometric is an equivalence relation on graphs which preserves recurrence/transience (this can be derived from Fact 4.5 and Remarks 4.7–4.8).

Fact 4.9 (e.g. [19] Theorem 2.17 and Proposition 2.18). Let $G_i := (V_i, E_i)$ (where $i = 1, 2$) be two connected graphs of uniformly bounded degree. Let $c^{(i)} := (c_e^{(i)})_{e \in E}$ ($i = 1, 2$) be some edge weights which are uniformly bounded from above and below by some positive constants.

(a) The random walk corresponding to $(V_1, E_1, c^{(1)})$ is transient iff SRW on G_1 is.

(b) If there exists a rough-embedding from G_1 to G_2 and SRW on G_1 is transient, then also SRW on G_2 is transient.

(c) If there exists a rough-isometry from G_1 to G_2 then SRW on G_1 is transient iff SRW on G_2 is transient.

We now recall the notion of 'lumping' (or 'short-circuiting') a network and the fact that this operation preserves transience.

Definition 4.10. Let $G = (V, E, (c_e)_{e \in E})$ be an infinite connected network on (V, E) with edge weights $(c_e)_{e \in E}$. Recall that for $v \in V$ we write $c_v := \sum_{u:u \sim v} c_{u,v}$. Let $A_1, A_2, \ldots$ be a partition of V into non-empty disjoint sets such that $\sum_{v \in A_i} c_v < \infty$ for all i. Consider the network $G_{(A_i)_{i \in \mathbb{N}}} := (\mathbb{N}, \tilde{E}, (\tilde{c}_e)_{e \in \tilde{E}})$ obtained from G by collapsing each A_i into a single state i and setting $\tilde{c}_{i,j} := \sum_{a_i \in A_i, a_j \in A_j} c_{a_i,a_j}$ for all $(i, j) \in \mathbb{N}^2$ (this operation is also called lumping A_i together, or short-circuiting A_i, for all i) and setting $\tilde{E} := \{\{i, j\} : \tilde{c}_{i,j} > 0\}$. We say that $G_{(A_i)_{i \in \mathbb{N}}}$ is a *factor* of G.

The following is a consequence of Rayleigh"s monotonicity principle.

Fact 4.11. In the setup of Definition 4.10, if G is transient then so is any factor of G.

4.2. *Comparison of a Markov chain to its additive symmetrization*

The connection between reversible networks, electrical networks and potential theory is classical (see e.g., [8] and [19, Ch. 2]). It was only in recent years that this connection was extended to the non-reversible setup in several extremely elegant works. The first progress on the non-reversible front was made by Doyle and Steiner [9] who derived an extremal characterization of the commute-time between two states, which shows that commute-times in the additive symmetrization of an irreducible chain cannot be smaller than the corresponding commute-times in the original chain.

Gaudillière and Landim [12] extended much of the classic potential theory to the non-reversible setup and derived several extremal characterizations for the capacity between two disjoint sets. In particular, they showed [12, Lemma 2.5] that the capacity between two disjoint sets of an irreducible Markov chain with a stationary measure π is at least as large as that of the additive symmetrization of the chain.

Theorem 4.12 ([12] Lemma 5.1). *Consider an irreducible Markov chain $\mathbf{X} = (X_k)_{k=0}^{\infty}$ on a countable state space V with a stationary measure π and transition probabilities given by P. Let $\mathbf{X}^s = (X_k^s)_{k=0}^{\infty}$ be the Markov chain corresponding to $S := \frac{1}{2}(P + P^*)$, where $P^*(x, y) := \frac{\pi(y)}{\pi(x)} P(y, x)$. Then, if $\mathbf{X}^s$ is transient then $\mathbf{X}$ is also transient.*

We shall make use of Theorem 4.12 only in the proof of Parts (i) and (iv) of Theorem 1. For the rest of our results, our method of proof allows for a direct comparison of the considered non-reversible chain with some reversible Markov chain which can be compared with the SRW on the considered graph.

5. Equivalence between SRW on G and the additive symmetrization of the k-NBRW

Let $G = (V, E)$ be a connected graph of bounded degree. In order to take advantage of Theorem 4.12 we first need to compare the additive symmetrization of the k-NBRW with SRW on G. In this section we do so by utilizing the machinery from Section 4.1–4.2. We will later show that in all of the cases considered in Theorems 1 and 2 the k-NBRW has a stationary measure which is pointwise uniformly bounded from above and from below. Hence by Fact 4.9, instead of the additive symmetrization of the k-NBRW, we may consider SRW on the graph supporting its transitions. We start with the case $k = 1$.

Recall that the *line graph* of G is defined as $G_E := (E, E')$, where $\{e, e'\} \in E'$ iff $e \cap e' \neq \varnothing$ and $e \neq e'$. Recall that $\overrightarrow{E} := \{(x, y) : \{x, y\} \in E\}$ is the collection of directed edges of G. Following [21] we define the *symmetric oriented line graph* of G to be $G_{\overrightarrow{E}} := (\overrightarrow{E}, F)$ the graph supporting the transitions of the NBRW on G. That is, $\{\overrightarrow{e}, \overrightarrow{f}\} \in F$ iff either

(a) $\overrightarrow{e}_+ = \overrightarrow{f}_-$ and $\mathbf{1}\{\deg(\overrightarrow{e}_+) = 1 \text{ or } \overrightarrow{e}_- \neq \overrightarrow{f}_+\} = 1$, or
(b) $\overrightarrow{e}_- = \overrightarrow{f}_+$ and $\mathbf{1}\{\deg(\overrightarrow{f}_+) = 1 \text{ or } \overrightarrow{f}_- \neq \overrightarrow{e}_+\} = 1$.

We shall make crucial use of the following simple observation:

Lemma 5.1. *Let G be a finite connected graph of bounded degree. The graph supporting the transitions of the additive symmetrization of the NBRW on G is $G_{\overrightarrow{E}}$. Moreover, both SRW on $G_{\overrightarrow{E}}$ and the additive symmetrization of the NBRW*

on G are reversible w.r.t. the counting measure on $\overrightarrow{E}$, and their minimal transition probability between neighboring directed edges is at least half the inverse of the maximal degree of G. In particular, either both are transient or both are recurrent.

Proof. The statement of the last sentence follows from the rest of the lemma via Fact 4.9 (by noting that when the reversible measure is the counting measure, the edge-weights are simply the transition probabilities). The rest can be verified by inspection. □

Before describing some relations between G, G_E and $G_{\overrightarrow{E}}$ we need a simple lemma. Its proof is deferred to Section A.2.

Lemma 5.2. *Let $G = (V, E)$ be an infinite locally finite connected graph of minimal degree at least* 2. *There exists a sink-free and source-free orientation of E, i.e., a map $f : E \to \overrightarrow{E}$ such that*:

  (i) *$f(\{u, v\})$ is either (u, v) or (v, u) for every $\{u, v\} \in E$, and*

  (ii) *every $v \in V$ is the tail of at least one $\overrightarrow{e} \in \overrightarrow{E}$ in the image of f and also the head of at least one $\overrightarrow{e}$ in the image of f (i.e. $V = \{f(e)_+ : e \in E\} = \{f(e)_- : e \in E\}$).*

Recall that for $\gamma = (a, b) \in \overrightarrow{E}$ its reversal (b, a) is denoted by γ^{r}. For a graph $H = (V_H, E_H)$ and $x, y \in V_H$ we denote by $\mathrm{dist}_H(x, y)$ the graph distance of x from y in H, defined as number of edges along of the shortest path from x to y in H.

Lemma 5.3. *Let G be an infinite connected graph of bounded degree and minimal degree at least* 2. *Then*

  (1) *The graphs G and G_E are roughly-isometric.*

  (2) *SRW on G_E is a factor of SRW on $G_{\overrightarrow{E}}$ (in the sense of Definition 4.10).*

  (3) *There exists a rough-embedding from G_E to $G_{\overrightarrow{E}}$.*

  (4) *If there exists some $K \in \mathbb{N}$ such that $\mathrm{dist}_{G_{\overrightarrow{E}}}(\gamma, \gamma^{\mathrm{r}}) \le K$ for every $\gamma \in \overrightarrow{E}$ then there exists a rough-isometry from G_E to $G_{\overrightarrow{E}}$.*

  (5) *If condition (1) holds then there exists some $K \in \mathbb{N}$ such that $\mathrm{dist}_{G_{\overrightarrow{E}}}(\gamma, \gamma^{\mathrm{r}}) \le K$ for every $\gamma \in \overrightarrow{E}$.*

Proof. For Part (2) observe that G_E is obtained from $G_{\overrightarrow{E}}$ by identifying (lumping) every $\overrightarrow{e} \in \overrightarrow{E}$ with its reversal.

Fix a sink-free and source-free orientation f (as in Lemma 5.2). For Part (1) observe that the map $e \to f(e)_+$ from E to V is clearly a 2-rough-isometry from G_E to G (for this we do not need to use the fact that f is sink-free and source-free). For Part (3) we now argue that f is a 2-rough-embedding from G_E to $G_{\overrightarrow{E}}$. In fact, we argue that $\mathrm{dist}_{G_E}(e, e') \le \mathrm{dist}_{G_{\overrightarrow{E}}}(f(e), f(e')) \le 2\,\mathrm{dist}_{G_E}(e, e')$ for all $e, e' \in E$. The first inequality is obvious. For the second inequality it suffices to consider the case that $\mathrm{dist}_{G_E}(e, e') = 1$. If $e \cap e' = \{v\}$ and $f(e)$ and $f(e')$ are not adjacent in $G_{\overrightarrow{E}}$, then either $v = f(e)_+ = f(e')_+$ or $v = f(e)_- = f(e')_-$. In the first case (by condition (ii) in the definition of a sink-free and source-free orientation) there must be some $\{v, v'\} \in E \setminus \{e, e'\}$ such that $f(\{v, v'\})_- = v$ (in the second case $f(\{v, v'\})_+ = v$) and so $f(e)$ and $f(e')$ are both adjacent to $f(\{v, v'\})$ in $G_{\overrightarrow{E}}$.

For Part (4) observe that under the assumption that $\mathrm{dist}_{G_{\overrightarrow{E}}}(\gamma, \gamma^{\mathrm{r}}) \le K$ for every $\gamma \in \overrightarrow{E}$, we have that the map f (defined above) is a rough-isometry.

For Part (5) note that if $v \in V$ is of degree at least 3, then similar reasoning as in the proof of Part (iii) yields that for every pair of adjacent vertices $u \sim v$ we have that (u, v) and (v, u) are of distance at most 3 in $G_{\overrightarrow{E}}$ (namely, if w, z are other neighbors of v then $((u, v), (v, w), (z, v), (v, u))$ is a path from (u, v) to (v, u) in $G_{\overrightarrow{E}}$). □

Corollary 5.4. *Let G be a connected infinite graph of bounded degree. Then SRW on G is transient iff the additive symmetrization of the NBRW on G is transient.*

Proof. We shall show that w.l.o.g. we may assume that the minimal degree of G is at least 2, and so the assertion of the corollary follows from Lemmas 5.1 and 5.3 in conjunction with Fact 4.9. We argue that by repeatedly deleting

all degree one vertices, while adding a self-loop at each remaining $v \in V$ such that at least one of its neighbors was deleted, we obtain a graph $G' = (V(G'), E(G'))$ of minimal degree at least 2 such that:

(i) SRW on G is transient iff SRW on G' is transient, and also

(ii) the additive symmetrization of the NBRW on G is transient iff the additive symmetrization of the NBRW on G' is transient.

