PHASE TRANSITION OF THE CONTACT PROCESS ON RANDOM REGULAR GRAPHS

JEAN-CHRISTOPHE MOURRAT, DANIEL VALESIN

ABSTRACT. We consider the contact process with infection rate λ on a random $(d+1)$-regular graph with n vertices, G_n. We study the extinction time τ_{G_n} (that is, the random amount of time until the infection disappears) as n is taken to infinity. We establish a phase transition depending on whether λ is smaller or larger than $\lambda_1(\mathbb{T}^d)$, the lower critical value for the contact process on the infinite, $(d+1)$-regular tree: if $\lambda < \lambda_1(\mathbb{T}^d)$, τ_{G_n} grows logarithmically with n, while if $\lambda > \lambda_1(\mathbb{T}^d)$, it grows exponentially with n. This result differs from the situation where, instead of G_n, the contact process is considered on the d-ary tree of finite height, since in this case, the transition is known to happen instead at the *upper* critical value for the contact process on $\mathbb{T}^d$.

1. INTRODUCTION

Let $G = (V, E)$ be a locally finite and connected graph. The contact process on G with parameter $\lambda > 0$ is the continuous-time Markov process $(\xi_t)_{t \geqslant 0}$ on the space of subsets of V whose transitions are given by:

$$(1.1) \quad \begin{array}{ll} \text{for every } x \in \xi_t, & \xi_t \to \xi_t \setminus \{x\} \quad \text{with rate 1,} \\ \text{for every } x \notin \xi_t, & \xi_t \to \xi_t \cup \{x\} \quad \text{with rate } \lambda \, |\{y \in \xi_t : \{x,y\} \in E\}|, \end{array}$$

where for a set A, we write $|A|$ to denote its cardinality. We identify subsets of V with elements of $\{0,1\}^V$ (through the indicator function).

The contact process is usually thought of as a model for the spread of an infection in a population. Vertices of the graph are interpreted as individuals, and states 0 and 1 indicate that an individual is healthy or infected, respectively. An infected site *recovers* with rate 1 and *infects* a neighbouring site with rate λ times the number of connecting edges (we will allow for multiple edges connecting a given pair of vertices). We refer the reader to [Li99] for a thorough introduction to the contact process, including the facts that we state without proof in this introduction.

We denote by $(\xi_t^A)_{t \geqslant 0}$ the contact process on G with initial configuration $\xi_0^A \equiv A$. If $A = \{x\}$, we write (ξ_t^x) instead of $(\xi_t^{\{x\}})$. We write $P_{G,\lambda}$ for a probability measure under which the contact process with parameter λ on G is defined.

The *extinction time* τ_G^A for the contact process on $G = (V, E)$ with initial configuration $A \subseteq V$ is defined by

$$\tau_G^A = \inf\{t : \xi_t^A = \varnothing\}.$$

Since the dynamics only allows for new infections to appear by transmission, the configuration $\varnothing$ is absorbing, so we have $\xi_t^A = \varnothing$ for any $t \geqslant \tau_G^A$. We write τ_G instead of τ_G^G and τ_G^x instead of $\tau_G^{\{x\}}$.

When the contact process is considered on infinite graphs, a central question is whether we have survival or extinction of the infection.

For a graph $G = (V, E)$, a finite set $A \subseteq V$ and $\lambda > 0$, define

$$p_{G,A,\lambda}^{\text{ext}} = P_{G,\lambda}\left[\tau_G^A < \infty\right] = P_{G,\lambda}\left[\exists t_0 : \xi_t^A = \varnothing \text{ for all } t \geqslant t_0\right],$$
$$p_{G,A,\lambda}^{\text{loc ext}} = P_{G,\lambda}\left[\exists t_0 : \xi_t^A \cap A = \varnothing \text{ for all } t \geqslant t_0\right],$$

the probabilities of *extinction* and *local extinction* of the contact process with parameter λ on G started from $\xi_0^A \equiv A$. It can be shown that we either have $p_{G,A,\lambda}^{\text{ext}} = 1$ for all A or $p_{G,A,\lambda}^{\text{ext}} < 1$ for all A; likewise, either $p_{G,A,\lambda}^{\text{loc ext}} = 1$ for all A or $p_{G,A,\lambda}^{\text{loc ext}} < 1$ for all A. We say that the contact process

- *dies out* if $p_{G,A,\lambda}^{\text{ext}} = 1$ for all A;
- *survives weakly* (or *globally but not locally*) if $p_{G,A,\lambda}^{\text{ext}} < 1$ and $p_{G,A,\lambda}^{\text{loc ext}} = 1$ for all A;
- *survives strongly* (or *locally*) if $p_{G,A,\lambda}^{\text{loc ext}} < 1$ for all A.

We let $\lambda_1(G) = \sup\{\lambda : p_{G,A,\lambda}^{\text{ext}} = 1\}$ and $\lambda_2(G) = \sup\{\lambda : p_{G,A,\lambda}^{\text{loc ext}} = 1\}$.

It is well known that for $G = \mathbb{Z}^d$, $0 < \lambda_1(\mathbb{Z}^d) = \lambda_2(\mathbb{Z}^d) =: \lambda_c(\mathbb{Z}^d) < \infty$. For $G = \mathbb{T}^d$, the infinite, $(d+1)$-regular tree (with $d \geqslant 2$), the situation is quite different, as we then have $0 < \lambda_1(\mathbb{T}^d) < \lambda_2(\mathbb{T}^d) < \infty$. In this case, if we take λ such that $\lambda_1(\mathbb{T}^d) < \lambda < \lambda_2(\mathbb{T}^d)$ and start the process from a single infection at a vertex x, then the infection has a chance of surviving, but it can only do so by propagating outwards from x; any finite neighbourhood of x only carries the infection for a finite amount of time, as required in the definition of weak survival.

The contact process on finite graphs (deterministic or random) has also been the subject of much investigation; below we will survey some of the past work that is most relevant to the object of interest of this paper. Let us start noting that for finite graphs, there is no question of survival or extinction: if G is finite, the infection almost surely disappears in finite time. One is thus interested in the behaviour of the infection before it disappears. The typical course of action goes as follows: we fix λ and the parameters that define the graph, then take a size parameter (e.g. the number of vertices) to infinity, so as to obtain a sequence of graphs G_n, and consider the asymptotic behaviour of some random quantity X_n associated with the contact process on G_n. Common choices for X_n are the extinction time τ_{G_n} and the average proportion of infected vertices before τ_{G_n}. Of particular interest are the cases in which the asymptotic behaviour of X_n depends sensitively on the choice of λ and the set of parameters that define the graph; this dependency can sometimes be related to a phase transition of the contact process on an infinite graph that is in some sense approximated by G_n.

This project has been first carried out for G_n equal to the subgraph of $\mathbb{Z}^d$ induced by the vertices contained in a box of side length n; see [CGOV84, Sc85, DL88, Ch94, DS88, Mo93, Mo99], or [Li99, Section I.3] for an overview. As suggested in the previous paragraph, the behaviour of τ_{G_n} exhibits a phase transition that mimics the phase transition of the contact process on $\mathbb{Z}^d$: τ_{G_n} grows logarithmically or exponentially with the volume of G_n respectively if λ is taken smaller or larger than $\lambda_c(\mathbb{Z}^d)$.

Let us now turn to finite trees. Let $\mathbb{T}_\ell^d$ denote the d-ary tree of height ℓ. In other words, $\mathbb{T}_\ell^d$ is a tree with a distinguished vertex o (called the root) so that o has degree d, all vertices at distance between 1 and $\ell - 1$ from o have degree $d + 1$ and all vertices at distance ℓ from o have degree 1. In [St01], bounds were obtained on $\tau_{\mathbb{T}_\ell^d}$ for different values of λ, and these bounds were improved in [CMMV13] to yield the result below.

Theorem 1.1. [CMMV13] *(a) For any $\lambda \in (0, \ \lambda_2(\mathbb{T}^d))$, there exists $c \in (0, \infty)$ such that, as $\ell \to \infty$,*

$$\frac{\tau_{\mathbb{T}^d_\ell}}{\log |\mathbb{T}^d_\ell|} \to c \ in \ probability.$$

(b) For any $\lambda \in (\lambda_2(\mathbb{T}^d), \infty)$, there exists $C \in (0, \infty)$ such that, as $\ell \to \infty$,

$$\frac{\log \tau_{\mathbb{T}^d_\ell}}{|\mathbb{T}^d_\ell|} \to C \ in \ probability.$$

Moreover, $\tau_{\mathbb{T}^d_\ell}$ divided by its expectation converges in distribution to the exponential distribution with parameter 1.

