Dependent random graphs and spatial epidemics

J. van den Berg[1], Geoffrey R. Grimmett[2], Rinaldo B. Schinazi[3]

Abstract. We extend certain exponential decay results of subcritical percolation to a class of locally dependent random graphs, introduced by Kuulasmaa as models for spatial epidemics on $\mathbf{Z}^d$. In these models, infected individuals eventually die (are removed) and are not replaced. We combine these results with certain continuity and re-scaling arguments in order to improve our knowledge of the phase diagram of a modified epidemic model in which new susceptibles are born at some positive rate. In particular, we show that, throughout an intermediate phase where the infection rate lies between a certain two critical values, no coexistence is possible for sufficiently small positive values of the recovery rate. This result provides a converse to results of Durrett–Neuhauser and Andjel–Schinazi. We show also that such an intermediate phase indeed exists for every $d \geq 1$ (i.e., that the two critical values mentioned above are distinct). An important technique is the general version of the BK inequality for disjoint occurrence, proved in 1994 by Reimer.

1. Introduction

The work reported in this paper originated as an attempt to understand better the phase diagram of the following epidemic model (studied by Durrett and Neuhauser (1991)). At any time t, a vertex $x \in \mathbf{Z}^d$ is in a state $\xi_t(x)$ which takes values in the set $\{0, 1, 2\}$, these states having the following interpretations:

0 means *dead* or removed,
1 means *susceptible* or healthy,
2 means *ill*.

The transition rates of $\xi_t(x)$, given the entire configuration $\xi_t = \big(\xi_t(y) : y \in \mathbf{Z}^d\big)$, are defined to be

---

AMS 1991 subject classification: 60K35.

Keywords: dependent percolation, interacting particle system, spatial epidemic, critical value, percolation.

[1]CWI, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands.

[2]Statistical Laboratory, 16 Mill Lane, Cambridge CB2 1SB, UK. Partially funded by the European Union under contracts CHRX-CT93-0411 and ERBFMRX-CT96-0075A, and by the Engineering and Physical Sciences Research Council under contract GR/L15426.

[3]Mathematics Department, University of Colorado, Colorado Springs CO 80933, USA. Partially funded by the National Science Foundation.

March 29, 1997

$$\begin{array}{ll} 1 \to 2 & \text{at rate } \alpha n_2(x), \\ 2 \to 0 & \text{at rate } \delta \text{ (taken to equal 1)}, \\ 0 \to 1 & \text{at rate } \beta, \end{array}$$

where $0 < \alpha < \infty$, $0 \le \beta \le \infty$, and $n_2(x)$ is the number of ill neighbours of x in ξ_t. We allow the value $\beta = \infty$, in which case the state of any vertex passes instantaneously through the value 0; i.e., recovery is immediate. We omit a formal definition of the Markov process ξ on the state space $\{0,1,2\}^{\mathbf{Z}^d}$, referring the reader to Chapter 1 of Liggett (1985).

There is some difficulty in achieving a useful definition of critical values for the process ξ when $0 < \beta < \infty$, owing to a lack of (provable) monotonicity. The following route is appealing but problematic. Writing $P_{\alpha,\beta}$ for the appropriate probability measure when the initial distribution is given by

$$\xi_0(u) = \begin{cases} 2 & \text{if } u = 0 \\ 1 & \text{if } u \neq 0, \end{cases}$$

we let

$$\theta(\alpha,\beta) = P_{\alpha,\beta}\big(|C| = \infty\big), \tag{1.1}$$

where C is the set of vertices which are in state 2 at some time. We may think of $\theta(\alpha,\beta)$ as the probability that the infection continues forever. We let

$$\alpha_c(\beta) = \sup\{\alpha : \theta(\alpha,\beta) = 0\}. \tag{1.2}$$

It may be proved that $\theta(\alpha,\beta)$ is a non-decreasing function of α when $\beta \in \{0,\infty\}$, whence (1.2) provides a useful definition of the critical point in these two extreme cases. Unfortunately, we do not know whether or not $\theta(\alpha,\beta)$ is monotonic in α when $0 < \beta < \infty$, and therefore we have insufficient information to describe the associated phase transitions.

When $\beta = \infty$, this model is (effectively) the well known contact process with possible states 1 and 2, and $\alpha_c(\infty)$ is the critical infection rate. For the basic properties of the contact model, see Liggett (1985, Chap. 6) and Bezuidenhout and Grimmett (1990).

