

Coalescing random walk on unimodular graphs

Eric Foxall^{*} Tom Hutchcroft[†] Matthew Junge[‡]

Abstract

Coalescing random walk on a unimodular random rooted graph for which the root has finite expected degree visits each site infinitely often almost surely. A corollary is that an opinion in the voter model on such graphs has infinite expected lifetime. Additionally, we deduce an adaptation of our main theorem that holds uniformly for coalescing random walk on finite random unimodular graphs with degree distribution stochastically dominated by a probability measure with finite mean.

Keywords: coalescing random walk; unimodular random graph; voter model.

AMS MSC 2010: 60K35.

Submitted to ECP on April 6, 2018, final version accepted on May 2, 2018.

1 Introduction

Coalescing random walk (CRW) starts with one particle at each vertex of a locally finite, connected and undirected graph. Each particle then performs a continuous time edge-driven random walk, jumping along each edge adjacent to its present location according to a unit intensity Poisson process. Call the process *site recurrent* if every site is visited infinitely often almost surely.

Griffeath proved that CRW is site recurrent on \mathbb{Z}^d [Gri78]. Benjamini et. al. recently showed CRW is site recurrent on any bounded degree graph [BFGG⁺16]. As for unbounded degree graphs, [BFGG⁺16, Theorem 2 (ii)] shows the process is site recurrent on any Galton-Watson tree whose offspring distribution has an exponential tail.

Our goal is to show that CRW is site recurrent on any *unimodular random rooted graph* for which the root has finite expected degree. This is a general class of random graphs that arises frequently in applications. A quick corollary is that opinions in the voter model on such graphs have infinite expected lifetime. Previous works relating to CRW and the voter model on random graphs include [BPS12, HP15]. In particular, [HP15] establishes that every pair of particles eventually coalesce a.s. in any *recurrent unimodular random rooted graph*.

Before we give the definition, we start with a few interesting examples of transient unimodular random graphs with unbounded, but finite expected degree.

^{*}Department of Mathematical and Statistical Sciences, University of Alberta E-mail: e.t.foxall@gmail.com

[†]Department of Pure Mathematics and Mathematical Statistics, University of Cambridge
E-mail: t.hutchcroft@maths.cam.ac.uk

[‡]Department of Mathematics, Duke University E-mail: jungem@math.duke.edu

- (i) The unimodular version of a Galton-Watson tree, which is obtained from a standard Galton-Watson tree by first adding an extra child at the root to obtain an *augmented Galton-Watson tree*, and then biasing the law of the resulting random graph by the reciprocal of the degree of the root [AL07, Example 1.1.]. Unimodular Galton-Watson trees always have finite expected degree regardless of whether the offspring distribution has finite mean. Note this is much more general than the exponential tail requirement needed in [BFGG⁺16, Theorem 2 (ii)].
- (ii) Various random graphs obtained from \mathbb{Z}^d or \mathbb{R}^d for $d \geq 3$, including: Random geometric graphs defined in terms of point processes in \mathbb{R}^d [Pen03, Rou15]; Supercritical long-range percolation clusters in \mathbb{Z}^d ; Graphs obtained from \mathbb{Z}^d by replacing each edge with a random number of parallel edges with finite mean in some translation-invariant way.
- (iii) Curien's planar stochastic hyperbolic triangulations, which are transient versions of the uniform infinite planar triangulation [Cur16]; Various random graphs obtained from hyperbolic space \mathbb{H}^d for $d \geq 2$, such as hyperbolic random geometric graphs built from point processes in \mathbb{H}^d [BPP14].

Furthermore, combining our result with a simple compactness argument also yields that, whenever G is a *finite* graph whose degree distribution is stochastically dominated by some integrable reference distribution μ , CRW "looks recurrent" on G from the perspective of most vertices, in a quantitative way that depends only on the reference distribution μ . See Corollary 1.4 for a precise statement.

Definitions and theorem statement

To state our theorem we develop the framework and definitions more carefully. As in [BC12], a *rooted graph* is a pair (G, ρ) with $G = (V, E)$ and $\rho \in V$. An isomorphism between two rooted graphs is a graph isomorphism that maps the roots of the graphs to each other. Let \mathcal{G}_\bullet denote the set of isomorphism classes of locally finite, connected rooted graphs with no loops or multiple edges. Similarly we have $\mathcal{G}_{\bullet\bullet}$, the set of isomorphism classes of bi-rooted graphs (G, x, y) . When there is no cause for confusion, we use (G, ρ) to refer interchangeably to an element of \mathcal{G}_\bullet , or to a representative of its class, and similarly for $\mathcal{G}_{\bullet\bullet}$.