In short, the asymptotic behaviour of the extinction time on $\mathbb{T}^d_\ell$ has a phase transition at the value $\lambda = \lambda_2(\mathbb{T}^d)$. It is worth mentioning that the above theorem remains true if $\mathbb{T}^d_\ell$ is replaced by $\hat{\mathbb{T}}^d_\ell$, defined as the subgraph of $\mathbb{T}^d$ induced by the vertex set given by a distinguished root vertex o and all vertices at graph distance at most ℓ from o ($\mathbb{T}^d_\ell$ is obtained from $\hat{\mathbb{T}}^d_\ell$ by removing one of the sub-trees emanating from the root of $\hat{\mathbb{T}}^d_\ell$).

The main goal of this paper is to investigate the asymptotic behaviour of the extinction time of the contact process on random $(d+1)$-regular graphs, which we now describe. Fix $n \in \mathbb{N}$, the number of vertices, with the restriction that $n(d+1)$ be even. Let us define a random graph G_n with vertex set $V_n = \{1, \ldots, n\}$. We endow each vertex with $d+1$ half-edges; attaching two half-edges (say, one belonging to vertex i and another to vertex j) produces an edge between i and j. Then, any perfect matching on the set of half-edges produces a graph (again we note that we allow for multiple edges between two vertices and also edges with both extremities at a single vertex). We choose one perfect matching uniformly at random from all the $(n(d+1) - 1) \cdot (n(d+1) - 3) \cdots 3 \cdot 1$ possibilities, and this produces the random edge set E_n. Obviously, G_n is a $(d+1)$-regular graph; by using an alternate construction of G_n that is described in Section 3, it is easy to show that, for any R, with probability tending to 1 as $n \to \infty$, the ball of radius R around a given vertex has no cycle, and is thus isomorphic to $\hat{\mathbb{T}}^d_R$.

We write $\mathbb{P}_n$ for the law of G_n, and $\mathbb{P}_{n,\lambda}$ for a probability measure under which both G_n and the contact process on G_n with parameter λ are defined. Our main result is

Theorem 1.2. *(a) For every $\lambda \in (0, \lambda_1(\mathbb{T}^d))$ there exists $C < \infty$ such that*

$$\lim_{n \to \infty} \mathbb{P}_{n,\lambda} \left[\tau_{G_n} < C \log(n) \right] = 1.$$

(b) For every $\lambda \in (\lambda_1(\mathbb{T}^d), \infty)$ there exists $c > 0$ such that

$$\lim_{n \to \infty} \mathbb{P}_{n,\lambda} \left[\tau_{G_n} > e^{cn} \right] = 1.$$

Let us stress the most important point: contrary to the situation in Theorem 1.1, the behaviour of the extinction time has here a phase transition at $\lambda_1(\mathbb{T}^d)$ instead of $\lambda_2(\mathbb{T}^d)$. Although surprising at first, this phenomenon can be understood as follows: when ℓ is large, the tree $\mathbb{T}^d_\ell$ *as seen from a vertex chosen uniformly at random* does not at all locally look like $\mathbb{T}^d$. For example, with non-vanishing probability, the random vertex is a leaf of $\mathbb{T}^d_\ell$. The notion of local limit of a sequence of graphs, as formalized in [BS01], captures such features, and one can check that $\mathbb{T}^d_\ell$ does not converge locally to $\mathbb{T}^d$, but rather to another infinite graph called the *canopy tree* $\mathbb{CT}^d$ (see [Be, Example 5.14] for details). Moreover, one can show that

$$(1.2) \qquad\qquad \lambda_1(\mathbb{CT}^d) = \lambda_2(\mathbb{CT}^d) = \lambda_2(\mathbb{T}^d)$$

(local survival for $\lambda > \lambda_2(\mathbb{T}^d)$ can be derived from [CMMV13, Theorem 1.6], while extinction for $\lambda \leqslant \lambda_2(\mathbb{T}^d)$ follows from [Li99, Proposition 4.57]). In light of this, the logarithmic-exponential phase transition for the extinction time on $\mathbb{T}_\ell^d$ does happen at the *lower* critical value of the limiting graph after all, but one must take into account that this limiting graph is the canopy tree.

It would be very interesting to find wider classes of graphs (such as the configuration model described below, or see [vdH] for an exposition) for which one can prove that the same phase transition occurs at the lower critical value of the limiting graph. This is reminiscent of the question of the locality of the percolation critical probability, see [BNP11], [Be, Section 5.2], [MT13] and [SXZ14].

The graph G_n is a particular case of the class of graphs known as the *configuration model*. Whereas for G_n we assumed that all degrees are taken equal to $d + 1$, in the configuration model one allows for a random degree sequence, typically i.i.d. from some fixed degree distribution. The contact process on these graphs has lately received attention ([CD09, MMVY12, MVY13]). In particular, it was proved that, if the degree distribution is a power law (and the graph is assumed to be connected), then the extinction time grows exponentially with the number of vertices *regardless of the value of λ*. In this sense, one can say that the contact process on the configuration model with a power law degree distribution has no phase transition: it is "always supercritical". Theorem 1.2 shows that, if the degrees are constant, then there is a phase transition. In fact, putting together Lemma 4.2 below with the main result of [MMVY12], we get

Theorem 1.3. *Let $\mathcal{G}_{n,d}$ be the set of connected graphs with n vertices and degree bounded by $d + 1$.*

(a) *For any $\lambda \in (0, \lambda_1(\mathbb{T}^d))$, there exists $C < \infty$ such that*
$$\lim_{n \to \infty} \inf_{G \in \mathcal{G}_{n,d}} P_{G,\lambda}\left[\tau_G < C \log n\right] = 1.$$

(b) *For any $\lambda \in (\lambda_c(\mathbb{Z}), \infty)$, there exists $c > 0$ such that*
$$\lim_{n \to \infty} \inf_{G \in \mathcal{G}_{n,d}} P_{G,\lambda}\left[\tau_G > e^{cn}\right] = 1.$$

In particular, if we take the contact process on the configuration model with a degree distribution with *bounded support*, the extinction time exhibits a phase transition on λ. A problem that we believe to be of much interest is whether such a phase transition also occurs if the degree distribution has unbounded support, but a light tail. Equally interesting is the corresponding question for supercritical Erdős-Rényi random graphs. In this direction, we refer the reader to the very recent remarkable results of [MS15], where existence of a phase transition for the contact process on certain graphs with unbounded degrees, such as Delaunay triangulations of $\mathbb{R}^d$, is shown.

We now comment on the proofs of the two parts of Theorem 1.2. Part (a) relies on the fact, stated in Lemma 4.2, that the contact process on any graph of degree bounded by $d + 1$ is stochastically dominated (in terms of the number of infected vertices) by a contact process with the same λ on $\mathbb{T}^d$. Although this is a very elementary and natural statement, we did not find it in the literature, so we give a proof, which relies on the concept of universal covering of a graph. Once Lemma 4.2 has been established, we invoke known bounds for the extinction time of the contact process on $\mathbb{T}^d$ in the extinction regime in order to conclude the proof of part (a).

As a side note, we remark that Lemma 4.2 also applies to infinite graphs, and as a consequence we prove:

Theorem 1.4. *If $\lambda \leqslant \lambda_1(\mathbb{T}^d)$, then the contact process with parameter λ on any graph with degree bounded by $d + 1$ dies out.*

The proof of part (b) of Theorem 1.2 is more involved, and relies on a geometric property of G_n (Theorem 3.4). The idea can be roughly summarised as follows. Assume that at time 0, there are δn infected vertices ($\delta > 0$ small) with the property that within distance r of each of these vertices, one can find a copy of $\mathbb{T}_\ell^d$, the copies (and the paths to them) being pairwise disjoint. Let us say that a set satisfying this property is *regenerative*. If ℓ can be chosen sufficiently large, then after some time, the infection will spread to $k\delta n$ vertices with k large, except for an event whose probability is exponentially small in n. The crucial geometric property of G_n that we need is that if k is sufficiently large, then with probability tending to 1, *any* subset of size $k\delta n$ contains a regenerative subset of size δn. This enables to iterate the argument. Since the probability of failure of a given step is exponentially small, the extinction time must be exponentially long.