While (i) is clear, (ii) requires some justification. Recall the notion of the induced chain on a set from Section 2. We denote the induced chain on the directed edges of G' w.r.t. the additive symmetrization of the NBRW on G by $\mathbf{Y}$. Clearly, $\mathbf{Y}$ is transient iff the additive symmetrization of the NBRW on G is transient, and so it suffices to show that $\mathbf{Y}$ is transient iff $\mathbf{Z}$ is transient, where $\mathbf{Z}$ is the additive symmetrization of the NBRW on G'. Observe that the graph $\widehat{G} = (V(\widehat{G}), E(\widehat{G}))$ supporting the transitions of $\mathbf{Y}$ is roughly-isometric to $\widehat{G}' = (V(\widehat{G}'), E(\widehat{G}'))$ the graph supporting the transitions of $\mathbf{Z}$ and that $\mathbf{Y}$ and $\mathbf{Z}$ are both reversible w.r.t. the counting measures on their state spaces, $\pi_{V(\widehat{G})}$ and $\pi_{V(\widehat{G}')}$, respectively.[3] (For the existence of the claimed rough-isometry, it is crucial that loops were added to G' at each remaining $v \in V$ such that at least one of its neighbors was deleted. The rough-isometry can be taken to be the inclusion map from $V(\widehat{G})$ to $V(\widehat{G}')$, which is a 2-rough-isometry.)

The edge-weights corresponding to $\mathbf{Z}$ are simply its transition probabilities and hence are clearly uniformly bounded from above and from below (as the maximal degree of G is finite; Recall that in general, for a Markov chain whose transition kernel P is reversible w.r.t. a measure π the edge-weight of edge $\{x, y\}$ can be taken to be $\pi(x)P(x, y)$). We claim that the same holds for $\mathbf{Y}$, which in conjunction with the fact that $\widehat{G}$ and $\widehat{G}'$ are roughly-isometric concludes the proof, using Parts (a) and (c) of Fact 4.9.

Taking $\pi_{V(\widehat{G})}$ for the stationary measure of $\mathbf{Y}$, the corresponding edges weights for $\mathbf{Y}$ are simply with the transition probabilities. Thus its enough to show that these are uniformly bounded from below (as they are bounded from above by 1). Note that any transition which is also allowed for the NBRW on G has transition probability which is at least half of the inverse of the maximal degree of G. Hence it suffices to verify this for transition probabilities from (b, a) to (v, a) and from (a, b) to (a, v) with $a \in A$ and $b, v \in V(G')$ such that $\{a, u\}, \{a, b\} \in E(G')$, where A is the set of vertices of $V(G')$ to which a self-loop was added in G'. In other words, A is the set of vertices of G which are connected in G to a degree one vertex via a path whose internal vertices all have degree two.

To see that such transitions (from (b, a) to (v, a) and from (a, b) to (a, v) with $a \in A$ and $b, v \in V(G')$ such that $\{a, u\}, \{a, b\} \in E(G')$) have probability bounded from below, consider the scenario that the additive symmetrization of the NBRW on G goes from (b, a) to (a, a') for some deleted vertex a' and then goes from (a, a') to (v, a). Similarly, it may move from (a, b) to (a', a) for some deleted vertex a' and then move from (a', a) to (a, v). Here by "deleted vertex" we mean a vertex belonging to $V \setminus V(G')$.

These scenarios show that the transition probabilities w.r.t. $\mathbf{Y}$ from (b, a) to (v, a) and from (a, b) to (a, v), with $a \in A$ and $b, v \in V(G')$, are indeed uniformly bounded from below, as desired. $\qquad \square$

We now treat the case $k > 1$. Recall the definitions of $\mathcal{P}_k^{\mathrm{e}}, \mathcal{P}_k^{\mathrm{v}}, \Omega_k^{\mathrm{e}}, \Omega_k^{\mathrm{v}}, P_{k,\mathrm{e}}$ and $P_{k,\mathrm{v}}$ from Section 3.1–3.2.

Definition 5.5. Let $G_k^{\mathrm{e}} := (\Omega_k^{\mathrm{e}}, E_k^{\mathrm{e}})$ and $G_k^{\mathrm{v}} := (\Omega_k^{\mathrm{v}}, E_k^{\mathrm{v}})$, where $\{a, b\} \in E_k^{\mathrm{e}}$ (respectively, $\in E_k^{\mathrm{v}}$) if and only if $P_{k,\mathrm{e}}(a, b) + P_{k,\mathrm{e}}(b, a) > 0$ (respectively, $P_{k,\mathrm{v}}(a, b) + P_{k,\mathrm{v}}(b, a) > 0$).

Observe that $G_1^{\mathrm{e}} = G_1^{\mathrm{v}} = G_{\overrightarrow{E}}$. More generally, under condition (3_k^{e}) (respectively, (3_k^{v})) we have that $\mathcal{P}_k^{\mathrm{e}} = \Omega_k^{\mathrm{e}}$ (respectively, $\mathcal{P}_k^{\mathrm{v}} = \Omega_k^{\mathrm{v}}$) and we may consider the graph $H_k^{\mathrm{e}} := (U_k^{\mathrm{e}}, F_k^{\mathrm{e}})$ (respectively, $H_k^{\mathrm{v}} := (U_k^{\mathrm{v}}, F_k^{\mathrm{v}})$) obtained from G_k^{e} (respectively, G_k^{v}) by identifying each $a \in \mathcal{P}_k^{\mathrm{e}}$ (respectively, $a \in \mathcal{P}_k^{\mathrm{v}}$) with its reversal. More precisely, $U_k^{\mathrm{e}} := \{\{a, a^{\mathrm{r}}\} : a \in \mathcal{P}_k^{\mathrm{e}}\}$ and $\{\{a, a^{\mathrm{r}}\}, \{b, b^{\mathrm{r}}\}\} \in F_k^{\mathrm{e}}$ iff $\sum_{w \in \{a, a^{\mathrm{r}}\}, w' \in \{b, b^{\mathrm{r}}\}} P_{\mathrm{e},k}(w, w') > 0$. The relations between G_k^{e} and H_k^{e} and between G_k^{v} and H_k^{v} are analogous to the one between $G_{\overrightarrow{E}}$ and G_E. Similarly, when G is of bounded degree and condition (3_k^{e}) (respectively, (3_k^{v})) holds, the relation between SRW on G_k^{e} (respectively, G_k^{v}) and the additive symmetrization of the edge (respectively, vertex) k-NBRW on G is analogous to the relation between SRW on $G_{\overrightarrow{E}}$

---

[3]Here we have used the fact that $\pi_{V(\widehat{G})}$ is the restriction to $V(\widehat{G})$ of $\pi_{\overrightarrow{E}}$ the counting measure on the directed edges of G, and that the additive symmetrization of the NBRW on G, whose restriction to $V(\widehat{G})$ is $\mathbf{Y}$, is reversible w.r.t. $\pi_{\overrightarrow{E}}$. (Recall from Section 2 that for an irreducible Markov chain which is reversible w.r.t. a measure π, for any recurrent subset B (i.e., one which is a.s. visited by the chain) the restriction of π to B is reversible for the induced chain on B.)

and the additive symmetrization of the NBRW described in Lemma 5.1, provided the stationary measure $\hat{\pi}$ w.r.t. we are taking the symmetrization satisfies $\sup_{x,y} \frac{\hat{\pi}(x)}{\hat{\pi}(y)} < \infty$. For the edge k-NBRW this will be proved in Lemma 7.1. For the vertex k-NBRW we are only able to prove this under the additional assumption that the graph is vertex-transitive (see the proof of Part (iv) of Theorem 1).

The following lemma is the $k > 1$ analog of Lemma 5.3.

Lemma 5.6. *Let G be a bounded degree graph satisfying condition* (3_k^{e}) *(resp.* (3_k^{v})*). Then*

(1) *The graphs G and H_k^{e} (respectively, H_k^{v}) are roughly-isometric.*

(2) *SRW on H_k^{e} (respectively, H_k^{v}) is a factor of SRW on G_k^{e} (respectively, G_k^{v}).*

(3) *If condition (2_k^{e}) (respectively, (2_k^{v})) holds there exists some $K \in \mathbb{N}$ such that $\mathrm{dist}_{G_k^{\mathrm{e}}}(\gamma, \gamma^{\mathrm{r}}) \leq K$ for every $\gamma \in \mathcal{P}_k^{\mathrm{e}}$ (respectively, $\mathrm{dist}_{G_k^{\mathrm{v}}}(\gamma, \gamma^{\mathrm{r}}) \leq K$ for every $\gamma \in \mathcal{P}_k^{\mathrm{v}}$) and so G_k^{e} and H_k^{e} (respectively, G_k^{v} and H_k^{v}) are roughly-isometric.*

Proof. The proofs of Parts (1) and (2) are completely analogous to those of Parts (1) and (2) of Lemma 5.3 and are thus omitted. The proof of Part (3) is analogous to that of Parts (3)–(5) of Lemma 5.3 and is thus omitted. (Note that in the proof of Part (3) of Lemma 5.3 we did not assume condition (2) holds, whereas here we assume condition (2_k^{e}) (respectively, (2_k^{v})) holds. This makes the proof of Part (3) of Lemma 5.6 easier, by eliminating the need to prove a generalization of Lemma 5.2 for $k > 1$.) □

6. An overview of our approach

In this section as a warm up and as motivation for what comes we prove Theorem 3. Recall that $\pi_{\overrightarrow{E}}$ is the counting measure on $\overrightarrow{E}$. Let $\overrightarrow{\pi} := \pi_{\overrightarrow{E}}$. Let B be the transition kernel of the NBRW $(S_n^{\mathrm{NB}})_{n=0}^{\infty}$. While B is non-reversible, it satisfies a symmetry similar to the detailed balance equation:

$$\forall a, b \in \overrightarrow{E}, \quad \overrightarrow{\pi}(a)B(a,b) = \overrightarrow{\pi}(b^{\mathrm{r}})B(b^{\mathrm{r}}, a^{\mathrm{r}}) = \overrightarrow{\pi}(b)B(b^{\mathrm{r}}, a^{\mathrm{r}}).$$

It follows by induction that for every $n \in \mathbb{N}$ and $a_0, a_1, \ldots, a_n \in \overrightarrow{E}$

$$\overrightarrow{\pi}(a_0)\mathrm{P}_{a_0}\big(S_1^{\mathrm{NB}} = a_1, \ldots S_n^{\mathrm{NB}} = a_n\big) = \overrightarrow{\pi}(a_n)\mathrm{P}_{a_n^{\mathrm{r}}}\big(S_1^{\mathrm{NB}} = a_{n-1}^{\mathrm{r}}, \ldots S_n^{\mathrm{NB}} = a_0^{\mathrm{r}}\big). \tag{6.1}$$

While $\overrightarrow{\pi}$ is redundant above, we include it in order to demonstrate a general principle. By summing over all $a_0, \ldots, a_n$ with $a_0 = a$ and $a_n = b$, we get for all $a, b \in \overrightarrow{E}$ and $n \in \mathbb{N}$ that

$$\overrightarrow{\pi}(a)B^n(a,b) = \overrightarrow{\pi}(b^{\mathrm{r}})B^n(b^{\mathrm{r}}, a^{\mathrm{r}}) = \overrightarrow{\pi}(b)B^n(b^{\mathrm{r}}, a^{\mathrm{r}}). \tag{6.2}$$

Observe that (6.2) implies that $\overrightarrow{\pi}$ is stationary for B and so also for B_p (where B_p is as in Theorem 3). Indeed, by (6.2) $\sum_b \overrightarrow{\pi}(b)B(b,a) = \overrightarrow{\pi}(a)\sum_b B(a^{\mathrm{r}}, b^{\mathrm{r}}) = \overrightarrow{\pi}(a)$.