The other extreme case $\beta = 0$ has been considered by Kuulasmaa (1982). This is an 'epidemic without recovery'. Kuulasmaa showed that $\theta(\alpha, 0)$ is non-decreasing in α, and that $0 < \alpha_c(0) < \infty$ when $d \ge 2$; it is easy to see that $\alpha_c(0) = \infty$ when $d = 1$.

It may be proved that $\alpha_c(\infty) \le \alpha_c(0)$, and one of our results is the strict inequality $\alpha_c(\infty) < \alpha_c(0)$, valid in any dimension. We prove this at Corollary 4.3 below, using a coupling of $P_{\alpha,0}$ and $P_{\alpha,\infty}$.

It is apparently difficult to study the above epidemic model for general α, β satisfying $0 < \alpha, \beta < \infty$, owing to the possibility of non-monotonicity. Therefore, we describe here certain approaches designed to understand the relevance of the critical values $\alpha_c(0)$, $\alpha_c(\infty)$. In particular, we are interested in whether the parameter-pair

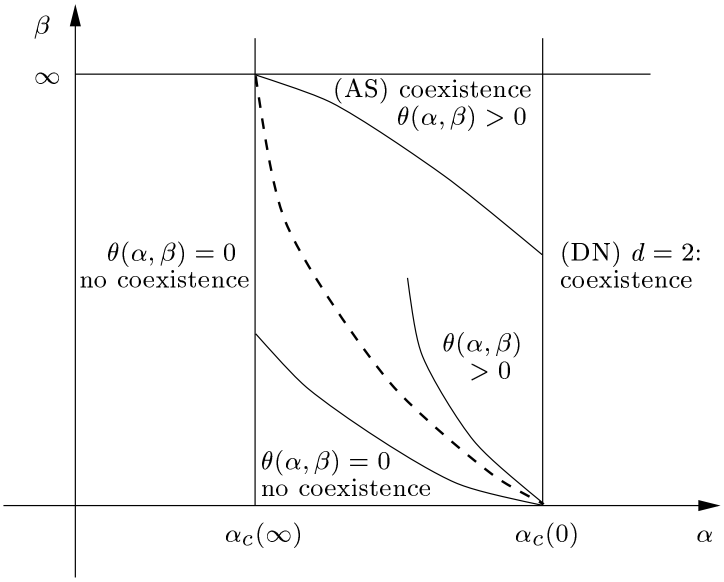

Fig. 1. A sketch of the phase diagram of the epidemic model with parameters α and β. The letters AS refer to Andjel–Schinazi (1996), and DN to Durrett–Neuhauser (1991). The dashed line indicates a critical curve separating the two phases characterised by 'coexistence' and 'no coexistence', but the existence of such a curve has not been proved. This sketch is valid for all dimensions $d \geq 1$, except where noted otherwise; we point out that $\alpha_c(0) = \infty$ when $d = 1$.

$(\alpha, \beta) = (\alpha_c(0), 0)$ lies on a (hypothetical) critical surface of the general model. Since we do not know whether there exists a unique critical curve, we restrict ourselves to a lesser question, namely whether arbitrarily small perturbations of (α, β), in the neighbourhood of $(\alpha_c(0), 0)$, may take the process into either of two phases, these phases being characterised by the occurrence or not of 'coexistence'. The property of coexistence has been studied fairly widely in the theory of infinite particle systems. In the present context, we say that 'coexistence occurs' if there exists an invariant measure ν for the process ξ such that ν is concentrated on configurations which include at least one vertex in each of the states 0, 1, and 2.

In Section 2, we introduce a class of *locally dependent* percolation models in which (directed) edges of the lattice have random states with a certain dependence structure. We present exponential-decay results for such systems, and note that these results generalise the corresponding statements for *independent* percolation. See Theorems 2.1 and 2.2 below; the corresponding results for independent percolation may be found in Grimmett (1989). In Section 3, we apply these result to a locally dependent random graph arising from the special case $\beta = 0$ of the epidemic model described above. In this way, we obtain exponential-decay theorems for epidemics without recovery.

In Section 4, we combine the above observations with certain continuity and re-scaling arguments in order to obtain a new result for the epidemic model with $0 < \beta < \infty$, namely the following. Durrett and Neuhauser (1991) showed that, if $\alpha > \alpha_c(0)$ and $\beta > 0$, coexistence occurs when $d = 2$. On the other hand, Andjel and Schinazi (1996) have proved that $\alpha > \alpha_c(\infty)$ if and only if there exists β such that coexistence occurs.