Coalescing random walk is well-defined on isomorphism classes of rooted graphs, since its distribution does not depend on the choice of representative. The same is true for random walk, with the understanding that the starting point of the walk may depend on the isomorphism class, and is mapped to an isomorphic location when passing between representatives. These two technicalities can be surmounted by choosing a root as the starting point of the walk.

A random rooted graph (G, ρ) is a random variable with values in \mathcal{G}_\bullet . Following [BC12] we use \mathbf{P} and \mathbf{E} to denote probability and expectation for a random rooted graph. We use \mathbf{P}_ρ^G and \mathbf{E}_ρ^G to denote the *quenched* probability and expectation of a process taking place on (G, ρ) , that is, the conditional law of the process given (G, ρ) . We use \mathbb{P} and \mathbb{E} to denote probabilities and expectations with respect to *annealed* measure, that is, the joint distribution of (G, ρ) and the process on (G, ρ) :

$$\mathbb{P} = \int \mathbf{P}_\rho^G d\mathbf{P}(G, \rho).$$

Before stating our theorem we give two definitions from [BC12]. Let (G, ρ) denote a random rooted graph and let $(X_n)_{n \geq 0}$ denote random walk on (G, ρ) with $X_0 = \rho$. Then (G, ρ) is *stationary* if

$$(G, X_0) = (G, X_1) \text{ in distribution.}$$

To introduce unimodularity we require something called the *mass-transport principle* or MTP for short. A random rooted graph (G, ρ) satisfies the MTP if for every Borel positive function $F: \mathcal{G}_{\bullet\bullet} \rightarrow \mathbb{R}_+$ we have

$$\mathbf{E} \left[\sum_{x \in V} F(G, \rho, x) \right] = \mathbf{E} \left[\sum_{x \in V} F(G, x, \rho) \right]. \tag{1.1}$$

A random rooted graph (G, ρ) is *unimodular* if it satisfies (1.1). We think of $F(G, \rho, x)$ as an amount of mass sent from ρ to x . The above says that the average mass sent from ρ to other vertices is the same as the average mass received by ρ . Note that every finite graph can be made into a unimodular random rooted graph by rooting it at a uniformly random vertex. Unimodular random rooted graphs were introduced by Benjamini and Schramm [BS01] and developed systematically by Aldous and Lyons [AL07].

With this notation we can state our theorem.

Theorem 1.1. *CRW is site recurrent \mathbb{P} -almost surely on any unimodular random rooted graph (G, ρ) for which $\mathbf{E}[\deg(\rho)] < \infty$.*

It is quite easy to construct counterexamples to show that the assumption $\mathbf{E}[\deg(\rho)] < \infty$ is necessary, for example by taking a canopy tree and replacing each edge at height n with a large number of parallel edges, see Example 3.1.

The proof of Theorem 1.1 uses the well-known duality between CRW and the voter model. The *voter model* is a process on (G, ρ) in which each site has a unique opinion which it spreads to neighbors at rate 1 (along each edge). Let ζ_t^ρ be the set of vertices at time t with the opinion started at ρ . The voter model is dual to CRW in the sense that ζ_t^ρ is equal in distribution to the set of particles in CRW that have coalesced and occupy site ρ at time t . We obtain Theorem 1.1 by proving a linear bound on the annealed second moment of $|\zeta_t^\rho|$.

Proposition 1.2. *Let (G, ρ) be a unimodular random rooted graph. Then $\mathbf{E}|\zeta_t^\rho|^2 \leq 1 + 2t\mathbf{E}[\deg(\rho)]$ for every $t \geq 0$.*

To prove Proposition 1.2 we start by defining the voter model and some important duality relationships in Section 2.1. We then use the MTP in Lemma 2.1 to show that the size of the voter model cluster *currently occupying* ρ , denoted $|\zeta_t^{(\rho)}|$, has the size-biased distribution of $|\zeta_t^\rho|$. A different duality relation at (2.2) relates the size of the voter model cluster occupying ρ to the total number of CRW particles coalesced with the one started at ρ . We use this in Proposition 2.4 to show that the expected number of other particles coalesced to a given particle is bounded by two times the integral of the vertex degrees along the particle’s random walk path. This is accomplished with a coupling that gives priority to the path followed by the particle started from ρ and ignores certain collisions. Using the stationarity and reversibility observed in Lemma 2.3, we obtain the bound $\mathbf{E}|\zeta_t^{(\rho)}| \leq 1 + 2t\mathbf{E}[\deg(\rho)]$ for all $t \geq 0$. The size-biasing relationship observed in Lemma 2.1 implies that $\mathbf{E}|\zeta_t^{(\rho)}| = \mathbf{E}[|\zeta_t^\rho|^2]$, which gives the desired second moment bound.