In Section 2, we give a summary of the notation we use, an exposition of the graphical construction of the contact process and some known results about the contact process on trees. Sections 3 and 4 are devoted to proving parts (b) and (a) of Theorem 1.2 respectively.

2. Notation and preliminary results

2.1. Summary of notation for sets and graphs.
For a set A, we write $|A|$ for the cardinality of A and $\mathbf{1}_A$ for the indicator function of A.

We will bother the reader with defining graphs, because we want to allow loops (edges between a vertex and itself) and multiple edges with the same endpoints, and this requires some care with the notation. Given sets V and E and a function

$$\zeta : E \to \{\text{subsets of } V \text{ with cardinality 1 or 2}\},$$

the triple $G = (V, E, \zeta)$ is a *graph*; elements of V are *vertices* and elements of E are *edges*. $e \in E$ is said to be an edge between vertices x and y if $\zeta(e) = \{x, y\}$; if such an edge exists, we say that x and y are neighbours. If $\zeta(e) = \{x\}$, we say e is a loop at x.

We may abuse notation and write $v \in G$ instead of $v \in V$ when convenient. Also, we will omit the presence of the function ζ, simply writing $G = (V, E)$ as is usual.

The *degree* of a vertex is the number of non-loop edges that contain it plus twice the number of loops that contain it. The *graph distance* between vertices x and y is the length of a shortest path between x and y, and is denoted by $\mathsf{dist}(x, y)$, or $\mathsf{dist}_G(x, y)$ when we want to make the graph explicit. A *cycle* is a sequence of vertices $x_0, \ldots, x_k$ so that $x_0, \ldots, x_{k-1}$ are all distinct, $x_k = x_0$ and for each i, x_i and x_{i+1} are neighbours.

The set $\vec{E}$ of oriented edges of G is defined as follows: for each non-loop $e \in E$, we add to $\vec{E}$ two oriented edges corresponding to the two possible orientations of e; for each loop in e, we also add to $\vec{E}$ two "oriented edges" from x to x (though this will be quite unimportant). A generic element of $\vec{E}$ will be denoted by $\vec{e}$. Note that, if there are k edges containing x and y, then there are $2k$ oriented edges containing x and y. We write $v_0(\vec{e})$ and $v_1(\vec{e})$ to denote the starting and ending vertex of $\vec{e}$, respectively. We also let $u(\vec{e}) \in E$ denote the unoriented edge to which $\vec{e}$ is associated.

A *rooted graph* is a pair (ρ, G) where G is a graph and $\rho \in G$. Given two rooted graphs (ρ, G) and (ρ', G') with $G = (V, E)$, $G' = (V', E')$, we say that $f : V \to V'$ is an *embedding* of (ρ, G) into (ρ', G') if

- $f(\rho) = \rho'$,
- f is injective,
- for every $u, v \in V$, the number of edges in G containing u and v is less than or equal to the number of edges of G' containing $f(u)$ and $f(v)$.

(Of course, if G and G' have no loops or multiple edges between vertices, the last property means that $f(u)$ and $f(v)$ are neighbours in G' whenever u and v are neighbours in G.) We say that (ρ', G') *embeds* (ρ, G) if there exists an embedding of (ρ, G) into (ρ', G'); that (ρ, G) and (ρ', G') are *isomorphic* if each embeds the other.

For the rest of the paper, we fix $d \geqslant 2$ and omit d from $\mathbb{T}^d$ and $\mathbb{T}_\ell^d$, thus writing $\mathbb{T}$ and $\mathbb{T}_\ell$ respectively. The root of $\mathbb{T}_\ell$ is denoted by o.

2.2. Graphical construction of the contact process.
Almost every paper on the contact process contains a brief exposition of its graphical construction, so this is by now quite a redundant addition. We nevertheless include it here because we allow our graphs to contain loops and parallel edges, so we need to be careful with the notation to avoid confusion.

Fix a graph $G = (V, E)$ and an infection rate $\lambda > 0$. Assume given the following families of independent Poisson point processes on $[0, \infty)$: $(T_{\vec{e}})_{\vec{e} \in \vec{E}}$, all with rate λ, and $(D_x)_{x \in V}$, all with rate 1. We now imagine G is embedded on the xy-plane and add marks on the vertical lines $(V \cup \vec{E}) \times [0, \infty)$ as follows: for each x and each $t \in D_x$, we add a so-called *recovery mark* at (x, t), and for each $\vec{e}$ and each $t \in T_{\vec{e}}$, we add a so-called *transmission arrow* at $(\vec{e}, t)$.

Given $x, y \in V$ and $0 \leqslant s \leqslant t$, we say that (x, s) and (y, t) are connected by an *infection path* (and write $(x, s) \leftrightarrow (y, t)$) if there exists a right-continuous function $\gamma : [s, t] \to V$ such that

- $\gamma(s) = x$, $\gamma(t) = y$,
- $(r, \gamma(r)) \notin D_{\gamma(r)}$ for all $r \in [s, t]$,
- $\gamma(r-) \neq \gamma(r)$ implies $r \in T_{\vec{e}}$ for some $\vec{e}$ with $v_0(\vec{e}) = \gamma(r-)$, $v_1(\vec{e}) = \gamma(r)$.

In other words, $(r, \gamma(r))_{s \leqslant r \leqslant t}$ must be a path from (x, s) to (y, t) which does not cross any recovery mark and is allowed to traverse transmission arrows.

For $x \in V$, let $\xi_t^x = \{y \in V : (x, 0) \leftrightarrow (y, t)\}$; for $A \subseteq V$, let $\xi_t^A = \cup_{x \in A} \xi_t^x$. Then, $(\xi_t^A)_{t \geqslant 0}$ is a Markov process on $\{0, 1\}^V$ with the same distribution as the process defined earlier in (1.1). The graphical construction has the advantage that all the contact processes $((\xi_t^A)_{t \geqslant 0})_{A \subseteq V}$ are defined in the same probability space and, if $A \subseteq B$, then $\xi_t^A \subseteq \xi_t^B$ for all t. Note also that $\xi_{s+t}^A = \cup_{x \in \xi_t^A} \{y : (x, t) \leftrightarrow (y, t+s)\}$.

2.3. Known results on the contact process on $\mathbb{T}$.
We now list some definitions and results about the contact process on $\mathbb{T}$ which we will need; all of them can be found in Section I.4 of [Li99]. The reader is encouraged to skip this subsection and refer back to it when each result is needed.

We start defining a "generation number" function $\ell : \mathbb{T} \to \mathbb{Z}$ as follows. Take an embedding $\kappa : \mathbb{Z} \to \mathbb{T}$ so that $\kappa(0) = o$. If $x \in \kappa(\mathbb{Z})$, set $\ell(x) = \kappa^{-1}(x)$. If $x \notin \kappa(\mathbb{Z})$, let y be the first vertex of $\kappa(\mathbb{Z})$ reached in a path from x to $\kappa(\mathbb{Z})$; set $\ell(x) = \ell(y) + \mathsf{dist}(x, y)$.

Now, given $\rho > 0$, as in [Li99] equation (4.23), page 87, we put

$$\phi_\lambda(\rho) = \lim_{t \to \infty} \mathbb{E}_{\mathbb{T}, \lambda} \left[\sum_{x \in \mathbb{T}} \xi_t^o(x) \cdot \rho^{\ell(x)} \right]^{\frac{1}{t}};$$

the limit is shown to exist (page 87). In Proposition 4.27(a) and (b) we respectively have

$$(2.1) \qquad \phi_\lambda(\rho) = \phi_\lambda\left(\frac{1}{\rho d}\right)$$

and

$$(2.2) \qquad \mathbb{E}_{\mathbb{T}, \lambda}\left[|\xi_t^o|\right] \leqslant C_\lambda \cdot \phi_\lambda(1)^t, \quad t \geqslant 0,$$

for some constant $C_\lambda > 0$. In Proposition 4.44(a), we have

$$(2.3) \qquad \frac{1}{\sqrt{d}} \leqslant \rho < \rho' \quad \text{and} \quad \phi_\lambda(\rho') \geqslant 1 \quad \Longrightarrow \quad \phi_\lambda(\rho) < \phi_\lambda(\rho').$$

As in equation (4.48), page 96, we define for $\lambda > 0$

$$(2.4) \qquad \beta(\lambda) = \lim_{n \to \infty} \mathbb{P}_{\mathbb{T},\lambda} \left[\exists t : \xi_t(\kappa(n)) = 1 \right]^{\frac{1}{n}} ;$$

again the limit is shown to exist (page 96). Regarding β, we will need the following facts, which can be found in Theorem 4.83, Corollary 4.87 and Theorem 4.130 of [Li99]:

$$(2.5) \qquad \beta(\lambda) < \frac{1}{\sqrt{d}} \quad \Longrightarrow \quad \phi_\lambda(\beta(\lambda)) = \phi_\lambda\left(\frac{1}{\beta(\lambda) \cdot d}\right) = 1,$$

$$(2.6) \qquad \beta(\lambda_1) = \frac{1}{d},$$

$$(2.7) \qquad \beta \text{ is strictly increasing in } [0, \lambda_2(\mathbb{T})].$$

3. Supercritical regime

The goal of this section is to prove part (b) of Theorem 1.2. We first state for later usage a classical large deviation result on binomial random variables, whose proof is standard.