(In this case, coexistence occurs for all sufficiently large values of β.) One of our main current results is that, if $\alpha_c(\infty) < \alpha < \alpha_c(0)$, then coexistence is not possible for β less than a certain non-trivial threshold (whose value depends on α); see Theorem 4.4 below. This identifies $\alpha_c(0)$ as a genuine critical value for the general epidemic model with recovery, thereby proving the existence of two distinct critical values, the other one being $\alpha_c(\infty)$. Proofs of our results are given in Sections 5–8. See Figure 1 for a sketch of the phase diagram of the model.

We finish this section with a little notation. As usual, $\mathbf{L}^d$ denotes the cubic lattice in d dimensions, where $d \geq 1$. For a vertex $u = (u_1, u_2, \ldots, u_d)$, we define

$$|u| = \sum_{i=1}^{d} |u_i|.$$

An edge from u to v in $\mathbf{L}^d$ is denoted by $\langle u, v\rangle$. When it is directed from u to v, we write $[u, v\rangle$. The *surface* of a set A of vertices (i.e., the set of vertices in A which are adjacent to some vertex outside A) is denoted by ∂A. We write $\Lambda_n = [-n, n]^d$, the box of $\mathbf{Z}^d$ having side length $2n$.

2. Exponential decay for locally dependent random graphs

In the usual bond percolation process on $\mathbf{L}^d$, the states of different edges are independent random variables. One of the main results for subcritical percolation is the exponential decay of the connectivity function (Aizenman and Barsky (1987), Menshikov (1986); see also Chapter 3 of Grimmett (1989)). We will state analogous results for a percolation-type model in which certain dependencies are introduced between the states of different edges. Such systems were studied by Kuulasmaa (1982) under the name 'locally dependent random graphs'.

We shall work on directed graphs rather than undirected graphs, and we begin by writing $\vec{\mathbf{L}}^d$ for the directed graph obtained from $\mathbf{L}^d$ by replacing each edge by a pair of oppositely directed edges (i.e., the edge $\langle u, v\rangle$ is replaced by an edge $[u, v\rangle$, together with an edge $[v, u\rangle$.

Next we introduce probabilities, in order to construct a random subgraph of $\vec{\mathbf{L}}^d$. Let $e_1, e_2, \ldots, e_d$ be the positive unit vectors of $\mathbf{Z}^d$, and write $\mathcal{E} = \{\gamma e_i : \gamma = \pm,\ 1 \leq i \leq d\}$. We may think of $\mathcal{E}$ as the set of neighbours of the origin of $\mathbf{L}^d$, or equivalently as the set of possible endpoints of edges directed out of the origin. Let $\mathcal{P}$ be the set of all probability measures on the set $2^{\mathcal{E}}$ of subsets of $\mathcal{E}$. That is, $\mathcal{P}$ contains all functions $\mu : 2^{\mathcal{E}} \to [0, 1]$ satisfying

$$\sum_{A \subseteq \mathcal{E}} \mu(A) = 1.$$

For $\mathcal{A} \subseteq 2^{\mathcal{E}}$, we write

$$\mu(\mathcal{A}) = \sum_{A \in \mathcal{A}} \mu(A),$$

and we call $\mathcal{A}$ *increasing* if

$$A' \in \mathcal{A}, \text{ whenever } A \subseteq A' \text{ and } A \in \mathcal{A}.$$

For $\mu, \nu \in \mathcal{P}$, we say that μ is (*stochastically*) *dominated* by ν (written $\mu \leq \nu$) if $\mu(\mathcal{A}) \leq \nu(\mathcal{A})$ for all increasing subsets $\mathcal{A}$ of $2^{\mathcal{E}}$. If $\mu(\mathcal{A}) < \nu(\mathcal{A})$ for all increasing $\mathcal{A} \subseteq 2^{\mathcal{E}}$ with $\mathcal{A} \neq \emptyset$, $\mathcal{A} \neq 2^{\mathcal{E}}$, then we say that μ is *strictly* dominated by ν.