Theorem 1.1 follows in an elementary way from duality and Proposition 1.2, so we give the proof now.

Proof of Theorem 1.1. To compress our notation we let E denote E_ρ^G , P denote P_ρ^G , and ζ_t denote ζ_t^ρ . Let $c = 1 + 2\mathbf{E}[\deg(\rho)]$. Fix $m > c$ and let $\epsilon = c/m$. Applying Markov’s inequality to the random variable $E|\zeta_t|^2 = \mathbf{E}[|\zeta_t|^2 | (G, \rho)]$, we obtain that

$$\mathbb{P}(E|\zeta_t|^2 \leq mt) \geq 1 - \epsilon > 0$$

for every $t \geq 1$. For integer $k \geq 1$ let $A_k = \{E|\zeta_k|^2 \leq mk\}$ so that $\mathbb{P}(A_k) \geq 1 - \epsilon$ for all k . By the reverse Fatou’s lemma applied to $\mathbf{1}(A_k)$, $\limsup_k \mathbb{P}(A_k) \leq \mathbb{P}(\limsup_k A_k)$, and thus

$$1 - \epsilon \leq \mathbb{P}(A_k \text{ occurs infinitely often}).$$

So, with \mathbb{P} -probability at least $1 - \epsilon$, $E|\zeta_{t_i}|^2 \leq mt_i$ for some sequence of integers $t_i \uparrow \infty$. By Lemma 2.2, we have that $E|\zeta_t| = 1$. The Cauchy-Schwarz inequality then yields

$$P(|\zeta_{t_i}| > 0) \geq \frac{(E|\zeta_{t_i}|)^2}{E|\zeta_{t_i}|^2} \geq \frac{1}{mt_i} \tag{1.2}$$

for each t_i . Let $i_0 = 0$ and for each $k \geq 1$ let i_k be minimal such that $t_{i_k} \geq 2t_{i_{k-1}}$. Let $s_k = t_{i_k}$. Since 0 is an absorbing state, $t \mapsto P(|\zeta_t| > 0)$ is non-increasing and we deduce that

$$\int_0^\infty P(|\zeta_t| > 0) dt \geq \frac{1}{m} \sum_{k=1}^\infty \frac{s_{k+1} - s_k}{s_{k+1}} \geq \frac{1}{m} \sum_{k=1}^\infty \frac{1}{2} = \infty \tag{1.3}$$

with \mathbb{P} -probability at least $1 - \epsilon$. Let B_t be the event that in CRW there is a particle at ρ at time t . By the duality relationship described in (2.1), $P(B_t) = P(|\zeta_t| > 0)$ and so

$$\mathbb{P} \left(\int_0^\infty P(B_t) dt = \infty \right) \geq 1 - \epsilon.$$

Since ϵ can be made small by choosing m large, the above has \mathbb{P} -probability 1. Moreover, on a fixed graph G , it is known (see for example [BFGG⁺16]) that $\int_0^\infty P(B_t) dt = \infty$ implies ρ is recurrent (i.e., is occupied infinitely often as $t \rightarrow \infty$) almost surely. Combining these observations, it follows that ρ is recurrent \mathbb{P} -a.s. Since any property that holds a.s. for the root of a unimodular graph holds for all vertices a.s. [AL07, Lemma 2.3], it follows that all vertices are recurrent \mathbb{P} -a.s. \square

We now give two corollaries of our theorem. The expression at (1.3) is the expected survival time of the opinion started at ρ in the voter model on G . From it we obtain the following corollary.

Corollary 1.3. *Let (G, ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then the quenched expected survival time of the opinion started at ρ is infinite a.s.*

As remarked previously, we can use compactness arguments to derive “uniform” versions of Theorem 1.1 that are already interesting in the context of finite graphs. We denote by $\sigma_t(v)$ the first time after t that v is visited by a particle in CRW.

Corollary 1.4. *Let μ be a probability measure on \mathbb{N} with finite mean. Then there exists a function $f_\mu(t, u) : [0, \infty) \times [0, \infty) \rightarrow [0, 1]$ with*

$$\lim_{u \rightarrow \infty} f_\mu(t, u) = 0 \quad \forall t \geq 0$$

such that whenever (G, ρ) is a unimodular random rooted graph such that the law of $\deg(\rho)$ is stochastically dominated by μ , we have that

$$\mathbb{P}(\sigma_t(\rho) \geq t + u) \leq f_\mu(t, u)$$

for every $t, u \geq 0$.