In order to exploit (6.2) in the proof of Theorem 3 we consider an auxiliary chain in which every $a \in \overrightarrow{E}$ is identified with a^{r}. The auxiliary chain is defined by looking at the p-BRW at random times at which for every $a \in \overrightarrow{E}$ the chain is equally likely to be at a or at a^{r}. The mechanism which allows us to construct such random times is the fact that the p-BRW backtracks at each step with probability p.

From the construction of the auxiliary chain it will be clear that it is recurrent iff the original chain is recurrent. The main idea is that the auxiliary chain is reversible and thus can be compared with SRW on G via the standard comparison techniques from Section 4. That is, it is recurrent iff the SRW on G is recurrent.

In Section 7 we show that in the setups of Theorems 1 and 2 either the k-NBRW or some related auxiliary chain enjoys similar symmetry relations as (6.1)–(6.2). Using similar reasoning as above this will allow us to find a stationary measure for the k-NBRW (or in the proof of Theorem 2, for a related auxiliary chain). We will then look at another auxiliary chain, obtained by looking at the k-NBRW (or the aforementioned related chain) at random times at which for every a it is equally likely to be at a or at a^{r}. The mechanism which allows us to construct this auxiliary chain is related to condition (2_k^{e}). Again, crucially, the auxiliary chain is reversible and thus amenable to comparison techniques.

Remark 6.1. In the setup of Theorem 1, analogs of (6.1)–(6.2) hold for the edge k-NBRW (see (7.4)) but not for the vertex k-NBRW (when $k > 1$). In general, such symmetry may fail if the walk may get stuck, because it is possible that a certain trajectory is possible in one direction but not in the opposite direction. In the setup of Theorem 2 we do not assume that the k-NBRW cannot get stuck. To overcome the aforementioned difficulty we work with a more symmetric auxiliary chain and exploit the fact that the group is Abelian.

Proof of Theorem 3. Let $G = (V, E)$ and $p \in (0, 1)$. Consider the lazy version of B_p whose transition kernel is $\frac{1}{2}(I + B_p)$ (where I is the identity operator). Denote this chain by $\mathbf{Y} := (Y_t)_{t=0}^{\infty}$. We may think of this chain as first picking a candidate for its next position according to B_p and then flipping a fair coin to decide whether to accept the candidate or to stay put. Denote the candidate that was picked at time t (for time $t + 1$) by Z_{t+1}.

Let $(x, y) \in \vec{E}$. Consider the case that the initial distribution is uniform on $\{(x, y), (y, x)\}$. Let $\tau_0 = -1$ and $e_0 := \{x, y\}$. Define inductively $\tau_n := \inf\{t > \tau_{n-1} : Z_{t+1} = Y_t^{\mathrm{r}}\}$ and $e_n := \{x_n, y_n\}$, where x_n and y_n are the endpoints of Z_{τ_n+1}. In words, τ_n is the first time after time τ_{n-1} at which the chain $\mathbf{Y}$ "attempts" to backtrack (i.e., the candidate directed edge it picked is the reversal of the current directed edge). Observe that given $(Y_i)_{i=0}^{\tau_n}$ and $(\tau_i)_{i=0}^{n}$ the conditional distribution of Y_{τ_n+1} is the uniform distribution on $\{(x_n, y_n), (y_n, x_n)\}$ (because either $Y_{\tau_n+1} = Z_{\tau_n+1}$ or $Y_{\tau_n+1} = Z_{\tau_n+1}^{\mathrm{r}}$, depending on the result of the coin toss).

Consider the process $\mathbf{Q} := (e_n)_{n=0}^{\infty}$ on E. To conclude the proof we show that

   (i) The process $\mathbf{Q}$ is a Markov chain. Its transition kernel W is symmetric (i.e., $W(a, b) = W(b, a)$ for all $a, b \in E$) and so $\mathbf{Q}$ is reversible w.r.t. the counting measure π_E on E.
   (ii) The p-BRW is recurrent iff $\mathbf{Q}$ is recurrent.
   (iii) The SRW on G is recurrent iff $\mathbf{Q}$ is recurrent.

We first prove (i). Let $B_{\mathrm{L}} := \frac{1}{2}(I + B)$ be the transition kernel of the lazy version of the NBRW. By (6.2) for every $\alpha, \beta \in \vec{E}$ and $n \in \mathbb{N}$ we have that

$$B_{\mathrm{L}}^n(\alpha, \beta) = \sum_{i=0}^{n} 2^{-n} \binom{n}{i} B^i(\alpha, \beta) = \sum_{i=0}^{n} 2^{-n} \binom{n}{i} B^i(\beta^{\mathrm{r}}, \alpha^{\mathrm{r}}) = B_{\mathrm{L}}^n(\beta^{\mathrm{r}}, \alpha^{\mathrm{r}}). \tag{6.3}$$

For every $\{x, y\} \in E$ let $D_{\{x,y\}} := \{(x, y), (y, x)\}$. Using the fact the Y_{τ_n+1} is uniformly distributed on D_{e_n} we get that for every $a, b \in E$ we have that

$$W(a, b) = \frac{1}{2} \sum_{i=0}^{\infty} \sum_{w \in D_a, w' \in D_b} B_{\mathrm{L}}^i(w, w')(1 - p)^i p. \tag{6.4}$$

By (6.3) W is indeed symmetric.

We now prove Claim (ii). Since $\mathbf{Y}$ is a lazy version of the p-BRW, it is recurrent iff the p-BRW is recurrent. Observe that above we constructed a coupling of $\mathbf{Y}$ and $\mathbf{Q}$ (in other words, while $\mathbf{Q}$ can be defined via (6.4) we constructed it using the process $\mathbf{Y}$). It is easy to see that in this coupling the expected number of (all) returns to the starting point of one chain is within a constant factor from that of the other (since every return of one chain has a constant probability of becoming a return for the other). Hence either both expectations are infinite or both are finite.

We now prove (iii). We first prove that w.l.o.g. we may assume that G has minimal degree at least 2. By repeatedly deleting degree one vertices we obtain a graph $H' = (V', E')$ with minimal degree two. For every $v \in V'$ let $n(v)$ be the number of degree one vertices of G which are connected to it in G via a path whose internal vertices are all of degree 2. We then consider the graph $\widehat{H} = (V', \widehat{E})$ obtained from H' by adding $n(v)$ self-loops at each $v \in V'$.

Clearly SRW on G is transient iff SRW on $\widehat{H}$ is transient. We argue that the same is true for the p-BRW. We now argue that for every directed edge (u, v) of H' we have that the expected number of visits to this edge by the p-BRW on G is the same as this expectation for the p-BRW on $\widehat{H}$ (when both p-BRWs start at the same directed edge). This is because the two p-BRWs can be coupled so that the former enters some path $(v_0, \ldots, v_\ell)$ such that

$\deg(v_0) \geq 3$, $\deg(v_\ell) = 1$ and $\deg(v_i) = 2$ for $i \in [\ell - 1]$ and ultimately leaves it from (v_1, v_0) to some (v_0, u) with $u \in V'$ (respectively, to some (v_0, v_1'), where $(v_0, v_1' \ldots, v_{\ell'}')$ is such that $\deg(v_{\ell'}') = 1$ and $\deg(v_i') = 2$ for $i \in [\ell - 1]$) iff the latter crosses the self-loop at v_0 corresponding to the path $(v_0, \ldots, v_\ell)$ and then (possibly after "backtracking" by crossing the same loop for some number of times) crosses to (v_0, u) (respectively, crosses to the self-loop which corresponds to the path $(v_0, v_1' \ldots, v_{\ell'}')$). We can further couple their trajectories together until the next time the former enters such a path, while the latter crosses a corresponding self-loop. Continuing in this fashion we obtain a coupling in which the two p-BRWs visit every directed edge (u, v) of H' the same number of times. In particular, one is recurrent iff the other is recurrent.

Recall that the *line graph* of G is defined as $G_E := (E, E')$, where $\{e, e'\} \in E'$ iff $e \cap e' \neq \varnothing$ and $e \neq e'$. By Lemma 5.3 SRW on G is recurrent iff SRW on G_E is recurrent. To conclude the proof of Claim (iii) we now show that $\mathbf{Q}$ is recurrent iff SRW on the line graph of G is recurrent. In order to do so we utilize the comparison technique from Lemmas 4.1 and 4.3. Namely, using the notation and terminology from Definition 4.2, by Lemmas 4.1 and 4.3 it suffices to show that there exists a flow of finite congestion from the network corresponding to $\mathbf{Q}$ to the network corresponding to SRW on G_E and vice-versa.

Let $\tilde{E} := E' \cup \{\{e, e\} : e \in E\}$. Denote by $\tilde{G} := (E, \tilde{E})$ the graph obtained from G_E by adding a self-loop at each site of G_E. Below we consider $\tilde{G}$ instead of G_E. Denote the transition kernel of SRW on $\tilde{G}$ by U. The networks corresponding to U and W can be represented, respectively, as $(E, \tilde{E}, (1)_{w \in \tilde{E}})$ and $(E, F, (c_w)_{w \in F})$, where

$$F := \{\{e, e'\} : e, e' \in E\} = \{\{e, e'\} : e, e' \in E \text{ such that } W(e, e') > 0\}$$

and $c_{e,e'} := W(e, e')$ for every $e, e' \in E$. We start with the easier direction. Observe that there exists a constant $M > 0$ such that $U(e, e') \leq MW(e, e')$ for all $e, e' \in E$. Thus we may define a W-U flow of finite congestion by mapping every (e, e') such that $\{e, e'\} \in \tilde{E}$ to itself (with weight 1).