Let $\mu \in \mathcal{P}$. We now define a random subgraph of $\vec{\mathbf{L}}^d$ as follows. Let $\{N_u : u \in \mathbf{Z}^d\}$ be a collection of independent random subsets of $\mathcal{E}$, each being chosen according to the probability measure μ. Our random subgraph has vertex set $\mathbf{Z}^d$, and its edge set is precisely the set of all directed edges $[u, v\rangle$ for which $v - u \in N_u$. We denote the ensuing directed random graph by Γ, and write P_μ for the associated probability measure.

Let $u_0, u_1, \ldots, u_m$ be the (distinct) vertices of a path in $\mathbf{L}^d$. We call this sequence an *open* (directed) *path* if $[u_i, u_{i+1}\rangle$ lies in Γ for $0 \leq i < m$. We write $u \to v$ if there exists an open path from u to v, and $u \to \infty$ if there exists an infinite open path from u.

The N_u have a natural interpretation in terms of a very general and simple epidemic model in which infected individuals eventually die (and are not replaced) and in which the notion of time is ignored (that is, we record only *which* individuals ever become ill, but not at which times this will happen): when an individual u becomes infected, he in turn sends germs to exactly those individuals in the set $u + N_u$ (consequently, he infects exactly those individuals in this set which have not been infected previously). With this interpretation, the event $\{u \to v\}$ is the set of configurations of the lattice such that, if initially only u is infected, then the infection will eventually reach v.

As a special case, assume for the moment that μ is product measure on $\mathcal{E}$ with density p. It is not difficult to see that $P_\mu(u \to v) = P_p(u \leftrightarrow v)$, where P_p is the measure associated with bond percolation (having density p), and $\{u \leftrightarrow v\}$ is the event that there exists an (undirected) open path from u to v.

Returning to the general case, we define the percolation and connectivity functions in the usual way:

$$\theta(\mu) = P_\mu(0 \to \infty), \quad \tau_\mu(u, v) = P_\mu(u \to v).$$

Next follows our main result for the above locally dependent percolation model.

Theorem 2.1. *Let $\mu, \nu \in \mathcal{P}$ and assume that μ is strictly dominated by ν. If $\theta(\nu) = 0$, then there exists a strictly positive constant $\gamma = \gamma(\mu, \nu)$ such that*

$$\tau_\mu(u, v) \leq e^{-\gamma|v-u|} \quad \text{for } u, v \in \mathbf{Z}^d.$$

As observed above, this conclusion generalises the exponential decay theorem for 'independent' percolation; see Theorem 3.4 of Grimmett (1989). We have presented the theorem for the lattice $\mathbf{L}^d$, but this is not essential.

Our proof of Theorem 2.1 is based upon Menshikov's proof for 'independent' percolation (see Grimmett (1989)). However, there are some key differences, of which the most significant lies in the use of a 'disjoint occurrence' inequality. Whereas the standard BK inequality suffices when the edge-states are independent, the more general version, conjectured by van den Berg and Kesten (1985) and proved by Reimer (1997), is necessary here.

Theorem 2.1 has an implication for the tail of the cluster-size distribution.

Theorem 2.2. *Let $\mu, \nu \in \mathcal{P}$, and assume that μ is strictly dominated by ν. If $\theta(\nu) = 0$, then there exists a strictly positive constant $\xi = \xi(\mu, \nu)$ such that the set $C = \{v : 0 \to v\}$ satisfies*

$$P_\mu(|C| \geq n) \leq e^{-\xi n} \quad \textit{for all } n \geq 0.$$

3. Application to the epidemic model without recovery

Theorems 2.1 and 2.2 may be applied to the epidemic model with $\beta = 0$, i.e., the model in which individuals are removed forever from the system once their period of illness is completed. We take as initial configuration

$$\xi_0(u) = \begin{cases} 2 & \text{if } u = 0 \\ 1 & \text{if } u \neq 0, \end{cases} \tag{3.1}$$

which is to say that illness is introduced at the origin into an ocean of susceptibles. Kuulasmaa (1982) made the interesting observation that the set C of individuals which are ever infected may be represented as the cluster at 0 in a certain locally dependent percolation model. He proved also that $\theta(\alpha, 0) = P_{\alpha,0}(|C| = \infty)$ is strictly positive and non-decreasing in α (> 0), and furthermore that, if $d \geq 2$, then the critical value

$$\alpha_c(0) = \sup\{\alpha : \theta(\alpha, 0) = 0\}$$

satisfies $0 < \alpha_c(0) < \infty$; cf. (1.1). We will show that Theorem 2.1, combined with Kuulasmaa's observations, yields the following result. We recall that $\Lambda_n = [-n, n]^d$.