Proof. Let M be the set of laws of unimodular random rooted graphs with root degree stochastically dominated by μ . Then M is compact with respect to the weak topology [Cur, Section 3.2.1], and since \mathcal{G}_\bullet is a Polish space when equipped with the local topology [Cur, Theorem 2], M is sequentially compact also.

Let $\sigma_t^{(r)}(\rho)$ be the first time after t that ρ is visited by a particle in the modified coalescing random walk in which particles are killed upon leaving the ball of radius r around ρ . Then for each $t, u \geq 0$, the function

$$\nu \mapsto \nu \left[P_\rho^G \left(\sigma_t^{(r)}(\rho) \geq t + u \right) \right]$$

is clearly continuous with respect to the weak topology on measures on \mathcal{G}_\bullet . Since $P_\rho^G(\sigma_t(\rho) \geq t + u) = \inf_{r \geq 1} P_\rho^G(\sigma_t^{(r)}(\rho) \geq t + u)$ it follows that the function

$$\nu \mapsto \nu [P_\rho^G(\sigma_t(\rho) \geq t + u)]$$

is an infimum of continuous functions and is therefore upper semi-continuous with respect to the weak topology on the space of probability measures on \mathcal{G}_\bullet . Thus, it obtains its maximum on M , and we denote this maximum by $f_\mu(t, u)$. Clearly $f_\mu(t, u)$ is decreasing in u .

Suppose for contradiction that $\inf_{u \geq 0} f_\mu(t, u) \geq \epsilon$ for some $\epsilon > 0$ and $t \in [0, \infty)$. Then for each $n \geq 1$ there exists $\nu_n \in M$ such that

$$\nu_n [P_\rho^G(\sigma_t(\rho) \geq t + m)] \geq \epsilon \text{ for all } 1 \leq m \leq n.$$

By compactness of M , the measures ν_n have a subsequential limit $\nu \in M$, and by upper semi-continuity we have that

$$\nu [P_\rho^G(\sigma_t(\rho) \geq t + m)] \geq \epsilon \text{ for all } m \geq 1.$$

This contradicts Theorem 1.1. □

2 Establishing Proposition 1.2

We follow the outline described in the introduction just preceding the statement of Proposition 1.2.

2.1 Voter model duality

We describe the graphical representation that allows us to construct and analyze coalescing random walk as well as a related *voter model*. We do this on a fixed graph; to obtain a reasonable joint measure \mathbb{P} of the random graph and CRW, it suffices to define the elements of the graphical representation (i.e., the processes $U_{(v,w)}$ given below) recursively over finite rooted graphs.

Let $G = (V, E)$ be an undirected, locally finite graph on which the continuous time edge-driven walk is non-explosive, i.e. makes at most finitely many jumps in any finite time interval a.s. This property always holds a.s. on unimodular random rooted graphs with finite expected degree [AL07, Corollary 4.4], and can therefore be safely assumed throughout our analysis. Double the edge set of G to form the set of directed edges $F = \{(v, w) : v, w \in E\}$. Define a family $\{U_{(v,w)} : (v, w) \in F\}$ of independent unit intensity Poisson point processes. Then a particle dropped at a spacetime location (v, s) follows a unique path $\Theta_t(v, s) : [s, \infty) \rightarrow V$ defined by updating to $\Theta_t(v, s) = u$, any time that $\Theta_{t-}(v, s) = w$ and $t \in U_{(w,u)}$; local finiteness and non-explosivity ensures that the jump times are a discrete subset of $[0, \infty)$, and hence that $\Theta_t(v, s)$ is well-defined for every $v \in V$, $s \in [0, \infty)$ and $t \geq s$ a.s. Coalescing random walk, labelled by ξ_t^v for the location of the particle that began at v , is defined by

$$\xi_t^v = \Theta_t(v, 0).$$

We let $\xi_t = \{w \in V : \exists v \text{ such that } \xi_t^v = w\}$ be the collection of all occupied sites at time t .

The related voter model ζ_t , initialized at any time $T > 0$, is defined on the time interval $[0, T]$ by letting

$$\zeta_t^v = \{w : \Theta_T(w, T - t) = v\}.$$

In particular, $\zeta_0^v = \{v\}$, and for $t > 0$, ζ_t^v and ζ_t^w are disjoint if $w \neq v$. The non-explosivity assumption ensures that $V = \bigcup_v \zeta_t^v$ for $t > 0$. Therefore $\{\zeta_t^v : v \in V\}$ is a partition of V ,

so we can think of the sets ζ_t^v as clusters vying for control of the territory V , with ζ_t^v being the cluster that started as $\{v\}$. Labelling clusters by their starting vertex, we could also define the voter model by $\zeta_t(v) = \Theta_T(v, T - t)$, with $\zeta_t(v)$ equal to the label of the cluster that contains v at time t .