We now construct a U-W flow Φ of finite congestion. Loosely speaking, the flow Φ assigns to every path the probability that the chain $\mathbf{Y}$ follows that path between two consecutive steps of $\mathbf{Q}$. This is not a precise description because $\mathbf{Y}$ is defined on directed edges, but we need to consider paths in $\tilde{G}$ (which involve undirected edges). Before defining Φ we first introduce some notation. Let

$$G' := \left(\vec{E}, \{\{a, b\} : a, b \in \vec{E}, B_{\mathrm{L}}(a, b) > 0\}\right)$$

be the graph supporting the transitions of B_{L}. Denote the collection of all directed paths of length ℓ in G' by $\mathcal{P}_\ell(G')$. For every path $\gamma := (e_0, \ldots, e_\ell)$ in $\tilde{G}$ which is not of the form $e_0 = \cdots = e_\ell$ there is at most one $h(\gamma) := (\vec{e_0}, \ldots, \vec{e_\ell}) \in \mathcal{P}_\ell(G')$ (here and below each f_i is a directed edge, i.e., an element of $\vec{E}$, not a vertex of G) such that

(1) $\{(\vec{e_i})_-, (\vec{e_i})_+\} = e_i$ for all $0 \leq i \leq \ell$,
(2) for all $0 \leq i < \ell$ if $e_i = e_{i+1}$ then $\vec{e_i} = \vec{e_{i+1}}$, and
(3) for all $0 \leq i < \ell$ if $e_i \neq e_{i+1}$ then $(\vec{e_{i+1}})_- = (\vec{e_i})_+$.

Denote the collection of all paths $\gamma = (e_0, \ldots, e_\ell)$ in $\tilde{G}$ not of the form $e_0 = \cdots = e_\ell$, which have such $h(\gamma) = (\vec{e_0}, \ldots, \vec{e_\ell})$ as above by $E^{(\ell)}$. For every path $\gamma = (e_0, \ldots, e_\ell) \in E^{(\ell)}$ let

$$\rho(e_0, \ldots, e_\ell) := \frac{(1-p)^\ell p}{2} \prod_{i \in [\ell]} B_{\mathrm{L}}(\vec{e_{i-1}}, \vec{e_i}), \quad \text{where } h(\gamma) = (\vec{e_0}, \ldots, \vec{e_\ell}).$$

Let $e, e' \in E$. For every $e_0 = e, e_1, \ldots, e_\ell = e' \in E$ with $(e_0, \ldots, e_\ell) \in E^{(\ell)}$ let $\Phi_{e,e'}((e_0, \ldots, e_\ell)) := \rho(e_0, \ldots, e_\ell)$. It is not hard to verify that by construction we have that $\Phi := (\Phi_{e,e'})_{\{e,e'\} \in F}$ is a U-W flow. It remains only to bound its congestion.

Observe that for all $(e_0, \ldots, e_\ell) \in E^{(\ell)}$ and all $1 \leq i < \ell$ we have that

$$\rho(e_0, \ldots, e_\ell) = \frac{2}{p}\rho(e_0, \ldots, e_i)\rho(e_i, \ldots, e_\ell) = \frac{2}{p}\rho(e_i, \ldots, e_1, e_0)\rho(e_i, \ldots, e_\ell), \tag{6.5}$$

if $e_0, e_1, \ldots, e_i$ are not all equal and also $e_i, e_1, \ldots, e_\ell$ are not all equal. To make this true without this restriction, if $e = e_0 = \cdots = e_i$ we define $\rho(e_0, \ldots, e_i) := \frac{(1-p)^i p}{2} 2^{-i}$. If $(e_0, \ldots, e_i) \notin E^{(i)}$ and is not of the form $e = e_0 = \cdots = e_i$, we set $\rho(e_0, \ldots, e_i) := 0$.

Let $\{e, e'\} \in \tilde{E}$. As in Definition 4.2 let $A_{e,e'}(\Phi)$ be the congestion of (e, e') w.r.t. Φ. Finally, there exists a constant $C_p > 0$ such that for every $\{e, e'\} \in \tilde{E}$ we have that

$$
A_{e,e'}(\Phi) \leq \sum_{\substack{(e_0,\ldots,e_\ell) \in E^{(\ell)}, \ell \in \mathbb{N} : (e,e')=(e_{i-1},e_i) \text{ for some } i \in [\ell]}} \ell^2 \rho(e_0, \ldots, e_\ell)
$$

$$
\left(\text{using (6.5) } \ell^2 \rho(e_0, \ldots, e_\ell) \leq \frac{8}{p} i^2 \rho(e_i, \ldots, e_1, e_0)(\ell - i)^2 \rho(e_i, \ldots, e_\ell) \right)
$$

for $i > 0$ such that $(e_{i-1}, e_i) = (e, e')$)

$$
\leq \frac{8}{p} \left(\sum_{\substack{e_0,\ldots,e_\ell \in E, \ell \in \mathbb{N} : (e_0,e_1)=(e,e')}} \ell^2 \rho(e_0, \ldots, e_\ell) \right)^2
$$

$$
\left(\text{writing } q_\ell := \sum_{\substack{e_0,\ldots,e_\ell \in E, : (e_0,e_1)=(e,e')}} \rho(e_0, \ldots, e_\ell) \text{ we have that } \sum_\ell q_\ell = 1/2 \right.
$$

$$
\left. \text{and so } \left(\sum_\ell q_\ell \ell^2 \right)^2 \leq \sum_\ell q_\ell \ell^4 = \sup_{(u,v) \in \overrightarrow{E}} \mathbb{E}_{(u,v)}\left[\tau_1^4 \right] \right)
$$

$$
\leq \frac{8}{p} \sup_{(u,v) \in \overrightarrow{E}} \mathbb{E}_{(u,v)}\left[\tau_1^4 \right] \leq C'_p.
$$

This concludes the proof of (iii).   □

7. Finding a stationary measure for the edge k-NBRW

Let $G = (V, E)$ be a graph. Let $s \geq k$. Recall that $\mathcal{P}_s$ is the collection of paths of length s in G. For $\gamma \in \mathcal{P}_s$ let

$$
m_i^{(k)}(\gamma) := \left| \{ j : i - k < j \leq i \text{ such that } \gamma_j = \gamma_i \} \right|.
$$

For $\gamma \in \mathcal{P}_s$ we denote the multi-set $\{ m_i^{(k)}(\gamma) : i \in [s-1] \text{ such that } \gamma_i = v \}$ by $A(v, k, \gamma)$ (it will be crucial in what comes that in this definition we do not consider $i \in \{0, s\}$).

For $\gamma \in \mathcal{P}_k^e$ let

$$
\pi_k^e(\gamma) := \prod_{i=1}^{k-1} \frac{1}{\deg(\gamma_i) - m_i^{(k)}(\gamma)}. \tag{7.1}
$$

For $s \geq k$ we define

$$
\mathcal{P}_k^e(s) := \left\{ \gamma \in \mathcal{P}_s : (\gamma_i, \gamma_{i+1}, \ldots, \gamma_{i+k}) \in \mathcal{P}_k^e \text{ for all } 0 \leq i \leq s - k \right\}. \tag{7.2}
$$

We say that two multi-sets are equal if each element appears in both with the same multiplicity.

Lemma 7.1. *Let $k \leq s$. Let $\gamma \in \mathcal{P}_k^e(s)$. Then also $\gamma^r \in \mathcal{P}_k^e(s)$. Moreover, $A(v, k, \gamma) = A(v, k, \gamma^r)$ (as multi-sets) for all v. Consequently, for all $\gamma \in \mathcal{P}_k^e$ we have that*

$$
\pi_k^e(\gamma) = \pi_k^e(\gamma^r). \tag{7.3}
$$

Proof. The claim that $\gamma \in \mathcal{P}_k^e(s)$ iff $\gamma^r \in \mathcal{P}_k^e(s)$ is obvious. It is also easy to see from (7.1) that (7.3) is immediate from the claim that $A(v, k, \gamma) = A(v, k, \gamma^r)$ for all v and all $\gamma \in \mathcal{P}_k^e(s)$ (by taking $s = k$). We now sketch the proof of the claim that $A(v, k, \gamma) = A(v, k, \gamma^r)$ for all v. First prove the case $k = s$ by induction on k. Then, for each fixed k prove the case $s > k$ by induction on s (with the base case being $s = k$).

In the first induction, in the induction step from $s = k = i$ to $s = k = i + 1$, compare $A(v, i + 1, \gamma)$ and $A(v, i + 1, \gamma^r)$ with $A(v, i, \alpha)$ and $A(v, i, \alpha^r)$, respectively, where α is obtained from γ by omitting from γ its last co-ordinate. In the second induction, in the induction step from (k, s) to $(k, s + 1)$ compare $A(v, k, \gamma)$ and $A(v, k, \gamma^r)$ with $A(v, k, \alpha)$ and $A(v, k, \alpha^r)$, respectively, where α is obtained from γ by omitting its last co-ordinate. $\square$

Corollary 7.2. *Assume that the edge k-NBRW on G cannot get stuck. Then for all $n \in \mathbb{N}$ and $\alpha, \beta, \alpha_0, \alpha_1, \ldots, \alpha_n \in \mathcal{P}_k^e$ we have that*

$$\pi_k^e(\alpha_0) \prod_{i=0}^{n-1} P_{k,e}(\alpha_i, \alpha_{i+1}) = \pi_k^e(\alpha_n) \prod_{i=0}^{n-1} P_{k,e}(\alpha_{n-i}^r, \alpha_{n-(i+1)}^r).$$

$$\pi_k^e(\alpha) P_{k,e}^n(\alpha, \beta) = \pi_k^e(\beta) P_{k,e}^n(\beta^r, \alpha^r). \tag{7.4}$$

In particular, π_k^e is stationary for $P_{k,e}$.

Proof. The first line of (7.4) is a consequence of Lemma 7.1. Indeed, consider $s = n + k$ and $\gamma := (\gamma_0, \ldots, \gamma_{n+k})$, where $(\gamma_i, \ldots, \gamma_{i+k}) = \alpha_i$ for all $i \in \{0, 1, \ldots, n\}$. Then

$$\pi_k^e(\alpha_0) \prod_{i=0}^{n-1} P_{k,e}(\alpha_i, \alpha_{i+1}) = \prod_{i=1}^{k+n-1} \frac{1}{\deg(\gamma_i) - m_i^{(k)}(\gamma)}$$

$$= \prod_{v \in \gamma} \prod_{m \in A(v,k,\gamma)} \frac{1}{\deg(v) - m} \overset{\text{by Lemma 7.1}}{=} \prod_{v \in \gamma^r} \prod_{m \in A(v,k,\gamma^r)} \frac{1}{\deg(v) - m}$$

$$= \prod_{i=1}^{k+n-1} \frac{1}{\deg(\gamma_i^r) - m_i^{(k)}(\gamma^r)} = \pi_k^e(\alpha_n) \prod_{i=0}^{n-1} P_{k,e}(\alpha_{n-i}^r, \alpha_{n-(i+1)}^r),$$

where each m in the third (respectively, fourth) term is taken with its multiplicity in the multi-set $A(v, k, \gamma)$ (respectively, $A(v, k, \gamma^r)$).

The second line of (7.4) follows from the first by summing over all $\alpha_0, \alpha_1, \ldots, \alpha_n \in \mathcal{P}_k^e$ such that $\alpha_0 = \alpha$ and $\alpha_n = \beta$. Note that here we are using the assumption that the walk cannot get stuck. Finally, $\sum_{\alpha \in \mathcal{P}_k^e} \pi_k^e(\alpha) P_{k,e}(\alpha, \beta) = \pi_k^e(\beta) \sum_{\alpha \in \mathcal{P}_k^e} P_{k,e}(\beta^r, \alpha^r) = \pi_k^e(\beta)$ and so π_k^e is indeed stationary for $P_{k,e}$. $\square$

8. Proof of Theorem 1

We shall need the following simple lemma whose proof is deferred to Section A.3.