Theorem 3.1. *If $\alpha < \alpha_c(0)$, there exist $\gamma = \gamma(\alpha)$ and $\delta = \delta(\alpha)$ satisfying $\gamma, \delta > 0$ such that, for all n,*

$$P_{\alpha,0}(C \cap \partial\Lambda_n \neq \emptyset) \leq e^{-\gamma n},$$
$$P_{\alpha,0}(|C| \geq n) \leq e^{-\delta n}.$$

Next we present an exponential decay theorem for the total time of the epidemic. Let T denote the supremum of the set of times at which the state of some vertex changes from 2 to 0.

Theorem 3.2. *There exists $\lambda = \lambda(\alpha)$, satisfying $\lambda > 0$ when $\alpha < \alpha_c(0)$, such that*

$$P_{\alpha,0}(T > t) \leq e^{-\lambda t} \quad \textit{for all } t \geq 0.$$

4. The intermediate phase of epidemics with recovery

We now turn to the epidemic with recovery (or re-birth), i.e., the case $\beta > 0$. As we wrote in the introduction, it is known that coexistence occurs:

$$\begin{gathered}\text{for every positive } \beta \text{ if } \alpha > \alpha_c(0) \text{ and } d = 2,\\ \text{for no } \beta \text{ if } \alpha < \alpha_c(\infty),\\ \text{for sufficiently large } \beta \text{ if } \alpha_c(\infty) < \alpha < \alpha_c(0).\end{gathered}$$

See Figure 1. Our understanding of the 'intermediate region', described in the third case, is far from complete, and this is the region in which we are interested here. Our first step is to show that the intermediate region exists, i.e., that $\alpha_c(\infty) < \alpha_c(0)$.

Consider the model with parameters α, β, and suppose we start as in (3.1) with one vertex ill (in state 2) and all others susceptible (in state 1). Let $\theta(\alpha, \beta)$ be the probability that infinitely many vertices become infected, as in (1.1). It is likely, at least for 'decent' graphs, that $\theta(\alpha, \beta)$ is non-decreasing in α and β, but no proof is known (except for the particular cases $\beta = 0$ and $\beta = \infty$, when monotonicity in α is easy to prove). This absence of monotonicity greatly complicates the proofs, since we have no natural or useful generic coupling of two models with different (general) α, β. Indeed it is the absence of such a coupling which adds much interest to the problem. In place of a coupling, we have the following useful comparison result. The proof may be found in Section 7.

Theorem 4.1. *Let $d \geq 1$. For each $\alpha, \beta > 0$, there exists $\epsilon = \epsilon(\alpha, \beta, d) > 0$ such that*

$$\theta(\alpha, \beta) \geq \theta(\alpha + \epsilon, 0).$$

Moreover, for fixed d, we have that ϵ is bounded away from 0 when: α is bounded away from 0 and from ∞, and β is bounded away from 0.

Corollary 4.2. *Let $d \geq 2$. If $0 < \beta \leq \infty$, there exists α satisfying $\alpha < \alpha_c(0)$ such that $\theta(\alpha, \beta) > 0$.*

Proof. Kuulasmaa (1982) has shown that $\alpha_c(0) < \infty$ if $d \geq 2$. We may choose α sufficiently close to (but strictly smaller than) $\alpha_c(0)$ such that $\epsilon = \epsilon(\alpha, \beta, d)$, given in Theorem 4.1, satisfies $\alpha + \epsilon > \alpha_c(0)$. Then $\theta(\alpha, \beta) > 0$, which yields the result. □

Corollary 4.3 *For every $d \geq 1$, we have that $\alpha_c(\infty) < \alpha_c(0)$.*

Proof. This is well-known when $d = 1$, in which case $\alpha_c(\infty) < \infty$ and $\alpha_c(0) = \infty$. When $d \geq 2$, it follows immediately from the case $\beta = \infty$ of Corollary 4.2. □

Remarks.

(i) Corollary 4.3 may be proved directly using the method of Menshikov (1986) and Aizenman–Grimmett (1991). It is not straightforward to achieve this, but the complications may be overcome. The basic idea is to find a family of processes indexed by a parameter γ taking values in $[0,1]$, such that (a) the case $\gamma=0$ (resp. $\gamma=1$) corresponds to $\beta=0$ (resp. $\beta=\infty$), and (b) the processes may be coupled in such a way that they are monotonic in γ. We do not know how to prove Corollary 4.2 (which is interesting in its own right) by that method.