From the graphical representation, looking backwards in time we see that the voter model is Markov with the following transition rule: for each $(v, w) \in F$, at rate 1, $\zeta_t(v) = \zeta_{t-}(w)$, or equivalently, the cluster containing w swallows v . There are various connections between ξ_t and ζ_t . The most basic is that

$$\zeta_T^v = \{w: \xi_T^w = v\}. \tag{2.1}$$

If we define $\zeta_t^{(v)} = \zeta_t^{\zeta_t^{(v)}}$ to be the cluster containing v , we note also that

$$\zeta_T^{(v)} = \{w: \xi_T^w = \xi_T^v\}, \tag{2.2}$$

and that $v \in \zeta_T^{\xi_T^v}$. If we think of the voter model as being constructed separately from CRW, these equalities are in distribution and hold for any fixed value of $T \geq 0$.

2.2 Proof of Proposition 1.2

Say that $(X_t)_{t \geq 0}$ is a continuous-time *edge-driven* random walk on a graph G if X_t moves along each edge at rate one.

We start by using the mass-transport principle to deduce a size-biasing property for the voter model.

Lemma 2.1. *Let (G, ρ) be a unimodular random graph on which the continuous time edge-driven random walk is non-explosive. Then $|\zeta_t^{(\rho)}|$ has the size-biased distribution of $|\zeta_t^\rho|$. That is, for every $t \geq 0$ and every integer $n \geq 0$,*

$$\mathbb{P}(|\zeta_t^{(\rho)}| = n) = n\mathbb{P}(|\zeta_t^\rho| = n).$$

Proof. The duality relation at (2.2) ensures that for any (G, ρ) we have $\rho \in \zeta_t^{\xi_t^\rho}$ when the voter model is initialized at time t . So, for $n \geq 1$ we have the disjoint union

$$\{|\zeta_t^{(\rho)}| = n\} = \bigcup_{x \in V(G)} \{|\zeta_t^x| = n, \rho \in \zeta_t^x\}. \tag{2.3}$$

Define $F: \mathcal{G}_{\bullet\bullet} \rightarrow \mathbb{R}_+$ by

$$F(G, \rho, x) = \mathbb{E}_G^\rho[\mathbf{1}(|\zeta_t^x| = n, \rho \in \zeta_t^x)].$$

Using the MTP as formulated at (1.1) and (2.3),

$$\mathbb{P}(|\zeta_t^{(\rho)}| = n) = \mathbf{E} \left[\sum_{x \in V(G)} F(G, \rho, x) \right] = \mathbf{E} \left[\sum_{x \in V(G)} F(G, x, \rho) \right]. \tag{2.4}$$

On the event $|\zeta_t^\rho| = n$, $\sum_{x \in V(G)} \mathbf{1}(x \in \zeta_t^\rho) = n$. Writing the indicator as a product and using Fubini's theorem,

$$\begin{aligned} \sum_{x \in V(G)} F(G, x, \rho) &= \sum_{x \in V(G)} \mathbb{E}_\rho^G[\mathbf{1}(|\zeta_t^\rho| = n, x \in \zeta_t^\rho)] \\ &= \mathbb{E}_\rho^G \left[\sum_{x \in V(G)} \mathbf{1}(|\zeta_t^\rho| = n) \mathbf{1}(x \in \zeta_t^\rho) \right] \\ &= n\mathbb{E}_\rho^G[\mathbf{1}(|\zeta_t^\rho| = n)]. \end{aligned}$$

Combining with (2.4),

$$\mathbb{P}(|\zeta_t^{(\rho)}| = n) = \mathbf{E}[nE_\rho^G[\mathbf{1}(|\zeta_t^{(\rho)}| = n)]] = n\mathbb{P}(|\zeta_t^{(\rho)}| = n),$$

as desired. \square

On bounded degree graphs $|\zeta_t^{(\rho)}|$ is a martingale. Indeed, it transitions as a nearest-neighbor simple random walk whose jump rate is equal to 2 times the number of (undirected) boundary edges of the cluster $\zeta_t^{(\rho)}$. A concern on unbounded degree graphs is that a lack of integrability might cause the size of the cluster to instead be a strict local martingale. The next lemma shows this is not the case in our setting.