Lemma 3.1. *Let* $\mathsf{Bin}(m, p)$ *denote a binomial random variable with parameters* $m \in \mathbb{N}$ *and* $p \in [0, 1]$ *(defined with respect to the measure* $\mathbb{P}$, *say). For every* $\delta \geqslant 0$,

$$\mathbb{P}[\mathsf{Bin}(m, p) \geqslant (p + \delta)m] \leqslant e^{-m\psi_p(\delta)},$$

where

$$
\begin{aligned}
\psi_p(\delta) &= \sup_\lambda \left[\lambda(p + \delta) - \log(1 - p + pe^\lambda) \right] \\
(3.1) \qquad &= (p + \delta) \log\left(\frac{p + \delta}{p}\right) + (1 - p - \delta) \log\left(\frac{1 - p - \delta}{1 - p}\right).
\end{aligned}
$$

We will need the following estimate:

Lemma 3.2. *For every* $\lambda > \lambda_1(\mathbb{T})$, *there exist* $R, p_0 > 0$ *and* $\alpha_0 > 1$ *such that, for* ℓ *large enough,*

$$P_{\mathbb{T}_\ell,\lambda} \left[|\xi^o_{R\ell}| \geqslant \alpha_0^\ell \right] \geqslant p_0.$$

Proof. Let $\mathbb{T}_\infty$ be the infinite, rooted, d-ary tree (that is, the infinite rooted tree in which the root o has degree d and all other vertices have degree $d + 1$). In [SS98] it is shown that for every $\lambda > \lambda_1(\mathbb{T})$, there exist $U > 0$, $\gamma > 1$ and $p > 0$ such that, for $\ell \in \mathbb{N}$ large enough,

$$(3.2) \qquad P_{\mathbb{T}_\infty,\lambda} \left[|\xi^o_{U\ell}| \geqslant \gamma^\ell \right] \geqslant p,$$

where $\beta(\lambda)$ is defined in (2.4). We remark that:
  (1) the above estimate follows from equation (2.6) in [SS98];
  (2) the running hypothesis in [SS98] is that $\lambda > \lambda_2(\mathbb{T})$, but (3.2) holds true more generally, for $\lambda > \lambda_1(\mathbb{T})$. Indeed, its proof depends on two ingredients. First, their Lemma 1, which in fact holds for *any* $\lambda > 0$. Second, an argument involving branching processes (see the paragraphs preceding their equation (2.6)) which is applicable as long as $d \cdot \beta(\lambda) > 1$. By (2.6) and (2.7), this is indeed the case when $\lambda > \lambda_1(\mathbb{T})$;
  (3) we can then choose γ as any constant between 1 and $d \cdot \beta(\lambda)$.

Using equation (4.76) in [Li99], one can readily show that there exists $S > 0$ such that, for all $t > 0$,

$$P_{\mathbb{T}_\infty,\lambda}\left[\xi_s^o(y) = 0 \text{ for all } s \leqslant t \text{ and } y \text{ with } \mathsf{dist}(o,y) \geqslant St\right]$$

$$\geqslant P_{\mathbb{T},\lambda}\left[\xi_s^o(y) = 0 \text{ for all } s \leqslant t \text{ and } y \text{ with } \mathsf{dist}(o,y) \geqslant St\right] > 1 - \frac{p}{2}.$$

Combining these two estimates and setting $R = S^{-1}$ and $\alpha_0 = \gamma^{R/U} > 1$ we get

$$\frac{p}{2} \leqslant P_{\mathbb{T}_\infty,\lambda}\left[|\xi_{R\ell}^o| \geqslant \alpha_0^\ell, \ \xi_s^o(y) = 0 \text{ for all } s \leqslant R\ell \text{ and } y \text{ with } \mathsf{dist}(o,y) \geqslant \ell\right]$$

$$\leqslant P_{\mathbb{T}_\ell,\lambda}\left[|\xi_{R\ell}^o| \geqslant \alpha_0^\ell\right]. \hspace{4cm} \square$$

We have the following consequence of the above result:

Lemma 3.3. *For every $\lambda > \lambda_1(\mathbb{T})$ and $r > 0$, there exist R, $\sigma > 0$ and $\alpha > 1$ such that for every ℓ large enough, the following holds. For any graph G with vertices x, y such that $\mathsf{dist}(x,y) \leqslant r$ and (y,G) embeds $(o,\mathbb{T}_\ell)$, we have*

$$P_{G,\lambda}\left[|\xi_{R\ell}^x| \geqslant \alpha^\ell\right] > \sigma.$$

Proof. Let R, p_0 and α_0 be as in the previous lemma; we have

$$P_{G,\lambda}\left[|\xi_{R(\ell+1)}^x| \geqslant \alpha_0^\ell\right] \geqslant P_{G,\lambda}\left[\xi_R^x(y) = 1\right] \cdot P_{G,\lambda}\left[|\xi_{R\ell}^y| \geqslant \alpha_0^\ell\right] \geqslant P_{G,\lambda}\left[\xi_R^x(y) = 1\right] \cdot p_0.$$

We can now find $\sigma > 0$ that bounds the right-hand side from below uniformly on $\mathsf{dist}(x,y) \in \{1,\ldots,r\}$ (one can guarantee that the event $\{\xi_R^x(y) = 1\}$ occurs by making finitely many prescriptions on Poisson processes in the graphical construction on a geodesic path connecting x and y and the time interval $[0,R]$). We then pick any $\alpha \in (1,\alpha_0)$; the result follows from observing that $\alpha_0^{\ell-1} > \alpha^\ell$ for ℓ large enough. $\square$

We say that a set of vertices $W \subseteq V_n$ is *ℓ-regenerative* if there exists a family $(G_v')_{v \in W}$ of subgraphs of G_n that are pairwise disjoint and such that for every $v \in W$, the following two conditions hold:

- G_v' contains v,
- there exists $x \in G_v'$ such that the distance in G_v' between x and v is 4 and (x,G_v') embeds $(o,\mathbb{T}_\ell)$.

The crucial geometric property of G_n that we need is

Theorem 3.4 (Finding large regenerative subsets). *For any $\ell \in \mathbb{N}$ and $\varepsilon > 0$ sufficiently small (depending on ℓ), the following holds with $\mathbb{P}_n$-probability tending to 1 as n tends to infinity. From every $W \subseteq V_n$ of cardinality at least εn, one can extract an ℓ-regenerative subset of cardinality at least $\frac{\varepsilon}{40d^6}n$.*

Remark 3.5. In what follows, in order to prevent the notation from getting too heavy, we will pretend that certain quantities, such as εn in the above proposition, are integers. It should be clear that in a correct but overscrupulous writing, one should take the pertinent integer parts or add or subtract 1 at certain places, but that this does not change our proofs in any relevant way.

We prove part (b) of Theorem 1.2 using Theorem 3.4, before turning to the proof of the latter.

Proof of part (b) of Theorem 1.2. We fix $\lambda > \lambda_1(\mathbb{T})$, and choose constants according to the following steps:

(1) let α, R, σ correspond to λ and $r = 4$, as in Lemma 3.3;
(2) take ℓ large enough, as required by Lemma 3.3, and also so that $\alpha^\ell > \frac{80d^6}{\sigma}$;
(3) take ε corresponding to ℓ as in Theorem 3.4.

Assume G_n satisfies the property stated in Theorem 3.4, namely

every $W \subseteq V_n$ with $|W| \geqslant \varepsilon n$ has an ℓ-regenerative subset of cardinality $\frac{\varepsilon}{40d^6}n$.

We will now prove that, for some constant $c > 0$ which does not depend on n,

(3.3)   for all $W \subseteq V_n$ with $|W| \geqslant \varepsilon n$, $P_{G_n}\left[|\xi^W_{R\ell}| \geqslant \varepsilon n\right] \geqslant 1 - e^{-cn}$.