(ii) It is somewhat annoying that we have not been able to prove the following natural 'counterpart' of Corollary 4.2: for every finite $\beta>0$, there exists an α satisfying $\alpha>\alpha_c(\infty)$ such that $\theta(\alpha,\beta)=0$. This problem is related to the phase diagram in Figure 1, and particularly to the behaviour of the process in a neighbourhood of the point $(\alpha_c(\infty),\infty)$.

(iii) Another open question is whether we can, in Corollary 4.2, replace the statement $\theta(\alpha,\beta)>0$ by the statement 'coexistence occurs'. More generally: is survival (in the sense that $\theta(\alpha,\beta)>0$) equivalent to coexistence? (We would allow some 'occasional' cases when this is false, such as when $\beta=0$.) The paper by Durrett and Neuhauser (1991) gives a partial answer, and indicates that a complete answer may be far from obvious.

We are now ready to state a converse to the results mentioned at the beginning of this section. Recall that $\alpha_c(0)=\infty$ when $d=1$.

Theorem 4.4. *Let $d\ge 1$. If $\alpha\in(\alpha_c(\infty),\alpha_c(0))$, there exists $\beta_c(\alpha)>0$ such that, when $\beta<\beta_c(\alpha)$, then $\theta(\alpha,\beta)=0$ and coexistence does not occur. More precisely, under these conditions, for any finite set A of vertices there exists, with probability one, a finite (random) time T_A after which there will never appear a 2 within A. Moreover, the quantities T_A may be chosen uniformly in the initial configuration ξ_0.*

The proof of Theorem 4.4 is based on a block construction together with Theorems 3.1 and 3.2. It is presented in Section 8, and is related to a block argument of Durrett and Schinazi (1993).

5. Proofs of Theorems 2.1 and 2.2

We first present a version of the generalised BK inequality. Let $X_1,X_2,\dots,X_n$ be independent random variables each taking values in some finite set S. Let A and B be events which are defined in terms of the X_i. We say that A and B are *perpendicular* (written $A\perp B$) if they are defined in terms of disjoint subcollections of the X_i; that is to say, if there exist $J,K\subseteq\{1,2,\dots,n\}$ with $K\cap L=\emptyset$ such that the indicator function 1_A of A (resp. 1_B of B) is a function of $(X_i:i\in J)$ (resp. $(X_i:i\in K)$) alone. The following theorem is equivalent to the general BK inequality proved by Reimer (1997) (see also van den Berg (1997)).

Theorem 5.1. *Let L and M be finite sets, and let A_l and B_m be events defined in terms of the X_i, for $l \in L$ and $m \in M$. Then*

$$P\left(\bigcup_{l,m:A_l \perp B_m} \{A_l \cap B_m\}\right) \le P\left(\bigcup_l A_l\right) P\left(\bigcup_m B_m\right).$$

This inequality will be applied to a locally dependent random graph in the following setting. Let V, W, V', W' be sets of vertices of a finite region Λ in $\mathbf{Z}^d$. We write $V \to W$ for the event that there exists a directed open path (of Λ) joining some vertex of V to some vertex of W. We say that two directed paths having vertex sequences $x_0, x_1, \dots, x_m$ and $y_0, y_1, \dots, y_n$ are *vertex-disjoint* if $x_i \neq y_j$ for all pairs (i, j) except possibly the pair $(i, j) = (m, n)$. That is, we do not require that the final vertices of the two paths be distinct. We write $\{V \to W\} \circ \{V' \to W'\}$ for the event that there are two vertex-disjoint directed open paths, one of which joins a vertex of V to a vertex of W, and the other of which joins a vertex of V' to a vertex of W'.

Let L be an index set for the set of all directed paths (in Λ) of $\vec{\mathbf{L}}^d$ from V to W (and similarly M indexes such paths from V' to W'). For $l \in L$ (resp. $m \in M$) let A_l (resp. B_m) be the event that the directed path indexed by l (resp. m) lies in Γ. It follows by Theorem 5.1 that

$$P_\mu\Big(\{V \to W\} \circ \{V' \to W'\}\Big) \le P_\mu(V \to W) P_\mu(V' \to W'). \tag{5.1}$$

Note here that the symbol $\circ$ requires *vertex-disjoint* paths, whereas in bond percolation it normally requires *edge-disjoint* paths.