Lemma 2.2. *Let (G, ρ) be a unimodular random rooted graph on which the continuous time edge-driven random walk is non-explosive. Then $E_\rho^G|\zeta_t^{(\rho)}| = 1$ for all $t \geq 0$ a.s.*

Proof. It follows from Lemma 2.1 that $\mathbb{E}|\zeta_t^{(\rho)}| = \sum_{n=0}^\infty \mathbb{P}(|\zeta_t^{(\rho)}| = n) = 1$. Thus, it suffices to show that, conditional on (G, ρ) , the process $|\zeta_t^{(\rho)}|$ is a supermartingale a.s., since then we have that $E_\rho^G|\zeta_t^{(\rho)}| \leq 1$ a.s., and since this random variable must integrate to 1 the claim follows.

Every non-negative local martingale is a supermartingale by Fatou’s lemma, and so it suffices to check that $|\zeta_t^{(\rho)}|$ is a local martingale. This follows easily after noting that the jump chain corresponding to the process $|\zeta_t^{(\rho)}|$ is a symmetric simple random walk absorbed at 0. \square

Lemma 2.3. *Let (G, ρ) be a unimodular random rooted graph with $\mathbf{E}[\deg(\rho)] < \infty$ and let $(X_t)_{t \geq 0}$ be a continuous-time edge-driven random walk with $X_0 = \rho$. Then, the measure of (G, ρ) is stationary and reversible for the random walk $(G, X_t)_{t \geq 0}$ on \mathcal{G}_\bullet .*

Proof. This follows from [AL07, Corollaries 4.3 and 4.4]. \square

The next result allows us to estimate the size of $\zeta_t^{(\rho)}$ in terms of the average degree of $\xi_t^{(\rho)}$ over time.

Proposition 2.4. *Let (G, ρ) be a fixed (nonrandom) element of \mathcal{G}_\bullet on which the continuous time edge-driven random walk is non-explosive, and let $(X_t)_{t \geq 0}$ be a continuous time edge-driven random walk on (G, ρ) with $X_0 = \rho$. Then*

$$E_\rho^G[|\zeta_t^{(\rho)}|] \leq 1 + 2 \int_0^t E_\rho^G[\deg(X_s)] ds. \tag{2.5}$$

Proof. Since the graph is fixed, use \mathbb{P} and \mathbf{E} to denote probability and expectation with respect to the process on (G, ρ) . First double up the edge set to $F = \{(x, y) : xy \in E\}$. Independently of (X_t) let $\{U_e : e \in F\}$ be a family of independent Poisson point processes with unit intensity on \mathbb{R}_+ , one for each directed edge. We think of the points

$$\{(e, t) : t \in U_e, e \in F\}$$

as arrows on the spacetime set $G \times \mathbb{R}_+$.

Define a version of coalescing random walk ξ_t on G using (X_t) and $\{U_e\}$ as follows. Recall that for $y \in V$, ξ_t^y is the position at time t of the particle that began at y . First, set $\xi_t^x = X_t$ for $t \geq 0$. Then, for $y \neq x$, let $\xi_0^y = y$ and follow these rules to determine ξ_t^y for $t > 0$.

1. If $\xi_t^y = X_t$, then $\xi_s^y = X_s$ for $s > t$.
2. If $\xi_{t-}^y = v \neq X_t$ and $t \in U_{(v,w)}$, then $\xi_t^y = w$.

In other words, particles remain stuck to (X_t) once they hit it, and otherwise follow the arrows given by the $\{U_e\}$. For $t \geq 0$ and $v \in V$ define the particle count N_t by

$$N_t(v) = |\{w : \xi_t^w = v\}|.$$

First, condition on $\gamma = \{X_t : t \geq 0\}$. For $v \in V$ let $S_v(\gamma) = \{t \geq 0 : X_t = v\}$. Then, for $t > 0$ let

$$U_\gamma(t) = \{((v, w), s) : (v, w) \in F, s \in U_{(v,w)} \cap S_w(\gamma) \cap [0, t]\},$$

and let $U_\gamma = \bigcup_{t>0} U_\gamma(t)$. The set U_γ consists of all arrows pointing towards the space-time path γ . Modify ξ_t so that it ignores U_γ , denoting the modified process by ξ_t^γ , with particle count N_t^γ . Note that $N_t(v) \leq N_t^\gamma(v)$ if $v \neq X_t$, and that conditional on γ , $\xi_t(\gamma)$ is independent of U_γ .