This will imply the statement of Theorem 1.2.

We fix W with $|W| \geqslant \varepsilon n$ and extract from it an ℓ-regenerative subset W' of cardinality $n' := \frac{\varepsilon}{40d^6}n$; we enumerate its elements, $W' = \{v_1, \ldots, v_{n'}\}$. By the definition of ℓ-regenerative sets, there exist pairwise disjoint subgraphs of G, $G'_{v_1}, \ldots, G'_{v_{n'}}$ such that $v_i \in G'_{v_i}$ and there exists $x_i \in G'_{v_i}$ such that $\mathsf{dist}(v_i, x_i) = 4$ and (x_i, G'_{v_i}) embeds $(o, \mathbb{T}_\ell)$.

For each v_i, let $(\zeta^{v_i}_t)_{t \geqslant 0}$ be the contact process on G'_{v_i}, started from only v_i infected, and built using the same family of Poisson processes as the original contact process on G. If $i \neq j$, then $(\zeta^{v_i}_t)$ and $(\zeta^{v_j}_t)$ are independent (since the G'_{v_i}'s are disjoint) and moreover, $\xi^W_t \supseteq \xi^{W'}_t \supseteq \cup_{v \in W'} \zeta^v_t$ for all t.

Define the events

$$E_i = \left\{|\zeta^{v_i}_{R\ell}| \geqslant \alpha^\ell\right\}, \qquad 1 \leqslant i \leqslant n'.$$

We then have $P_{G_n}[E_i] \geqslant \sigma$. Thus, by standard large deviation estimates for binomial random variables (see Lemma 3.1), we have

$$P_{G_n}\left[\sum_{i=1}^{n'} \mathbf{1}_{E_i} \geqslant \frac{\sigma}{2}n'\right] \geqslant 1 - e^{-c(\varepsilon, \sigma)n}.$$

Finally, if the event in the above probability occurs, we have

$$|\xi^W_{R\ell}| \geqslant \alpha^\ell \cdot \sum_{i=1}^{n'} \mathbf{1}_{E_i} \geqslant \alpha^\ell \cdot \frac{\sigma}{2}n' > \varepsilon n.$$

This completes the proof. $\square$

We now give some definitions and preliminary results for the proof of Theorem 3.4. We say that a set of vertices $W \subseteq V$ is *well separated* if the following two properties hold:

- for every $v \in W$, the 3-neighbourhood of v is cycle-free (for $r \in \mathbb{N}$, the r-neighbourhood of a vertex $v \in V$ is the set of vertices w whose distance to v is at most r),
- for every two distinct $v, w \in W$, the 3-neighbourhoods of v and w are disjoint.

The key step of the proof lies in the following proposition.

Proposition 3.6 (Key estimate)**.** *For every $\ell \in \mathbb{N}$, there exists $c > 0$ such that for every $\varepsilon > 0$ suficiently small (depending on ℓ) and every n large enough,*

$$\mathbb{P}\left[\begin{array}{l} \text{the set } \{1, \ldots, \varepsilon n\} \text{ is well separated but} \\ \text{has no } \ell\text{-regenerative subset of size } \varepsilon n/5 \end{array}\right] \leqslant (c\varepsilon^{\frac{6}{5}})^{\varepsilon n}.$$

Remark 3.7. In order to prove Theorem 3.4, we will use a union bound over all sets of cardinality εn. Roughly speaking, there are about $\varepsilon^{-\varepsilon n}$ such sets ($\varepsilon \ll 1$). So the crucial point in Proposition 3.6 is that the exponent $\frac{6}{5}$ appearing there is strictly larger than 1.

Before proving this key estimate, we introduce some terminology. A *semi-graph* $g = (V, \mathcal{E}, \mathcal{H})$ is a triple consisting of a set V of vertices, a set $\mathcal{E}$ of edges between points of V, and a set $\mathcal{H}$ of *half-edges*, each half-edge being attached to some vertex in V. Given two half-edges h and h', we write $h + h'$ to denote the edge obtained by "gluing together" the half-edges h and h' (that is to say, if h is attached to a vertex v and h' to a vertex w, then $h + h'$ is an edge connecting v and w). The distance in the semi-graph g is simply the distance in the graph $(V, \mathcal{E})$.

Consider the semi-graph $g = (V_n, \mathcal{E}, \mathcal{H})$ such that $\mathcal{E} = \varnothing$ and each vertex in V_n has exactly $d + 1$ half-edges. We can construct a random regular graph with distribution $\mathbb{P}_n$ by the following recursive procedure. Take an arbitrary half-edge h in $\mathcal{H}$ (call it the *elected* half-edge); take a half-edge h' uniformly at random in $\mathcal{H} \setminus \{h\}$; add the edge $h + h'$ to the set $\mathcal{E}$, remove the half-edges h and h' from $\mathcal{H}$; repeat until the set of half-edges is empty. (Recall that we assume $(d + 1)n$ to be even, so that there is an even number of half-edges to begin with). The resulting random graph has law $\mathbb{P}_n$. A feature of this procedure that will be crucial in our reasoning is that at each step, we have the freedom to choose the elected half-edge as we wish among the half-edges of $\mathcal{H}$.

It will be convenient to write operations on sets such as those done in the above construction in the more symbolic form

$$(3.4) \qquad \mathcal{E} \leftarrow \mathcal{E} \cup \{h + h'\}, \qquad \mathcal{H} \leftarrow \mathcal{H} \setminus \{h, h'\}.$$

Proof of Proposition 3.6. Fix $\ell \in \mathbb{N}$. Also fix $\varepsilon \in (0, 1)$ and $n \in \mathbb{N}$; for the moment we do not specify their choice. We take the semi-graph $g = (V_n, \mathcal{E}, \mathcal{H})$ such that $\mathcal{E} = \varnothing$ and such that every site has exactly $(d + 1)$ half-edges. We denote by $\mathcal{W}$ the set of vertices $\{1, \dots, \varepsilon n\}$. As long as there is a half-edge h attached to a vertex at distance 2 or less from $\mathcal{W}$, we choose this as the elected half-edge, pick another half-edge h' uniformly at random from $\mathcal{H} \setminus \{h\}$, and do the operations in (3.4). When there is no longer any such half-edge, the graph is still incomplete, but we can already decide if $\mathcal{W}$ is well separated, since we have revealed all the 3-neighbourhoods of vertices in $\mathcal{W}$ (and nothing more). If the set is not well separated, we can stop. On the event that the set is well separated, we continue our construction of the graph, with the aim of showing that $\mathcal{W}$ will contain a good subset of size $\varepsilon n/5$ with high probability.

Let $\mathcal{F}$ be the set of vertices still having $(d + 1)$ half-edges at this point. We call elements of $\mathcal{F}$ *fresh* vertices (they are still unseen by our construction of the graph). Elements of $\mathcal{W}$ will be called *seeds*. Recall that since we consider only the event that $\mathcal{W}$ is well separated, for every seed ρ, the 3-neighbourhood of ρ, rooted at ρ, is isomorphic to $(o, \mathbb{T}_3)$; moreover, every vertex x at distance 3 from ρ has exactly d half-edges attached to it at this point.

Below we will give a procedure to continue electing half-edges and thus continue the construction of the graph; at any point in the construction, we say that a seed ρ is *active* if there are at least 3 half-edges attached to vertices at distance 3 from ρ (as of now, there are $(d + 1)d^3$ such half-edges).

A seed that is not active will be called *quiet*. As of now, every seed is active. We let ρ be any active seed (as of now, this only means that we take $\rho \in \mathcal{W}$), and run the pass described below.

The pass started from the active seed ρ.

Let $\overline{V}$ be the set of vertices at distance 3 or less from ρ and let $\overline{\mathcal{E}}$ be the set of edges with both extremities in $\overline{V}$. Let $\overline{\mathcal{H}}$ be the set of half-edges attached to vertices of $\overline{V}$; by the definition of an active seed, $\overline{\mathcal{H}}$ has at least 3 elements.

✠⟩ Let $v \in \overline{V}$ be at distance strictly less than $\ell + 4$ from ρ as measured in the graph $(\overline{V}, \overline{\mathcal{E}})$, and such that there is a half-edge $h \in \overline{\mathcal{H}}$ attached to it. If such a v does not exist, then declare that the pass is a *success* and stop. Otherwise, pick h' uniformly at random in $\mathcal{H} \setminus \{h\}$, and let v' be the vertex to which it is attached. If $v' \notin \mathcal{F}$, then declare that a *collision* occurs (say that it is a *short collision* if moreover, $v' \in \overline{V}$; that it is a *long collision* otherwise).