Proof of Theorem 2.1. We introduce an auxiliary parameter p as follows. Let $\nu \in \mathcal{P}$, and let Γ be a graph sampled according to the measure P_ν. We colour each vertex u of $\mathbf{Z}^d$ *black* with probability p, otherwise we colour it *white*. Vertex colours are independent of one another and of the graph Γ. From Γ we obtain another graph Γ_p according to the following rule. The graph Γ_p has vertex set $\mathbf{Z}^d$, and the directed edge $[u, v\rangle$ lies in Γ_p if and only if: $[u, v\rangle$ lies in Γ and u is black. Thus, Γ_p is obtained from Γ by deleting all edges emanating from white vertices. It is clear that Γ_p is a locally dependent percolation model distributed according to the measure P_{ν_p}, where $\nu_p \in \mathcal{P}$ is given by

$$\text{for } A \in \mathcal{E}, \quad \nu_p(A) = \begin{cases} p\nu(A) & \text{if } A \neq \emptyset \\ 1 - p(1 - \nu(\emptyset)) & \text{if } A = \emptyset. \end{cases}$$

Let B be an increasing event, i.e., a set of directed subgraphs of $\vec{\mathbf{L}}^d$ which is closed under the operation of increasing the edge-set. A vertex u of $\mathbf{Z}^d$ is called *pivotal* for B if B occurs when u is black, but B does not occur when u is white. We write $N(B)$ for the (random) number of pivotal vertices for B. The usual Russo formula is easily adapted to obtain that, if B is increasing and finite-dimensional, then

$$\begin{aligned} \frac{d}{dp} P_{\nu_p}(B) &= \frac{d}{dp} E\big(P_{\nu_p}(B \mid \Gamma)\big) = E\Big(E\big(N(B) \mid \Gamma\big)\Big) \\ &= \frac{1}{p} E\Big(E\big(N(B) 1_B \mid \Gamma\big)\Big) = \frac{1}{p} E\big(N(B) \mid B\big) P_{\nu_p}(B), \end{aligned}$$

where E is the expectation operator of P_{ν_p}. Therefore

$$\frac{d}{dp}\{\log P_{\nu_p}(B)\} = \frac{1}{p}E\big(N(B) \mid B\big). \tag{5.2}$$

We shall use this formula in the following way. Assume that $\theta(\nu) = 0$. We shall show that (5.2) implies the existence of a constant $\gamma = \gamma(p, \nu)$, satisfying $\gamma > 0$ if $p < 1$, such that

$$P_{\nu_p}(0 \to \partial\Lambda_n) \le e^{-\gamma n} \quad \text{for } n \ge 1, \tag{5.3}$$

where $\Lambda_n = \{u \in \mathbf{Z}^d : |u| \le n\}$ as usual. This implies the claim of the theorem, as follows. If μ is strictly dominated by ν then, by an argument using the continuity of $\nu_p(A)$ for $A \subseteq \mathcal{E}$, and the fact that $\nu_1 = \nu$, we may find p (< 1) such that $\mu \le \nu_p$. For this value of p, we have that

$$P_\mu(0 \to u) \le P_{\nu_p}(0 \to u) \le e^{-\gamma|u|}$$

as required.

We turn now to the proof of (5.3). Let $A_n = \{0 \to \partial\Lambda_n\}$. Assume that $\theta(\nu) = 0$, and further that $P_\nu(A_n) > 0$ and $p > 0$ (the claim is trivial otherwise). We shall follow the proof of Theorem 3.4 of Grimmett (1989), and we shall refer to the notation as well as the equation and lemma numbers of the last reference. Let $N(A_n)$ be the number of pivotal vertices for A_n. The idea is to find a lower bound for $E(N(A_n) \mid A_n)$, to substitute this into (5.2) with $B = A_n$, and to integrate. The lower bound is obtained in very much the same way as in Grimmett (1989) but with some differences.