Lemma 8.1. *Let P be the transition kernel of an irreducible Markov chain. Let $\hat{P} := \sum_{i=0}^m p_i P^i$ where $\sum_{i=0}^m p_i = 1$ and $m \in \mathbb{N}$. If the greatest common denominator (gcd) of $\{i \in [m] : p_i > 0\}$ is 1 then P is recurrent iff $\hat{P}$ is recurrent.*

8.1. *Proof of Parts (i) and (ii) of Theorem 1*

Proof. The assertion of Part (i) follows from Corollary 5.4 in conjunction with Theorem 4.12. We now prove Part (ii). Let $G = (V, E)$ be a connected graph of bounded degree. Assume that condition (2) holds. Let R be as in condition

(2). Let d be the maximal degree in G. We first argue that by condition (2) for[4] $M := 4R + 1$ and $p := (d - 1)^{-M}/[2(M + 1)]$ for every $a \in \vec{E}$ we have that

$$\sum_{i=1}^{M} B^i(a, a^{\mathrm{r}}) \geq 2p(M + 1). \tag{8.1}$$

Indeed, by condition (2) for all $a \in \vec{E}$ there is some $i = i(a) \leq M$ such that for some $a_0 = a, a_1, \ldots, a_i = a^{\mathrm{r}}$ we have that $B(a_{j-1}, a_j) > 0$ for all $j \in [i]$. Then $\prod_{j \in [i]} B(a_{j-1}, a_j) \geq (d - 1)^{-M} = 2p(M + 1)$.

Let $D := \frac{1}{M+1} \sum_{i=0}^{M} B^i$ and $D_{\mathrm{L}} := \frac{1}{2}(I + D)$. Let $\mathbf{X} := (X_n)_{n=0}^{\infty}$ be the Markov chain corresponding to D_{L}. By Lemma 8.1 we have that the NBRW is transient iff $\mathbf{X}$ is transient.

By (8.1) we may generate $\mathbf{X}$ as follows: given $X_n = a$ the chain first picks a candidate Z_{n+1} for X_{n+1} according to D (i.e., $Z_{n+1} = b$ with probability $D(a, b)$) and then flips a fair coin in order to decide if the candidate is accepted (i.e., with probability $1/2$ we set $X_{n+1} = Z_{n+1}$ and otherwise we set $X_{n+1} = X_n$). Furthermore, we may generate $\mathbf{X}$ and the Z_n's such that for some $\xi_0, \xi_1, \ldots$ i.i.d. Bernoulli(p) random variables for all $n \geq 0$ we have that $Z_{n+1} = X_n^{\mathrm{r}}$ whenever $\xi_n = 1$ (and possibly also if $\xi_n = 0$). Indeed, the probability that $Z_{n+1} = X_n^{\mathrm{r}}$, given X_n is at least p. Given X_n we may generate Z_{n+1} and ξ_n simultaneously using $U_n \sim U[0, 1]$ (where $U[0, 1]$ denotes the uniform distribution on $[0, 1]$ and $U_1, U_2, \ldots$ are i.i.d. $U[0, 1]$) as follows: Z_{n+1} may taken different values $a_1 := X_n^{\mathrm{r}}, a_2, \ldots, a_i$ with probabilities $q_1, \ldots, q_i$ (where the a_j's and the q_j's depend on X_n). Set $\xi_n = \mathbf{1}\{U_n \leq p\}$. Set $Z_{n+1} = a_i$ if $U_n \in [\sum_{j=1}^{i-1} q_j, \sum_{j=1}^{i} q_j)$ (where $\sum_{j=1}^{0} q_j$ is defined to be 0). Then since regardless of the value of X_n we have that $q_1 \geq p$ we get that $Z_{n+1} = a_1$ whenever $\xi_n = 1$.

Fix some $a = (x_0, y_0) \in \vec{E}$. Consider the case that the initial distribution of $\mathbf{X}$ is the uniform distribution on $\{a, a^{\mathrm{r}}\}$. Let $e_0 = \{x_0, y_0\}$ and $\tau_0 = -1$. We define inductively $\tau_{n+1} := \inf\{t > \tau_n : \xi_{t-1} = 1\}$ and $e_n := \{x_n, y_n\}$, where x_n and y_n are the end-points of Z_{τ_n+1}. Consider the process $\mathbf{Q} := (e_n)_{n=0}^{\infty}$ on E given by $Q_n = e_n$ for all $n \geq 0$. To conclude the proof it suffices to show that

(i) The process $\mathbf{Q}$ is a Markov chain. Its transition kernel W is symmetric. Hence $\mathbf{Q}$ is reversible w.r.t. the counting measure π_E on E.

(ii) The NBRW is recurrent iff $\mathbf{Q}$ is recurrent.

(iii) SRW on G is recurrent iff $\mathbf{Q}$ is recurrent.

We first prove part (i). Consider the transition kernel K on $\vec{E}$ given by

$$K(a, b) = \begin{cases} (1 - p)^{-1} D(a, b) & \text{if } b \neq a^{\mathrm{r}}, \\ (1 - p)^{-1}(D(a, a^{\mathrm{r}}) - p) & \text{if } b = a^{\mathrm{r}}. \end{cases}$$

It follows from (6.2) that $K(a, b) = K(b^{\mathrm{r}}, a^{\mathrm{r}})$ for all $a, b \in \vec{E}$. Hence $K_{\mathrm{L}} := \frac{1}{2}(I + K)$ also satisfies this identity. Finally, the symmetry of W follows from the last identity since for every $a, b \in E$ we have

$$W(a, b) := \frac{1}{2} \sum_{i=0}^{\infty} \sum_{w \in D_a, w' \in D_b} K_{\mathrm{L}}^i(w, w')(1 - p)^i p,$$

where $D_{\{x, y\}} := \{(x, y), (y, x)\}$. The proofs of (ii)–(iii) above are almost identical to those of (ii)–(iii) from the proof of Theorem 3 (and also to that of (iii) in the proof of Part (iii) of Theorem 1 below) and are thus omitted. □

---

[4]We have not attempted to pick the smallest possible M.

8.2. *Proof of Part* (iii) *of Theorem* 1

Proof. Assume that conditions (2_k^e) and (3_k^e) hold. Then there exist some $M \in \mathbb{N}$ and $p \in (0, 1)$ such that for every $a \in \mathcal{P}_k^e$ we have that

$$\sum_{i=1}^{M} P_{k,e}^i(a, a^r) \geq 2p(M + 1). \tag{8.2}$$

Let $D := \frac{1}{M+1} \sum_{i=0}^{M} P_{k,e}^i$ and $D_L := \frac{1}{2}(I + D)$. Let $\mathbf{X} := (X_n)_{n=0}^{\infty}$ be the Markov chain corresponding to D_L. By Lemma 8.1 the edge k-NBRW is transient iff $\mathbf{X}$ is transient. By (8.2) we may generate $\mathbf{X}$ as follows: given $X_n = a$ the chain first picks a candidate Z_{n+1} for X_{n+1} according to D and then flips a fair coin in order to decide if the candidate is accepted. Furthermore, one may generate $Z_1, Z_2, \ldots$ such that for some $\xi_0, \xi_1, \ldots$ i.i.d. Bernoulli(p) random variables for all $n \geq 0$ we have that $Z_{n+1} = X_n^r$ whenever $\xi_n = 1$ (and possibly also if $\xi_n = 0$). The proof of this claim is identical to the corresponding claim in the proof of Part (ii) of Theorem 1.

Fix some $a = \mathcal{P}_k^e$. Consider the case that the initial distribution of $\mathbf{X}$ is the uniform distribution on $\{a, a^r\}$. Let $e_0 = \{a, a^r\}$ and $\tau_0 = -1$. We define inductively $\tau_{n+1} := \inf\{t > \tau_n : \xi_{t-1} = 1\}$ and $e_n := \{Z_{\tau_n+1}, Z_{\tau_n+1}^r\}$. Consider the process $\mathbf{Q} := (Q_n)_{n=0}^{\infty}$ on $U_k^e := \{\{a, a^r\} : a \in \mathcal{P}_k^e\}$ given by $Q_n := e_n$ for all $n \geq 0$. To conclude the proof it suffices to show that

(i) The process $\mathbf{Q}$ is a Markov chain. Its transition kernel W satisfies that for all $a, b \in \mathcal{P}_k^e$

$$\pi_k^e(a) W(\{a, a^r\}, \{b, b^r\}) = \pi_k^e(b) W(\{b, b^r\}, \{a, a^r\}),$$

where π_k^e is as in (7.1). Hence $\mathbf{Q}$ is reversible w.r.t. the measure $\pi(\{a, a^r\}) := \pi_k^e(a)$.
(ii) The edge k-NBRW is recurrent iff $\mathbf{Q}$ is recurrent.
(iii) SRW on G is recurrent iff $\mathbf{Q}$ is recurrent.

We first prove part (i). Consider the transition kernel K on $\mathcal{P}_k^e$ given by

$$K(a, b) = \begin{cases} (1-p)^{-1} D(a, b) & \text{if } b \neq a^r, \\ (1-p)^{-1}(D(a, a^r) - p) & \text{if } b = a^r. \end{cases}$$

It follows from (7.4) that K is satisfies $\pi_k^e(a) K(a, b) = \pi_k^e(b) K(b^r, a^r)$ for all $a, b \in \mathcal{P}_k^e$. Hence also $K_L := \frac{1}{2}(I + K)$ satisfies $\pi_k^e(a) K_L(a, b) = \pi_k^e(b) K_L(b^r, a^r)$ for all $a, b \in \mathcal{P}_k^e$. The assertion of (i) now follows from the fact that for every $a, b \in \mathcal{P}_k^e$ we have that

$$W(\{a, a^r\}, \{b, b^r\}) := \frac{1}{2} \sum_{i=0}^{\infty} \sum_{w \in \{a,a^r\}, w' \in \{b,b^r\}} K_L^i(w, w')(1-p)^i p.$$

The proof of (ii) is analogous of that of claim (ii) from the proof of Theorem 3, and thus omitted.