- If it is the second time during the pass that a collision occurs, then declare that the pass is a *failure*, do the updates

(3.5) $$\mathcal{E} \leftarrow \mathcal{E} \cup \{h + h'\}, \qquad \mathcal{H} \leftarrow \mathcal{H} \setminus \{h, h'\},$$

  and stop the pass.
- If it is the first time during the pass that a collision occurs, do the updates

(3.6) $$\mathcal{E} \leftarrow \mathcal{E} \cup \{h + h'\}, \qquad \mathcal{H} \leftarrow \mathcal{H} \setminus \{h, h'\}, \qquad \overline{\mathcal{H}} \leftarrow \overline{\mathcal{H}} \setminus \{h, h'\},$$

  and go back to ✠⟩.
- If a collision does not occur (that is, if $v' \in \mathcal{F}$), then do the updates

$$\mathcal{E} \leftarrow \mathcal{E} \cup \{h + h'\}, \quad \mathcal{H} \leftarrow \mathcal{H} \setminus \{h, h'\}, \quad \overline{\mathcal{H}} \leftarrow \overline{\mathcal{H}} \setminus \{h\},$$
$$\overline{V} \leftarrow \overline{V} \cup \{v'\}, \quad \mathcal{F} \leftarrow \mathcal{F} \setminus \{v'\}, \quad \overline{\mathcal{E}} \leftarrow \overline{\mathcal{E}} \cup \{h + h'\},$$

  and finally add the remaining half-edges of v' to $\overline{\mathcal{H}}$. Then go back to ✠⟩.

Observation 3.8. In (3.6), we only need to remove h' from $\overline{\mathcal{H}}$ in case of a short collision; if the collision is long, then h' is not an element of $\overline{\mathcal{H}}$ at this point.

We then iterate this pass, **always starting with an active seed that has not been the starting point of any previous pass**, until every seed is quiet or has already been used as the starting point of a pass.

We now list some important observations concerning the effect of passes.

Observation 3.9 (The effect of a successful pass). If the pass is a success, we define $G'_\rho = (\overline{V}, \overline{\mathcal{E}})$. This subgraph of G_n is a tree satisfying:

(a) ρ is in the vertex set,
(b) there exists $x \in \overline{V}$ such that $\mathsf{dist}_{G'_\rho}(\rho, x) = 4$ and (x, G'_ρ) embeds $(o, \mathbb{T}_\ell)$.

Let us verify (b); to this end, we introduce some notation. Given a vertex $x \in \overline{V}$, we let $\overline{V}_x$ be the set of vertices $y \in \overline{V}$ such that the unique path in G'_ρ from y to ρ contains x, then let $\overline{\mathcal{E}}_x$ be the set of edges of $\overline{\mathcal{E}}$ with both extremities in $\overline{V}_x$, and finally let $G'_{\rho,x} = (\overline{V}_x, \overline{\mathcal{E}}_x)$. Now, if there were no collisions in the pass, then there are at least three distinct vertices x_1, x_2, x_3 such that $\mathsf{dist}_{G'_\rho}(x_i, \rho) = 4$ and (x_i, G'_{ρ,x_i}) is isomorphic to $(o, \mathbb{T}_\ell)$ for each i, so (b) holds. If one collision occurred in the pass, then at most two of these subtrees is compromised by missing edges, so that there still exists some x with $\mathsf{dist}_{G'_\rho}(x, \rho) = 4$ and $(x, G'_{\rho,x})$ isomorphic to $(o, \mathbb{T}_\ell)$.

Apart from satisfying the properties listed above, G'_ρ does not intersect any subgraph $G'_{\hat{\rho}}$ obtained from previous passes.

Observation 3.10 (Many fresh sites). During one pass, at most

$$c_\ell := (d+1)d^3 + (d+1)d^4 + \cdots + (d+1)d^{\ell+4}$$

iterations of the instruction ✠ are performed, so in particular at most c_ℓ vertices are removed from the set $\mathcal{F}$ of fresh vertices. Initially, there are

$$n - |\mathcal{W}|(1 + (d+1) + (d+1)d + (d+1)d^2)$$

fresh vertices. Since we start with $|\mathcal{W}|$ active seed vertices, we can run the pass no more than $|\mathcal{W}|$ times. At any given time, there are thus at least $n - c'_\ell|\mathcal{W}|$ fresh vertices, where

$$c'_\ell := 1 + (d+1) + (d+1)d + \cdots + (d+1)d^{\ell+4}.$$

Observation 3.11 (Many passes). It may occur, due to half-edges being removed in collisions, that a seed becomes inactive before it is used in a pass (let us say that it is then *ruined*). Recall that for a seed ρ,

- there are initially $(d+1)d^3$ half-edges attached to vertices at distance 3 from ρ;
- ρ remains active as long as at least 3 of these half-edges are still present;
- any given pass which starts from some other seed ρ' can cause the removal of at most two of these half-edges.

In conclusion, it takes at least $\frac{(d+1)d^3-3}{2}$ passes started from other seeds to ruin the seed ρ. Taking this into account we see that, if we run the pass t times, then at most

$$t + t \cdot \frac{2}{(d+1)d^3 - 3}$$

seeds become inactive (namely, those that are being explored by each pass and those that are ruined by collisions). In other words, since we start with $|\mathcal{W}| = \varepsilon n$ seeds, we can run the pass at least

$$\left(1 + \frac{2}{(d+1)d^3 - 3}\right)^{-1} |\mathcal{W}| \geqslant \frac{4}{5}\varepsilon n$$

times without running out of active seeds (since $d \geqslant 2$).

Let us now estimate the probability that a pass is a failure. By observation 3.10, the probability to find a collision while running an iteration of instruction ✠ (that is, at each moment in the pass in which we pick a half-edge uniformly at random) is bounded above by

$$\frac{dc'_\ell|\mathcal{W}|}{(d+1)(n - c'_\ell|\mathcal{W}|)} \leqslant \frac{c'_\ell\varepsilon}{1 - c'_\ell\varepsilon}.$$

Hence, the probability that a double collision is found during a pass (or in other words: the pass fails due to the occurrence of two collisions) is smaller than

$$\mathbb{P}\left[\mathsf{Bin}\left(c_\ell, \frac{c'_\ell\varepsilon}{1 - c'_\ell\varepsilon}\right) \geqslant 2\right],$$

where we used the fact that during one pass, we iterate instruction ✠ at most c_ℓ times. Hence, we can choose ε sufficiently small (depending on d and ℓ) so that

(3.7)   the probability for a pass to fail is less than $c'\varepsilon^2$,

where $c' = (c_\ell c'_\ell)^2$.

The probability that at least $\frac{3}{5}\varepsilon n$ passes fail due to double collisions is thus bounded by

$$\mathbb{P}\left[\mathsf{Bin}\left(|\mathcal{W}|, c'\varepsilon^2\right) \geqslant \frac{3}{5}\varepsilon n\right] = \mathbb{P}\left[\mathsf{Bin}\left(\varepsilon n, c'\varepsilon^2\right) \geqslant \frac{3}{5}\varepsilon n\right],$$

using (3.7) and the fact that we run at most $|\mathcal{W}|$ passes. Applying Lemma 3.1 with $m = \varepsilon n = |\mathcal{W}|$, $p = c'\varepsilon^2$ and $p + \delta = \frac{3}{5}$, we see that this probability is bounded by

$$\exp\left\{-|\mathcal{W}|\left[\frac{3}{5}\log\left(\frac{3}{5c'\varepsilon^2}\right) + \frac{2}{5}\log\left(\frac{2}{5(1 - c'\varepsilon^2)}\right)\right]\right\}.$$

Hence, the probability above is bounded by $(c\varepsilon^{\frac{6}{5}})^{|\mathcal{W}|}$ for some constant c independent of ε.

To sum up, we have shown that in our procedure, we

- run at least $\frac{4}{5}\varepsilon n$ passes;
- have at most $\frac{3}{5}\varepsilon n$ failed passes with probability larger than $1 - (c\varepsilon^{\frac{6}{5}})^{|\mathcal{W}|}$.