Since particles do not escape from X_t , it follows that

$$N_t(X_t) = 1 + \sum_{((v,w),s) \in U_\gamma(t)} N_s(v) + \sum_{s \leq t : X_s \neq X_{s-}} N_{s-}(X_s).$$

Using the above inequality we deduce the upper bound

$$N_t(X_t) \leq 1 + \sum_{((v,w),s) \in U_\gamma(t)} N_s^\gamma(v) + \sum_{s \leq t : X_s \neq X_{s-}} N_{s-}^\gamma(X_s).$$

Almost surely, $U_\gamma(t)$ and $\{s \leq t : X_s \neq X_{s-}\}$ are finite, and the above is a sum of finitely many terms. Conditioning, then noting that $\xi_t(\gamma)$ and thus N_t^γ are conditionally independent of U_γ given γ , we obtain

$$\begin{aligned} \mathbb{E}[N_t(X_t) \mid \gamma, U_\gamma] &\leq 1 + \sum_{((v,w),s) \in U_\gamma(t)} \mathbb{E}[N_s^\gamma(v) \mid \gamma] \\ &\quad + \sum_{s \leq t : X_s \neq X_{s-}} \mathbb{E}[N_{s-}^\gamma(X_s) \mid \gamma]. \end{aligned} \tag{2.6}$$

First we estimate $\mathbb{E}[N_s^\gamma(v) \mid \gamma]$ for $v \neq X_s$. Note that, conditioned on γ , $N_s^\gamma(v) = |\zeta_s|$ where $\{\zeta_r : r \in [0, s]\}$ is a voter model with $\zeta_0 = \{v\}$. Looking backwards in time, this voter model has the transitions:

1. If $w \in \zeta_{r-}$ and $r \in U_{(w,z)}$ with $z \neq X_{s-r}$ then $\zeta_r = \zeta_{r-} \setminus \{w\}$.
2. If $w \in \zeta_{r-}$ and $r \in U_{(z,w)}$ with $z \neq X_{s-r}$ then $\zeta_r = \zeta_{r-} \cup \{z\}$.
3. If $w \in \zeta_{r-}$ and $w = X_{s-r}$ then $\zeta_r = \zeta_{r-} \setminus \{w\}$.

Transitions 1. and 2. respectively increase and decrease $|\zeta_r|$ by 1 at the same rate, while transition 3. decreases $|\zeta_r|$, so $|\zeta_r|$ is a supermartingale. Therefore,

$$\mathbb{E}[N_s^\gamma(v) \mid \gamma] = \mathbb{E}[|\zeta_s|] \leq \mathbb{E}[|\zeta_0|] = 1.$$

Next we estimate $N_{s-}^\gamma(X_s)$. However, $N_{s-}^\gamma(X_s) = |\zeta_s|$ for the same process, except with $\zeta_0 = \{X_0\}$. Therefore $\mathbb{E}[N_{s-}^\gamma(X_s) \mid \gamma] \leq 1$ as well. Plugging into (2.6),

$$\mathbb{E}[N_t(X_t) \mid \gamma, U_\gamma] \leq 1 + |U_\gamma(t)| + |\{s \leq t : X_s \neq X_{s-}\}|. \tag{2.7}$$

Conditioned on γ , $\{|U_\gamma(t)| : t \geq 0\}$ is a counting process with time-varying intensity $\deg(X_t)$, so

$$\mathbb{E}[|U_\gamma(t)| \mid \gamma] = \int_0^t \deg(X_s) ds.$$

Without conditioning on γ , $\{|\{s \leq t: X_s \neq X_{s-}\}|: t \geq 0\}$ is a counting process with adapted intensity $\deg(X_t)$. Taking $\mathbb{E}[\cdot \mid \gamma]$ in (2.7) gives

$$\mathbb{E}[N_t(X_t) \mid \gamma] \leq 1 + \int_0^t \deg(X_s) ds + |\{s \leq t: X_s \neq X_{s-}\}|,$$

then taking $\mathbb{E}[\cdot]$, noting the jump rate of X_s is $\deg(X_s)$ and using Fubini's theorem, we find

$$\mathbb{E}[N_t(X_t)] \leq 1 + 2 \int_0^t \mathbb{E}[\deg(X_s)] ds.$$

With respect to our construction, $\xi_t^\rho = X_t$ for $t \geq 0$. Therefore, the result follows from the duality relation between coalescing random walk and the voter model, since

$$\zeta_t^{(\rho)} \stackrel{d}{=} \{x: \xi_t^x = \xi_t^\rho\}. \quad \square$$

Finally we complete the proof of Proposition 1.2.