We now prove (iii). Recall that $H_k^e := (U_k^e, F_k^e)$ is the graph obtained from G_k^e (from Definition 5.5) by identifying each $a \in \mathcal{P}_k^e$ with its reversal. More precisely (as above) $U_k^e := \{\{a, a^r\} : a \in \mathcal{P}_k^e\}$ and $\{\{a, a^r\}, \{b, b^r\}\} \in F_k^e$ iff $\sum_{w \in \{a,a^r\}, w' \in \{b,b^r\}} P_{e,k}(w, w') > 0$. Let U be the transition kernel of SRW on $\tilde{H} := (U_k^e, \tilde{F})$, where

$$\tilde{F} := \left\{ \{\{a, a^r\}, \{b, b^r\}\} : a, b \in \mathcal{P}_k^e \text{ such that } \sum_{w \in \{a,a^r\}, w' \in \{b,b^r\}} D_L(w, w') > 0 \right\}.$$

By Lemma 5.6 the SRW on G is recurrent iff the SRW on H_k^e is recurrent. Since H_k^e and $\tilde{H}$ are roughly-isometric, to conclude the proof of Claim (iii) we now show that $\mathbf{Q}$ is recurrent iff the SRW on $\tilde{H}$ is recurrent. By Lemmas 4.1 and 4.3 it suffices to show that there exist a U-W flow and a W-U flow of finite congestion. Similarly to the proof of Claim (iii) in the proof of Theorem 3, the W-U flow obtained by mapping every (e, e') such that $\{e, e'\} \in \tilde{F}$ to itself (with weight 1) is of finite congestion.

We now construct a U-W flow Φ of finite congestion. Loosely speaking, the flow Φ assigns to every path the probability that the chain $\mathbf{X}$ follows that path between two consecutive steps of $\mathbf{Q}$. This is not a precise description because $\mathbf{X}$ is defined on $\mathcal{P}_k^{\mathrm{e}}$ rather than on U_k^{e}. Before defining Φ we first introduce some notation. Let $\tilde{G} = (\mathcal{P}_k^{\mathrm{e}}, \tilde{E})$ be the graph supporting the transitions of D_{L} (i.e. $\{a, b\} \in \tilde{E}$ iff $D_{\mathrm{L}}(a, b) > 0$). Denote the collection of all directed paths of length ℓ in $\tilde{G}$ by $\mathcal{P}_{\ell}(\tilde{G})$. For every path $(a_0, \ldots, a_{\ell}) \in \mathcal{P}_{\ell}(\tilde{G})$ let

$$\sigma(a_0, \ldots, a_{\ell}) := \prod_{i \in [\ell]} K_{\mathrm{L}}^i(a_{i-1}, a_i).$$

For every $e_0, \ldots, e_{\ell} \in U_k^{\mathrm{e}}$ let

$$J(e_0, \ldots, e_{\ell}) := \left\{ (a_0, \ldots, a_{\ell}) \in \mathcal{P}_{\ell}(\tilde{G}) : \{a_i, a_i^{\mathrm{r}}\} = e_i \text{ for all } 0 \le i \le \ell \right\},$$

$$\rho(e_0, \ldots, e_{\ell}) = \frac{\pi(e_0)}{2} \sum_{(a_0, \ldots, a_{\ell}) \in J(e_0, \ldots, e_{\ell})} \sigma(a_0, \ldots, a_{\ell})(1 - p)^{\ell} p.$$

As oppose to the situation in the proof of Theorem 3, if $J(e_0, \ldots, e_{\ell})$ is not empty, then it may contain more than one path even if it is not the case that $e_0 = \cdots = e_{\ell}$. If $J(e_0, \ldots, e_{\ell}) = \varnothing$ then we set $\rho(e_0, \ldots, e_{\ell}) = 0$.

Let $e, e' \in U_k^{\mathrm{e}}$ be such that $W(e, e') > 0$. For every $e_0 = e, e_1, \ldots, e_{\ell} = e' \in U_k^{\mathrm{e}}$ let $\Phi_{e,e'}((e_0, \ldots, e_{\ell})) := \rho(e_0, \ldots, e_{\ell})$. It is not hard to verify that by construction we have that $\Phi := (\Phi_{e,e'})_{e,e':W(e,e')>0}$ is a W-U flow. It remains only to bound its congestion.

Observe that for all every $e_0, \ldots, e_{\ell} \in U_k^{\mathrm{e}}$ and $0 \le i \le \ell$ we have that[5]

$$\rho(e_0, \ldots, e_{\ell}) \le \frac{2}{p\pi(e_i)} \rho(e_0, \ldots, e_i)\rho(e_i, \ldots, e_{\ell}) \le C_0 \rho(e_i, \ldots, e_1, e_0)\rho(e_i, \ldots, e_{\ell}).$$

The remainder of the proof is analogous to the remainder of the proof of Theorem 3 after (6.5). $\square$

8.3. *Proof of Part* (iv) *of Theorem* 1

Proof. Let $G = (V, E)$ be a vertex-transitive graph satisfying conditions (2_k^{v}) and (3_k^{v}). Assume that the SRW on G is transient. Assume towards a contradiction that the vertex k-NBRW is recurrent. Since G is transient, by Part (3) of Lemma 5.6 the SRW on G_k^{v} is also transient. We shall show that $\sup_{\alpha, \beta \in \mathcal{P}_k^{\mathrm{v}}} \frac{\pi(\alpha)}{\pi(\beta)} < \infty$. Using Theorem 4.12 and Fact 4.9 we get that the vertex k-NBRW is transient, in contradiction to our assumption.

By conditions (2_k^{v}) and (3_k^{v}) the vertex k-NBRW is irreducible. By recurrence, its stationary measure π is unique up to a multiplication by a constant factor. Moreover, by irreducibility and recurrence, for all $\alpha, \beta \in \mathcal{P}_k^{\mathrm{v}}$ we have that [6]

$$\frac{\pi(\alpha)}{\pi(\beta)} = \lim_{t \to \infty} \frac{\sum_{i=0}^{t} P_{k,\mathrm{v}}^i(\alpha, \alpha)}{\sum_{i=0}^{t} P_{k,\mathrm{v}}^i(\beta, \beta)}.$$

Every automorphism ϕ of G can be extended into a probability preserving bijection of $\mathcal{P}_k^{\mathrm{v}}$, by defining $\phi(\alpha) := (\phi(\alpha_0), \ldots, \phi(\alpha_k))$ for $\alpha = (\alpha_0, \ldots, \alpha_k) \in \mathcal{P}_k^{\mathrm{v}}$. It is 'probability preserving' in the sense that $P_{k,\mathrm{v}}^i(\alpha, \beta) = P_{k,\mathrm{v}}^i(\phi(\alpha), \phi(\beta))$ for all $\alpha, \beta \in \mathcal{P}_k^{\mathrm{v}}$ and $i \in \mathbb{N}$. In particular,

$$\frac{\pi(\alpha)}{\pi(\phi(\alpha))} = \lim_{t \to \infty} \frac{\sum_{i=0}^{t} P_{k,\mathrm{v}}^i(\alpha, \alpha)}{\sum_{i=0}^{t} P_{k,\mathrm{v}}^i(\beta, \beta)} = 1. \tag{8.3}$$

---

[5]As opposed to the situation in the proof of Theorem 3, the first inequality is no longer an equality. E.g., the left most term can be 0 while the middle term is positive. Also, it is possible that middle term may contain contributions coming from opposite orientations e_i.

By conditions (2_k^v) and (3_k^v) there exist some $M \in \mathbb{N}$ and $p > 0$ such that if $\alpha, \beta \in \mathcal{P}_k^v$ and $\alpha_0 = \beta_0$ then $\sum_{i=1}^{M} P_{k,v}^i(\alpha, \beta) > p$ and consequently for $c = p/M$ we have that

$$\pi(\beta) = \frac{1}{M} \sum_{i=1}^{M} \sum_{\gamma \in \mathcal{P}_k^v} \pi(\gamma) P_{k,v}^i(\gamma, \beta) \geq \frac{1}{M} \sum_{i=1}^{M} \pi(\alpha) P_{k,v}^i(\alpha, \beta) \geq c\pi(\alpha). \tag{8.4}$$

Finally, by (8.3) and the fact that G is vertex-transitive, $\sup_{\alpha, \beta \in \mathcal{P}_k^v} \frac{\pi(\alpha)}{\pi(\beta)} = \sup_{\alpha, \beta \in \mathcal{P}_k^v, \alpha_0 = \beta_0} \frac{\pi(\alpha)}{\pi(\beta)}$, and so by (8.4) $\sup_{\alpha, \beta \in \mathcal{P}_k^v} \frac{\pi(\alpha)}{\pi(\beta)} \leq 1/c$, as desired. $\qquad \square$

9. Proof of Theorem 2

Let $G = H(S)$ be as in Theorem 2. As H is Abelian, we denote the group's operation by '+' and denote the inverse of $h \in H$ by $-h$. For $h \in H$ and $k \in \mathbb{N}$ we define kh inductively to be $h + (k-1)h$. We denote the identity element by 0. Recall the notion of an induced chain from Section 2. To exploit the symmetry of G we consider the following auxiliary chain.

Definition 9.1. Let $G = H(S)$ be the Cayley graph of an infinite finitely generated Abelian group $H \neq \mathbb{Z}$ w.r.t. a finite symmetric set of generators S. Fix some $s \in S$ of infinite order (there must be such $s \in S$). Let $k \in \mathbb{N}$. For every $h \in H$ let

$$\alpha_h := (h, h + s, \ldots, h + ks).$$

Let $\mathcal{A} := \{\alpha_h, \alpha_h^r : h \in H\}$. Let $\pi_{\mathcal{A}} \equiv 1$ be the counting measure on $\mathcal{A}$. Let $\mathbf{Y}^e := (Y_t^e)_{t=0}^{\infty}$ (respectively, $\mathbf{Y}^v := (Y_t^v)_{t=0}^{\infty}$) be the induced chain on $\mathcal{A}$ w.r.t. the edge (respectively, vertex) k-NBRW. Denote its transition matrix by $P_{\mathcal{A},e}$ (respectively, $P_{\mathcal{A},v}$).

Lemma 9.2. *In the setup of Definition 9.1 we have that $\pi_{\mathcal{A}}$ is stationary for both $\mathbf{Y}^e$ and $\mathbf{Y}^v$. Moreover, for all $a, b \in \mathcal{A}$ we have that*

$$P_{\mathcal{A},e}(a, b) = P_{\mathcal{A},e}(b^r, a^r) \quad \text{and} \quad P_{\mathcal{A},v}(a, b) = P_{\mathcal{A},v}(b^r, a^r). \tag{9.1}$$

We defer the proof to Section A.4. The main reason for which the assertion of the lemma holds is that Cayley graphs of finitely generated Abelian groups satisfy that for every pair of vertices x, y there is an automorphism that maps x to y and y to x. Namely, first translate by $-x$, then map every a to $-a$ and finally translate by $+y$. A more precise explanation is that Lemma 9.2 is a consequence of symmetry under negation, along with translation invariance and the fact that $\mathcal{A}$ is closed under negation, i.e., $-\mathcal{A} := \{-\gamma : \gamma \in \mathcal{A}\} = \mathcal{A}$, where for $\gamma = (\gamma_0, \ldots, \gamma_k)$ we define $-\gamma := (-\gamma_0, \ldots, -\gamma_k)$. By 'symmetry under negation' (respectively, 'translation invariance') we mean that for the k-NBRW the transition probability from γ to $\tilde{\gamma}$ equals its transition probability from $-\gamma$ to $-\tilde{\gamma}$ (respectively, from $\gamma + h := (\gamma_0 + h, \ldots, \gamma_k + h)$ to $\tilde{\gamma} + h$, for all $h \in H$).