This is exactly what we need to conclude, as was explained in Observation 3.9. $\square$

Corollary 3.12. *For every $\ell \in \mathbb{N}$ and $\varepsilon > 0$ sufficiently small,*

$$\mathbb{P}\left[\begin{array}{l} \textit{from every well-separated set } W \subseteq V_n \textit{ of size } \varepsilon n, \\ \textit{one can extract an } \ell\textit{-regenerative subset of size } \varepsilon n/5 \end{array}\right] \xrightarrow{n\to\infty} 1.$$

Proof. Take c, ε corresponding to ℓ as in Proposition 3.6. For large enough n, the probability for a given well-separated set $W \subseteq V_n$ of size εn not to contain any ℓ-regenerative subset of size $\varepsilon n/5$ is smaller than $(c\varepsilon^{\frac{6}{5}})^{\varepsilon n}$. We prove the corollary by a union bound, counting the total number of sets $W \subseteq V_n$ of size εn. This number is

$$\binom{n}{\varepsilon n} \leqslant \frac{n^{\varepsilon n}}{(\varepsilon n)!}.$$

Using the fact that $\log n! \geqslant n \log n - n$ (which can be proved by induction on n), we see that

$$\log\binom{n}{\varepsilon n} \leqslant \varepsilon n \log n - \varepsilon n \log(\varepsilon n) + \varepsilon n = \varepsilon n\left(1 + \log\left(\frac{1}{\varepsilon}\right)\right).$$

By reducing the value of $\varepsilon > 0$ if necessary, we can ensure that

$$(c\varepsilon^{\frac{6}{5}})^{\varepsilon n} \exp\left[\varepsilon n\left(1 + \log\left(\frac{1}{\varepsilon}\right)\right)\right]$$

tends to 0 as n tends to infinity, and this proves the result. $\square$

Proof of Theorem 3.4. The number of cycles in G_n of size bounded above by 3 remains tight as n tends to infinity (in order to see this, it suffices to check that the expectation of the number of such cycles remains bounded). In particular, with probability tending to 1, there are less than $\sqrt{n}$ cycles of size bounded by 3. We assume from now on that the event that there are no more than $\sqrt{n}$ cycles of size bounded by 3 is realized.

For any site $v \in V$, the size of its 6-neighbourhood is no more than

$$1 + (d + 1) + (d + 1)d + \cdots + (d + 1)d^5 \leqslant 4d^6.$$

We observe that if $m \in \mathbb{N}$ and U is a set of m vertices, we can extract from U a subset of $m/(4d^6)$ vertices whose 3-neighbourhoods are pairwise disjoint. To see this, let x_1 be an arbitrary point of U. Then, the set U' obtained by removing from U all points that belong to the 6-neighbourhood of x_1 has size at least $m - 4d^6$. Moreover, the 3-neighbourhood of any point of U' is disjoint from the 3-neighbourhood of x_1. We then let x_2 be an arbitrary point of U', and continue this procedure until $x_{m/(4d^6)}$ is defined.

Now assume that W is a set of εn vertices. By the above, we can find $W' \subseteq W$ with at least $\varepsilon n/(4d^6)$ vertices so that the 3-neighbourhoods of vertices in W' are all disjoint. As there are no more than $\sqrt{n}$ cycles of size bounded above by 3, assuming n is large enough that $\sqrt{n} < \varepsilon n/(8d^6)$, we can further obtain $W'' \subseteq W'$

with $|W''| > \varepsilon n/(8d^6)$ vertices so that W'' is well-separated. The desired conclusion then follows immediately from Corollary 3.12. $\qquad\square$

4. SUBCRITICAL REGIME

We start from the following consequence of well-known estimates for the subcritical contact process on trees.

Proposition 4.1. *For any $\lambda < \lambda_1(\mathbb{T})$, there exists $C > 0$ such that*
$$\lim_{n\to\infty} \sup_{A\subseteq\mathbb{T}:|A|=n} P_{\mathbb{T},\lambda}\left[\tau_{\mathbb{T}}^A > C\log n\right] = 0.$$

Proof. The proof will depend on the fact that, for any $\lambda < \lambda_1(\mathbb{T})$, there exist $c_0, C_0 > 0$ such that

(4.1) $$\mathbb{E}_{\mathbb{T},\lambda}\left[|\xi_t^o|\right] \leqslant C_0 e^{-c_0 t}, \quad t \geqslant 0.$$

This is proved as follows. For $\lambda < \lambda_1(\mathbb{T})$, (2.6) and (2.7) give $\beta(\lambda) < \frac{1}{d}$. Together with (2.5), this shows that $\phi_\lambda(\rho) = 1$ for some $\rho > 1$. Then, using (2.3), we get $\phi_\lambda(1) < 1$, and then (2.2) gives the desired equation (4.1).

Let us now show how (4.1) completes the proof of our proposition. Noting that $|\xi_t^A| = |\cup_{x\in A} \xi_t^x| \leqslant \sum_{x\in A} |\xi_t^x|$, we have
$$P_{\mathbb{T},\lambda}\left[|\xi_t^A| \neq \varnothing\right] \leqslant E_{\mathbb{T},\lambda}\left[|\xi_t^A|\right] \leqslant C_0|A|e^{-c_0 t}.$$

The proof is completed by taking $C = 2/c_0$. $\qquad\square$

Part (a) of Theorem 1.2 is a consequence of the above proposition and the following result:

Lemma 4.2. *For any finite graph $G = (V, E)$ with degree bounded by $d + 1$, $A \subseteq V$ and $t > 0$,*
$$P_{G,\lambda}\left[\tau_G^A > t\right] \leqslant \sup_{B\subseteq\mathbb{T}:|B|=|A|} P_{\mathbb{T},\lambda}\left[\tau_{\mathbb{T}}^B > t\right].$$

Proof. Since the contact process is unaffected by the presence of edges that start and end at the same vertex (loops), we assume that G has none.

We will now recall the concept of universal covering of the graph G; we will construct from G a new graph $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ with certain desirable properties. We start fixing a reference vertex $x \in V$. We say that a sequence $\gamma = (\vec{e}_1, \ldots, \vec{e}_n)$ of oriented edges of $\vec{E}$ is a *non-backtracking path from x* if $v_0(\vec{e}_1) = x$ and, for $1 \leqslant i < n$, $v_1(\vec{e}_i) = v_0(\vec{e}_{i+1})$ and $u(\vec{e}_i) \neq u(\vec{e}_{i+1})$ (recall that, for an oriented edge $\vec{e}$, $v_0(\vec{e})$, $v_1(\vec{e})$ and $u(\vec{e})$ respectively denote the starting vertex, ending vertex and undirected edge associated to $\vec{e}$). Let $\mathcal{V}$ be the set of all non-backtracking paths from x, including an empty path which we denote by o. For any $\gamma, \gamma' \in \mathcal{V}$ with $\gamma = (\vec{e}_1, \ldots, \vec{e}_n)$ and $\gamma' = (\vec{e}_1, \ldots, \vec{e}_n, \vec{e}_{n+1})$, we connect γ and γ' by an edge; this defines the edge set $\mathcal{E}$ of $\mathcal{T}$. Finally, put $\psi(o) = x$ and, for $\gamma = (\vec{e}_1, \ldots, \vec{e}_n)$, put $\psi(\gamma) = v_1(\vec{e}_n)$, the ending vertex of the path γ. It is now easy to check that $\mathcal{T}$ and ψ satisfy the properties:

(a) $\mathcal{T}$ is a tree with degree bounded by $d + 1$;
(b) for every $\gamma \in \mathcal{V}$, letting $x = \psi(\gamma) \in V$, we have

for all $y \in V$, $\quad |\{$edges of E with extremities x and $y\}|$
$$= |\{\gamma' \in \mathcal{V} : \gamma \text{ and } \gamma' \text{ are neighbours in } \mathcal{T},\ \psi(\gamma') = y\}|.$$

(in case x is connected to each of its neighbours by a single edge, this just says that ψ maps the neighbourhood of γ bijectively to the neighbourhood of x).

For $v \in V$, the set $\psi^{-1}(v)$ is called the *fiber* of v. Define the set of configurations of $\{0,1\}^{\mathcal{V}}$ that have at most one particle per fiber,

$$\Omega_{\mathcal{T}} = \left\{ \zeta \in \{0,1\}^{\mathcal{V}} : \sum_{\gamma \in \psi^{-1}(v)} \zeta(\gamma) \in \{0,1\} \text{ for all } v \in V \right\}.$$

Define the projection $\pi : \Omega_{\mathcal{T}} \to \{0,1\}^V$ by $[\pi(\zeta)](v) = \sum_{\gamma \in \psi^{-1}(v)} \zeta(\gamma)$ for $v \in V$. We abuse notation and, for a set $B \subseteq \mathcal{V}$, we write $B \in \Omega_{\mathcal{T}}$ if $\mathbf{1}_B \in \Omega_{\mathcal{T}}$.