Proof of Proposition 1.2. Lemma 2.3 ensures that $\mathbb{E}[\deg(X_t)] = \mathbb{E}[\deg(\rho)]$ for $t \geq 0$. Apply this to Proposition 2.4, so when we take $\mathbf{E}[\cdot]$ of (2.5) and apply Fubini's theorem we obtain

$$\mathbb{E}[|\zeta_t^{(\rho)}|] \leq 1 + 2t\mathbf{E}[\deg(\rho)]$$

for $t \geq 0$. We conclude the proof by observing that, by Lemma 2.1, we have $\mathbb{E}[|\zeta_t^\rho|^2] = \mathbb{E}[|\zeta_t^{(\rho)}|]$. \square

3 Counterexample when $\mathbf{E}[\deg(\rho)] = \infty$

Example 3.1. We now briefly discuss an example showing that the assumption of finite expected degree is needed for the validity of Theorem 1.1. Let (T, ρ) be the binary canopy tree. This graph arises as the Benjamini-Schramm limit of balls in the 3-regular tree, and can be described as follows: Every vertex v of T has a height $h(v) \in \{0, 1, \dots\}$. Vertices of height zero are leaves that each have a single height 1 neighbour, while vertices at height $n \geq 1$ have two neighbours at height $n - 1$ and a single neighbour at height $n + 1$. The isomorphism class of the rooted graph (T, v) depends only on $h(v)$, and it is well-known that (T, ρ) is unimodular if ρ is chosen to have a random height with $\mathbb{P}(h(\rho) = n) = 2^{-n-1}$.

Replace each edge of (T, ρ) between height n and $n+1$ with 3^n parallel edges to obtain a unimodular random rooted graph (G, ρ) . Observe that $\mathbf{E} \deg(\rho) = \sum_{n=1}^\infty h(n)3^n = \infty$. If Z_n is a discrete time simple random walk on G , then $h(Z_n)$ is a biased random walk on $\{0, 1, \dots\}$ that moves up with probability $3/5$ when it is not at height zero. From any starting point, the expected number of discrete time steps that the walk spends at any given height is bounded by a constant, and it follows that the expected time that the continuous time walk spends at height n is exponentially small in n , uniformly in the starting point. We deduce from this together with the Markov property that the continuous time walk on G is explosive, and moreover that the probability that it has not exploded by time t is exponentially small in t , uniformly in the starting point.

Now suppose that we start a continuous time random walk at every site of G . We run the walks independently, and particles do not coalesce. The natural coupling ensures, for each vertex v of G , the probability that v is occupied by at least one particle at time t in this model is at least the probability that v is occupied at time t in the CRW. On the other hand, we have by time reversal that the expected number of particles occupying v at time t in this model is equal to the probability that the random walk started at v has not exploded by time t . Since this probability decays exponentially with t , we deduce that the expected occupation time of v is finite, and hence that the CRW on G is not site recurrent.

References

- [AL07] David Aldous and Russell Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12:1454–1508, 2007. MR-2354165
- [BC12] Itai Benjamini and Nicolas Curien. Ergodic theory on stationary random graphs. *Electron. J. Probab.*, 17(93):1–20, 2012. MR-2994841
- [BFGG⁺16] Itai Benjamini, Eric Foxall, Ori Gurel-Gurevich, Matthew Junge, and Harry Kesten. Site recurrence for coalescing random walk. *Electron. Commun. Probab.*, 21:12 pp., 2016. MR-3522593
- [BPP14] Itai Benjamini, Elliot Paquette, and Joshua Pfeffer. Anchored expansion, speed, and the hyperbolic poisson voronoi tessellation. *arXiv preprint arXiv:1409.4312*, 2014.
- [BPS12] Martin T. Barlow, Yuval Peres, and Perla Sousi. Collisions of random walks. *Ann. Inst. Henri Poincaré Probab. Stat.*, 48(4):922–946, 2012. MR-3052399
- [BS01] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6(23):13, 2001. MR-1873300
- [Cur] Nicolas Curien. Random graphs: The local convergence point of view. Unpublished lecture notes. Available at <https://www.math.u-psud.fr/~curien/cours/cours-RG-V3.pdf>.
- [Cur16] Nicolas Curien. Planar stochastic hyperbolic triangulations. *Probab. Theory Related Fields*, 165(3–4):509–540, 2016. MR-3520011
- [Gri78] David Griffeath. Annihilating and coalescing random walks on \mathbb{Z}^d . *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 46(1):55–65, 1978. MR-0512333
- [HP15] Tom Hutchcroft and Yuval Peres. Collisions of random walks in reversible random graphs. *Electron. Commun. Probab.*, 20:6 pp., 2015. MR-3399814
- [Pen03] Mathew Penrose. *Random geometric graphs*, volume 5 of *Oxford Studies in Probability*. Oxford University Press, Oxford, 2003. MR-1986198
- [Rou15] Arnaud Rousselle. Recurrence or transience of random walks on random graphs generated by point processes in \mathbb{R}^d . *Stochastic Process. Appl.*, 125(12):4351–4374, 2015. MR-3406589