Proof of Theorem 2. We only prove the equivalence between the SRW and the vertex k-NBRW as the analysis of the edge k-NBRW is analogous. Using the notation from Definition 9.1 there exists some $p \in (0, 1)$ such that for every $a \in \mathcal{A}$ we have that

$$P_{\mathcal{A},v}(a, a^r) = p \tag{9.2}$$

(the constant p is independent of a by the aforementioned translation invariance and symmetry under negation) Let $\mathbf{X} := (X_n)_{n=0}^{\infty}$ be the Markov chain corresponding to the transition kernel $P' := \frac{1}{2}(I + P_{\mathcal{A},v})$. Since $\mathcal{A}$ is a.s. visited by the vertex k-NBRW we have that the vertex k-NBRW is transient iff $P_{\mathcal{A},v}$ is transient, which occurs iff $\mathbf{X}$ is transient.

By (9.2) we may generate $\mathbf{X}$ as follows: given $X_n = a$ the chain first picks a candidate Z_{n+1} for X_{n+1} according to P' and then flips a fair coin in order to decide if the candidate is accepted. Furthermore, one may generate $Z_1, Z_2, \ldots$ such that for some $\xi_0, \xi_1, \ldots$ i.i.d. Bernoulli(p) random variables, for all $n \geq 0$ we have that $Z_{n+1} = X_n^{\mathrm{r}}$ iff $\xi_n = 1$.

Fix some $a \in \mathcal{A}$. Consider the case that the initial distribution of $\mathbf{X}$ is the uniform distribution on $\{a, a^{\mathrm{r}}\}$. Let $e_0 = \{a, a^{\mathrm{r}}\}$ and $\tau_0 = -1$. We define inductively $\tau_{n+1} := \inf\{t > \tau_n : \xi_{t-1} = 1\}$ and $e_n := \{Z_{\tau_n+1}, Z_{\tau_n+1}^{\mathrm{r}}\}$. Consider the process $\mathbf{Q} := (e_n)_{n=0}^{\infty}$ on $\mathcal{B} := \{\{a, a^{\mathrm{r}}\} : a \in \mathcal{A}\}$. To conclude the proof it suffices to show that

(i) The process $\mathbf{Q}$ is a Markov chain. Its transition kernel W is symmetric. Hence $\mathbf{Q}$ is reversible w.r.t. the counting measure on $\mathcal{B}$.

(ii) The vertex k-NBRW is recurrent iff $\mathbf{Q}$ is recurrent.

(iii) SRW on G is recurrent iff $\mathbf{Q}$ is recurrent.

We first prove part (i). Consider the transition kernel K on $\mathcal{A}$ given by

$$K(a, b) := \begin{cases} (1-p)^{-1} P_{\mathcal{A}, \mathrm{v}}(a, b) & \text{if } b \neq a^{\mathrm{r}}, \\ (1-p)^{-1}(P_{\mathcal{A}, \mathrm{v}}(a, a^{\mathrm{r}}) - p) & \text{if } b = a^{\mathrm{r}}. \end{cases}$$

It follows from (9.1) that $K(a, b) = K(b^{\mathrm{r}}, a^{\mathrm{r}})$ for all $a, b \in \mathcal{A}$. Hence $K_{\mathrm{L}} := \frac{1}{2}(I + K)$ satisfies this identity as well. The assertion of (i) now follows from the fact that for every $a, b \in \mathcal{A}$ we have that

$$W(\{a, a^{\mathrm{r}}\}, \{b, b^{\mathrm{r}}\}) := \frac{1}{2} \sum_{i=0}^{\infty} \sum_{w \in \{a, a^{\mathrm{r}}\}, w' \in \{b, b^{\mathrm{r}}\}} K_{\mathrm{L}}^i(w, w')(1-p)^i p.$$

The proofs of (ii) is analogous to that of claim (ii) from the proof of Theorem 3 and is thus omitted. We now prove (iii). Let U be the transition matrix of SRW on G. Observe that we may identify $\mathcal{B}$ with H via the map $\{\alpha_h, \alpha_h^{\mathrm{r}}\} \to h$. Hence it suffices to construct a U-W flow and a W-U flow of finite congestion. The details are similar to those from the proof of Part (iii) of Theorem 1 and are thus omitted. □

Appendix: Technical proofs

A.1. *Proof of Proposition* 1.1

Proposition. *Let $G = (V, E)$ be a d-regular connected graph for some $d \geq 3$. Then the SRW on G is recurrent iff the NBRW on G is recurrent.*

Proof. Fix some $o \in V$. Let T be the universal cover of G. Namely, $T = (V', E')$ is an infinite d-ary tree rooted at o, whose ℓth level is labelled by the collection of all non-backtracking paths in G of length ℓ started from o. The children of a vertex $\gamma = (o, v_1, \ldots, v_\ell)$ in T are precisely the vertices which are labelled by the possible extensions of γ to a non-backtracking path of length $\ell + 1$ (i.e. all paths of the form $(o, v_1, \ldots, v_{\ell+1})$ with $v_{\ell+1} \neq v_{\ell-1}$ and $\{v_\ell, v_{\ell+1}\} \in E$). For every $\gamma = (o, v_1, \ldots, v_\ell) \in V'$ let $\psi(\gamma) = v_\ell$.

Let $\mathbf{Y} := (Y_i)_{i=0}^{\infty}$ (respectively, $\mathbf{W} := (W_i)_{i=0}^{\infty}$) be a SRW (respectively, NBRW) on T started from o. Observe that $(\psi(Y_i))_{i=0}^{\infty}$ (respectively, $(\psi(W_i))_{i=0}^{\infty}$) is a SRW (respectively, NBRW) on G started from o. Clearly, the SRW (respectively, NBRW) on G is recurrent iff a.s. $|\{i : \psi(Y_i) = o\}| = \infty$ (respectively, $|\{i : \psi(W_i) = o\}| = \infty$). We may generate $\mathbf{W}$ from $\mathbf{Y}$ simply as the ray along which $\mathbf{Y}$ diverges (that is, W_i is the last vertex in the ith level of T to be visited by $\mathbf{Y}$). As $\{i : \psi(W_i) = o\} \subseteq \{i : \psi(Y_i) = o\}$, we get that if the NBRW on G is recurrent then so is the SRW.

Conversely, assume that the SRW on G is recurrent. We will show that every time at which the SRW on T reaches some $v \in V'$ with $\psi(v) = o$, with probability (uniformly) bounded from below v belongs to the ray through which the walk diverges.

Denote the ith level of T by $\mathcal{L}_i$. Let $\sigma_0 = 0 = \rho_0$ and define inductively for all $i \geq 0$,

$$\tau_i := \sup\{t : Y_t \in \mathcal{L}_{\rho_i}\}, \qquad \sigma_{i+1} := \inf\{t > \tau_i : \psi(Y_t) = o\} \quad \text{and} \quad \rho_{i+1} := \mathrm{dist}_T(Y_{\sigma_{i+1}}, o).$$

It is not hard to see that a.s. σ_i is finite for all i (by the condition that the SRW is recurrent, and the fact that every level of T is a.s. visited only finitely many times) and that for all i, conditioned on $((\tau_j, \sigma_{j+1}, \rho_{j+1}))_{j=0}^{i-1}$, the probability that $\tau_i = \sigma_i$ is at least $\frac{d-2}{d-1}$. Thus a.s. there are infinitely many i's such that $\tau_i = \sigma_i$. Note that for such i's we have that $\psi(Y_{\sigma_i}) = o$ and $Y_{\sigma_i} \in \{\psi(W_n) : n \in \mathbb{Z}_+\}$.

A.2. Existence of a sink-free and source-free orientation – proof of Lemma 5.2

Lemma. *Let $G = (V, E)$ be an infinite locally finite connected graph of minimal degree at least 2. There exists a sink-free and source-free orientation of E.*

Proof. Label V by $\mathbb{N}$. We start by picking an infinite bi-infinite path $\gamma^{(1)} = (\gamma_i^{(1)} : i \in \mathbb{Z})$ containing the vertex k_1 with label 1 and pick the orientation on it so that $\{\gamma_i^{(1)}, \gamma_{i+1}^{(1)}\}$ gets the orientation $(\gamma_i^{(1)}, \gamma_{i+1}^{(1)})$ for all i in $\mathbb{Z}$. Let k_2 be the vertex with the minimal label which is not part of $\gamma^{(1)}$. If there exists a finite path $\gamma^{(2)} = (\gamma_0^{(2)}, \ldots, \gamma_\ell^{(2)})$ with $k_2 \in \{\gamma_j^{(2)} : j \in [\ell-1]\} \subseteq V \setminus \{\gamma_i^{(1)} : i \in \mathbb{Z}\}$ and $\gamma_0^{(2)}, \gamma_\ell^{(2)}$ in $\gamma^{(1)}$ then pick one such path and pick the orientation on it so that $\{\gamma_{i-1}^{(2)}, \gamma_i^{(2)}\}$ gets the orientation $(\gamma_{i-1}^{(2)}, \gamma_i^{(2)})$ for all $i \in [\ell]$. If no such path exists, then there exists an infinite path $\gamma^{(2)} = (\gamma_0^{(2)}, \gamma_1^{(2)}, \ldots)$ with $k_2 \in \{\gamma_j^{(2)} : j \in \mathbb{N}\} \subseteq V \setminus \{\gamma_i^{(1)} : i \in \mathbb{Z}\}$ and $\gamma_0^{(2)}$ in γ. Pick one such path and pick the orientation on it so that $\{\gamma_{i-1}^{(2)}, \gamma_i^{(2)}\}$ gets the orientation $(\gamma_{i-1}^{(2)}, \gamma_i^{(2)})$ for all $i \in \mathbb{N}$.

Set A_m to be the collection of all vertices belonging to the union of the first m paths of the construction. If $A_m = V$ set $A_i = A_m$ for all $i \geq m$. Otherwise, at stage $m + 1$ let k_{m+1} be the vertex with the minimal label which is not part of A_m. If there exists a finite path $\gamma^{(m+1)} = (\gamma_0^{(m+1)}, \ldots, \gamma_\ell^{(m+1)})$ with $k_{m+1} \in \{\gamma_j^{(m+1)} : j \in [\ell-1]\} \subseteq V \setminus A_m$ and $\gamma_0^{(m+1)}, \gamma_\ell^{(m+1)}$ in A_m then pick one such path and pick the orientation on it so that $\{\gamma_{i-1}^{(m+1)}, \gamma_i^{(m+1)}\}$ gets the orientation $(\gamma_{i-1}^{(m+1)}, \gamma_i^{(m+1)})$ for all $i \in [\ell]$. If no such path exists, then there exists an infinite path $\gamma^{(m+1)} = (\gamma_0^{(m+1)}, \gamma_1^{(m+1)}, \ldots)$ with $k_{m+1} \in \{\gamma_j^{(m+1)} : j \in \mathbb{N}\} \subseteq V \setminus A_m$ and $\gamma_0^{(m+1)}$ in A_m. Pick one such path and pick the orientation on it so that $\{\gamma_{i-1}^{(m+1)}, \gamma_i^{(m+1)}\}$ gets the orientation $(\gamma_{i-1}^{(m+1)}, \gamma_i^{(m+1)})$ for all $i \in \mathbb{N}$. Note that $A_\infty := \bigcup_{i \in \mathbb{N}} A_i = V$ and that by construction every vertex has at least one edge oriented towards it and one away from it, as can be seen by considering the stage m in which the vertex first belonged to A_m, noting that the vertex must have been an internal vertex of the path $\gamma^{(m)}$. Finally, the construction is concluded by picking an arbitrary orientation for the remaining edges.