Given $B \in \Omega_{\mathcal{T}}$, we define the *constrained contact process on* $\mathcal{T}$, $(\zeta_t^B)_{t \geqslant 0}$, as follows. We set $\zeta_0^B = B$ and let ζ evolve as a contact process on $\mathcal{T}$ with the restriction that we suppress every transition which would result in a configuration not in $\Omega_{\mathcal{T}}$, that is, births on vertices belonging to fibers already containing infected vertices. Formally, (ζ_t) has generator

$$\mathcal{L}f(\zeta) = \sum_{\gamma \in \mathcal{V}} \left[f(\zeta^{0 \to \gamma}) - f(\zeta) \right] + \lambda \cdot \sum_{\substack{\{\gamma,\gamma'\} \in \mathcal{E}: \\ \zeta(\gamma)=1}} \left[\mathbf{1}_{\{\zeta^{1 \to \gamma'} \in \Omega_{\mathcal{T}}\}} \cdot \left(f(\zeta^{1 \to \gamma'}) - f(\zeta) \right) \right],$$

where $\zeta^{i \to \gamma'}$ is the configuration obtained by modifying ζ so that γ' is set to state i. Noting that G is finite, and hence ζ has at most $|V|$ infected vertices at any given time, it is not hard to see that the above generator indeed gives rise to a Feller process on $\Omega_{\mathcal{T}}^{[0,\infty)}$. In fact, ζ can be constructed from a Harris system on $\mathcal{T}$, with the extra care of ignoring any transmission mark which would cause a fiber to become doubly occupied. Also, using property (b) of $\mathcal{T}$ and ψ stated above, it is easy to show that $(\pi(\zeta_t^B))_{t \geqslant 0}$ has the same distribution as the contact process on G started from $\pi(B)$.

Now, if we start from $A \subseteq V$, we can choose an arbitrary $B \in \Omega_{\mathcal{T}}$ such that $\pi(B) = A$ and conclude that τ_G^A has the same distribution as $\inf\{t : \zeta_t^B = \varnothing\}$. By seeing $\mathcal{T}$ (and hence B) as a subset of $\mathbb{T}$, we have that ζ^B is stochastically dominated by the contact process on $\mathbb{T}$ started from B infected, and hence $\inf\{t : \zeta_t^B = \varnothing\}$ is stochastically dominated by $\tau_{\mathbb{T}}^B$. This completes the proof. $\qquad \square$

Let G be a graph and x a vertex of G. Given the contact process on G started from a single infection at x, $(\xi_t^x)_{t \geqslant 0}$, define

$$\kappa_G^x = \sup\{\text{dist}(x,y) : y \in \xi_t^x \text{ for some } t \geqslant 0\}.$$

Lemma 4.3. *For any finite graph G with degree bounded by $d+1$, any vertex x of G and $k > 0$,*

$$\mathbb{P}\left[\kappa_G^x > k\right] \leqslant \mathbb{P}\left[\kappa_{\mathbb{T}}^o > k\right].$$

Proof. Repeating the construction in the previous lemma, we note that, if $\psi(\gamma) = x$ and $\psi(\gamma') = y$, then $\text{dist}_{\mathcal{T}}(\gamma,\gamma') \geqslant \text{dist}_G(x,y)$, and we then see that κ_G^x is stochastically dominated by $\sup\{\text{dist}_{\mathcal{T}}(\gamma,\gamma') : \gamma' \in \zeta_t^{\{\gamma\}} \text{ for some } t \geqslant 0\}$. The latter is in turn stochastically dominated by $\kappa_{\mathbb{T}}^o$, completing the proof. $\qquad \square$

Proof of Theorem 1.4. In order to show that the contact process on G dies out, it suffices to show, for any vertex x, that $\mathbb{P}\left[\tau_G^x < \infty\right] = 1$. Denote by $B_G(x,r)$ the subgraph of G induced by the set of vertices at graph distance at most r from x. Then,

$$\mathbb{P}\left[\tau_G^x < \infty\right] = \lim_{r \to \infty} \mathbb{P}\left[\kappa_G^x \leqslant r\right] = \lim_{r \to \infty} \mathbb{P}\left[\kappa_{B_G(x,r+1)}^x \leqslant r\right] \geqslant \lim_{r \to \infty} \mathbb{P}\left[\kappa_{\mathbb{T}}^o \leqslant r\right] = 1.$$

$\qquad \square$

References

[Be] I. Benjamini. *Coarse geometry and randomness – École d'Été de Probabilités de Saint-Flour 2011.* Lecture Notes in Mathematics **2100**, Springer (2013).

[BNP11] I. Benjamini, A. Nachmias, Y. Peres. *Is the critical percolation probability local?* Probab. Theory Related Fields **149** (1-2), 261–269 (2011).

[BS01] I. Benjamini, O. Schramm. *Recurrence of distributional limits of finite planar graphs.* Electron. J. Probab. **6** (23), 13 pp. (2001).

[CGOV84] M. Cassandro, A. Galves, E. Olivieri, M. Vares. *Metastable behavior of stochastic dynamics: a pathwise approach.* J. Statist. Phys. **35** (5-6), 603-634 (1984).

[CD09] S. Chatterjee, R. Durrett. *Contact processes on random graphs with power law degree distributions have critical value 0.* Ann. Probab. **37**, 2332 - 2356 (2009).

[Ch94] J.W. Chen. *The contact process on a finite system in higher dimensions.* Chinese J. Contemp. Math. 15, 13-20 (1994).

[CMMV13] M. Cranston, T. Mountford, J.C. Mourrat, D. Valesin, *The contact process on finite trees revisited.* ALEA Vol XI, 385–408 (2014).

[DL88] R. Durrett, X. Liu. *The contact process on a finite set.* Ann. Probab. **16**, 1158-1173 (1988).

[DS88] R. Durrett, R.H. Schonmann. *The contact process on a finite set II.* Ann. Probab. **16**, 1570-1583 (1988).

[Li99] T. Liggett, *Stochastic interacting systems: contact, voter and exclusion processes.* Grundlehren der mathematischen Wissenschaften **324**, Springer (1999).

[MT13] S. Martineau, V. Tassion. *Locality of percolation for abelian Cayley graphs.* Preprint, arXiv:1312.1946 (2013).

[MS15] L. Ménard, A. Singh. *Percolation by cumulative merging and phase transition for the contact process on random graphs.* Preprint, arXiv:1502.06982 (2015).

[Mo93] T. Mountford. *A metastable result for the finite multidimensional contact process.* Canad. Math. Bull. **36** (2), 216-226 (1993).

[Mo99] T. Mountford. *Existence of a constant for finite system extinction.* J. Statist. Phys. **96** (5-6), 1331-1341 (1999).

[MMVY12] T. Mountford, J.C. Mourrat, D. Valesin, Q. Yao. *Exponential extinction time of the contact process on finite graphs.* Preprint, arXiv:1203.2972 (2012).

[MVY13] T. Mountford, D. Valesin, Q. Yao. *Metastable densities for the contact process on power law random graphs.* Electronic Journal of Probability, Vol 18 1-36 (2013).

[vdH] vd Hofstad, R. *Random graphs and complex networks.* Available at http://www.win.tue.nl/∼rhofstad/

[SS98] M. Salzano, R. Schonmann. *A new proof that for the contact process on homogeneous trees, local survival implies complete convergence.* Ann. Probab. **26**, 1251-1258 (1998).

[Sc85] R. Schonmann. *Metastability for the contact process.* J. Statist. Phys. **41** (3-4), 445-464 (1985).

[SXZ14] H. Song, K.-N. Xiang, S.-C.-H. Zhu. *Locality of percolation critical probabilities: uniformly nonamenable case.* Preprint, arXiv:1410.2453 (2014).

[St01] A. Stacey. *The contact process on finite homogeneous trees.* Probab. Theory Related Fields. **121** (4), 551-576 (2001).

(Jean-Christophe Mourrat) ENS Lyon, CNRS, 46 allée d'Italie, 69007 Lyon, France

(Daniel Valesin) University of Groningen, Nijemborgh 9, 9747AG Groningen, The Netherlands