A.3. Proof of Lemma 8.1

Lemma. *Let P be the transition kernel of an irreducible Markov chain. Let $\hat{P} := \sum_{i=0}^m p_i P^i$ where $\sum_{i=0}^m p_i = 1$, $m \in \mathbb{N}$. If the greatest common denominator (gcd) of $\{i \in [m] : p_i > 0\}$ is 1 then P is recurrent iff $\hat{P}$ is recurrent.*

Proof. Clearly $\hat{P}$ is also irreducible. Let $\hat{P}^i$ be $(\hat{P})^i$. Let x be an arbitrary state. Clearly, $\sum_{i \geq 0} \hat{P}^i(x, x) \leq (1 - p_0)^{-1} \sum_{i \geq 0} P^i(x, x)$. Conversely, consider the random walk $(S_t)_{t \in \mathbb{Z}_+}$ on $\mathbb{Z}_+$ started from 0 whose transition kernel is given by $K(i, i + j) = p_j$. Let q_ℓ be the probability that $S_t = \ell$ for some t. It is standard that the assumption that the gcd of $\{i \in [m] : p_i > 0\}$ is 1 implies that there is some $\delta > 0$ and $N_0 \in \mathbb{N}$ such that for all $\ell \geq N_0$ we have that $q_\ell \geq \delta$ (we shall use the lemma only in the case when $p_i > 0$ for all $i \in [m]$, for which the proof of the last claim is particularly simple). Finally, it is easy to see that $\sum_{i \geq 0} \hat{P}^i(x, x) \geq \sum_{i \geq 0} q_i P^i(x, x) \geq \delta \sum_{i \geq N_0} P^i(x, x)$. Thus $\sum_{i \geq 0} \hat{P}^i(x, x)$ diverges iff $\sum_{i \geq 0} P^i(x, x)$ diverges. $\qquad\square$

A.4. Proof of Lemma 9.2

Lemma. *In the setup of Definition 9.1 we have that $\pi_{\mathcal{A}}$ is stationary for both $\mathbf{Y}^e$ and $\mathbf{Y}^v$. Moreover, for all $a, b \in \mathcal{A}$ we have that*

$$P_{\mathcal{A},e}(a, b) = P_{\mathcal{A},e}(b^r, a^r) \quad and \quad P_{\mathcal{A},v}(a, b) = P_{\mathcal{A},v}(b^r, a^r).$$

Proof. Let $-(g_1, \ldots, g_k) := (-g_1, \ldots, -g_k)$. Observe that $-\alpha_h = -\alpha^{\mathrm{r}}_{-h-sk}$ and that $-\alpha^{\mathrm{r}}_h = \alpha_{-h-sk}$. By translation $P_{\mathcal{A},\mathrm{e}}(\alpha_a, \alpha_b) = P_{\mathcal{A},\mathrm{e}}(\alpha_{a+c}, \alpha_{b+c})$, $P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_a, \alpha^{\mathrm{r}}_b) = P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_{a+c}, \alpha^{\mathrm{r}}_{b+c})$, $P_{\mathcal{A},\mathrm{e}}(\alpha_a, \alpha^{\mathrm{r}}_b) = P_{\mathcal{A},\mathrm{e}}(\alpha_{a+c}, \alpha^{\mathrm{r}}_{b+c})$ and $P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_a, \alpha_b) = P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_{a+c}, \alpha_{b+c})$, for all $a, b, c \in H$. The same holds w.r.t. $P_{\mathcal{A},\mathrm{v}}$. Also, $P_{\mathcal{A},\mathrm{e}}(\alpha_x, \alpha_y) = P_{\mathcal{A},\mathrm{e}}(-\alpha_x, -\alpha_y)$ and $P_{\mathcal{A},\mathrm{v}}(\alpha_x, \alpha_y) = P_{\mathcal{A},\mathrm{v}}(-\alpha_x, -\alpha_y)$. To see this, consider an arbitrary path $(\alpha_x = a_0, a_1, \ldots, a_r = \alpha_y)$ with $a_i \in \Omega^{\mathrm{e}}_k$ (respectively, $\in \Omega^{\mathrm{v}}_k$) for all i. Then $(-\alpha_x = -a_0, -a_1, \ldots, -a_r = -\alpha_y)$ is a path from $-\alpha_x$ to $-\alpha_y$. Moreover, the probability the edge (respectively, vertex) k-NBRW started from α_x follows the first path equals the probability the edge (respectively, vertex) k-NBRW started from $-\alpha_x$ follows the second path. Indeed, this follows from the fact that $x \to -x$ is an automorphism of G which extends to a probability preserving bijection of Ω^{e}_k (respectively, Ω^{v}_k) in the sense that $P_{k,\mathrm{e}}(a, b) = P_{k,\mathrm{e}}(-a, -b)$ for all $a, b \in \Omega^{\mathrm{e}}_k$ (respectively, $P_{k,\mathrm{v}}(a, b) = P_{k,\mathrm{v}}(-a, -b)$ for all $a, b \in \Omega^{\mathrm{v}}_k$).

Hence

$$P_{\mathcal{A},\mathrm{e}}(\alpha_a, \alpha_b) = P_{\mathcal{A},\mathrm{e}}(\alpha_{-ks}, \alpha_{b-a-ks}) = P_{\mathcal{A},\mathrm{e}}(-\alpha_{-ks}, -\alpha_{b-a-ks})$$
$$= P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_0, \alpha^{\mathrm{r}}_{a-b}) = P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_b, \alpha^{\mathrm{r}}_a)$$

(where 0 is the identity element of H) and similarly $P_{\mathcal{A},\mathrm{v}}(\alpha_a, \alpha_b) = P_{\mathcal{A},\mathrm{v}}(\alpha^{\mathrm{r}}_b, \alpha^{\mathrm{r}}_a)$.

Similarly,

$$P_{\mathcal{A},\mathrm{e}}(\alpha_a, \alpha^{\mathrm{r}}_b) = P_{\mathcal{A},\mathrm{e}}(\alpha_{-ks}, \alpha^{\mathrm{r}}_{b-a-ks}) = P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_0, \alpha_{a-b}) = P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_b, \alpha_a).$$

The identities $P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_a, \alpha_b) = P_{\mathcal{A},\mathrm{e}}(\alpha^{\mathrm{r}}_b, \alpha_a)$, $P_{\mathcal{A},\mathrm{v}}(\alpha_a, \alpha^{\mathrm{r}}_b) = P_{\mathcal{A},\mathrm{v}}(\alpha_b, \alpha^{\mathrm{r}}_a)$ and $P_{\mathcal{A},\mathrm{v}}(\alpha^{\mathrm{r}}_a, \alpha_b) = P_{\mathcal{A},\mathrm{v}}(\alpha^{\mathrm{r}}_b, \alpha_a)$ are proved in the same manner. □

Acknowledgements

References

[1] D. Aldous and J. Fill Reversible Markov chains and random walks on graphs. Unfinished manuscript.

[2] N. Alon, I. Benjamini, E. Lubetzky and S. Sodin. Non-backtracking random walks mix faster. *Commun. Contemp. Math.* **9** (4) (2007) 585–603. MR2348845

[3] A. Ben-Hamou and J. Salez. Cutoff for nonbacktracking random walks on sparse random graphs. *Ann. Probab.* **45** (3) (2017) 1752–1770. MR3650414

[4] N. Berestycki E. Lubetzky, Y. Peres and A. Sly. Random walks on the random graph. *Ann. Probab.* **46** (1) (2018) 456–490. MR3758735

[5] C. Bordenave, M. Lelarge and L. Massoulié. Non-backtracking spectrum of random graphs: Community detection and non-regular Ramanujan graphs. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science – FOCS 2015* 1347–1357. IEEE Computer Soc., Los Alamitos, CA, 2015. MR3473374

[6] C. Derman. A solution to a set of fundamental equations in Markov chains. *Proc. Amer. Math. Soc.* **5** (1954) 332–334. MR0060757

[7] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible Markov chains. *Ann. Appl. Probab.* **3** (3) (1993) 696–730. MR1233621

[8] P. G. Doyle and J. L. Snell. *Random Walks and Electric Networks. Carus Mathematical Monographs* **22**. Mathematical Association of America, Washington, DC, 1984. MR0920811

[9] P. G. Doyle and J. Steiner. Commuting time geometry of ergodic Markov chains. arXiv preprint, 2011. Available at arXiv:1107.2612.

[10] M. Dyer, L. A. Goldberg, M. Jerrum and R. Martin. Markov chain comparison. *Probab. Surv.* **3** (2006) 89–111. MR2216963

[11] R. Fitzner and R. van der Hofstad. Non-backtracking random walk. *J. Stat. Phys.* **150** (2) (2013) 264–284. MR3022459

[12] A. Gaudillière and C. Landim. A Dirichlet principle for non reversible Markov chains and some recurrence theorems. *Probab. Theory Related Fields* **158** (1–2) (2014) 55–89. MR3152780

[13] T. E. Harris. Transient Markov chains with stationary measures. *Proc. Amer. Math. Soc.* **8** (1957) 937–942. MR0091564

[14] M. Kempton. A non-backtracking Pólya's theorem. *J. Comb.* **9** (2) (2018) 327–343. MR3763648

[15] C.-H. Lee, X. Xu and D. Y. Eun. Beyond random walk and Metropolis–Hastings samplers: Why you should not backtrack for unbiased graph sampling. In *ACM SIGMETRICS Performance Evaluation Review* 319–330, **40**. ACM, New York, 2012.

[16] D. A. Levin and Y. Peres. *Markov Chains and Mixing Times*. J. G. Propp and D. B. Wilson (Eds) American Mathematical Society, Providence, RI, 2017. Second edition. With contributions by Elizabeth L. Wilmer and a chapter on "Coupling from the past". MR3726904

[17] E. Lubetzky and Y. Peres. Cutoff on all Ramanujan graphs. *Geom. Funct. Anal.* **26** (4) (2016) 1190–1216. MR3558308

[18] R. Lyons, R. Pemantle and Y. Peres. Ergodic theory on Galton–Watson trees: Speed of random walk and dimension of harmonic measure. *Ergodic Theory Dynam. Systems* **15** (3) (1995) 593–619. MR1336708

[19] R. Lyons and Y. Peres. *Probability on Trees and Networks. Cambridge Series in Statistical and Probabilistic Mathematics* **42**. Cambridge University Press, New York, 2016. MR3616205

[20] S. Northshield. Cogrowth of regular graphs. *Proc. Amer. Math. Soc.* **116** (1) (1992) 203–205. MR1120509

[21] R. Ortner and W. Woess. Non-backtracking random walks and cogrowth of graphs. *Canad. J. Math.* **59** (4) (2007) 828–844. MR2338235