The crucial lemma is Lemma 3.12 of Grimmett (1989). Let $y_0, y_1, \dots$ be the pivotal vertices for A_n, taken in their natural order from 0 to $\partial\Lambda_n$. We have that $y_0 = 0$. Let $\rho_i = |y_i - y_{i-1}|$. As at the bottom of page 49, and by (5.1) above,

$$\begin{aligned} P_{\nu_p}(\{\rho_1 > r_1\} \cap A_n) &\le P_{\nu_p}(\{0 \to \partial\Lambda_n\} \circ \{\Delta 0 \to \partial\Lambda_{r_1+1}\}) \\ &\le P_{\nu_p}(A_n) f_p(r_1 + 1), \end{aligned}$$

where ΔS denotes the external boundary of a set S of vertices (i.e., the set of vertices y of $\mathbf{Z}^d \setminus S$ such that y is adjacent in $\mathbf{L}^d$ to some vertex in S), and where $f_p(m) = P_{\nu_p}(\Delta 0 \to \partial\Lambda_m)$. Therefore

$$P_{\nu_p}(\rho_1 > r_1 \mid A_n) \le f_p(r_1 + 1), \quad r_1 \ge 0.$$

Next we turn to a proof of a version of (3.13) of Grimmett (1989). The proof given there has a minor imperfection which has been corrected in Grimmett (1996). Fix a vertex u of Λ_n, and define D_u to be the set of all vertices attainable from 0 along paths of Γ_p not using edges emanating from u; we turn D_u into a graph by adding all edges of Γ_p emanating from $D_u \setminus \{u\}$. Fix positive integers k (≥ 2), $r_1, r_2, \dots r_k$, such that $\sum_{i=1}^k r_i \le n$, and let B_u be the event that the following statements hold:

(a) u lies in D_u, and u is black,

(b) D_u contains no vertex of $\partial\Lambda_n$,

(c) the pivotal vertices for the event $\{0 \to u\}$ are, taken in order, $y_0, y_1, \dots, y_{k-2}$, where $|y_{i-1} - y_i| = r_i$ for $1 \le i < k-1$, and furthermore $|y_{k-2} - u| = r_{k-1}$.

We now define the event $B = \bigcup_u B_u$. If $A_n \cap B$ occurs, then there exists a unique (random) u such that B_u occurs. Now, arguing as in Grimmett (1989, 1996), but using Theorem 5.1 in place of the usual BK inequality, we obtain that for each positive integer k, and all positive integers $r_1, \dots, r_k$ with $\sum_{i=1}^k r_i \le n$,

$$P_{\nu_p}\Big(\rho_k > r_k,\ \rho_i = r_i \text{ for } 1 \le i < k \,\Big|\, A_n\Big) \le f_p(r_k+1)\, P_{\nu_p}\Big(\rho_i = r_i \text{ for } 1 \le i < k \,\Big|\, A_n\Big).$$

This leads as before to the inequality

$$E\big(N(A_n) \mid A_n\big) \ge \frac{n}{1 + \sum_{i=1}^n f_p(i)} - 1. \tag{5.4}$$

Now

$$f_p(m) \le 2d P_{\nu_p}(A_{m-1}) \quad \text{for } m \ge 1. \tag{5.5}$$

We substitute (5.5) into (5.4) and (5.2) to find that $g_p(n) = P_{\nu_p}(A_n)$ satisfies

$$\frac{d}{dp}\{\log g_p(n)\} \ge \frac{n}{1 + 2d\sum_{i=0}^n g_p(i)} - 1.$$

The proof now proceeds as before, and (5.3) follows. □

Proof of Theorem 2.2. The proof is a straightforward adaptation to this setting of the first part of Theorem 1 of Kesten (1981). We omit all details. The overall argument is roughly as follows. We have from Theorem 2.1 that the radius of C has an exponentially decaying tail. This implies that the chance of an 'easy-way crossing' of the box $[0, 3N]^{d-1} \times [0, N]$ may be made as close to zero as desired by choosing N sufficiently large. We call the point 0 *good* if one of the boxes $[0, 3N]^i \times [0, N] \times [0, 3N]^{d-i-1}$, $0 \le i \le d-1$, has an 'easy-way crossing'. This definition of 'good' may be extended, by translation, to any point of the form $(k_1N, k_2N, \dots, k_dN)$ for integers $k_1, k_2, \dots, k_d$. Now, if $|C| \ge n$, then there exists a connected cluster of 'good points on a renormalised lattice', where the scale of the renormalisation is N, this latter cluster having size at least An for some constant A. The number of possible renormalised clusters of cardinality m is less than B^m for some B. Combining these estimates, and choosing N sufficiently large, we obtain an exponential estimate for the tail of C. For the details, see Kesten (1981). □

6. Proofs of Theorems 3.1 and 3.2

Let $\beta = 0$. Following Kuulasmaa (1982) we may construct the epidemic process as follows. Let $(T_u : u \in \mathbf{Z}^d)$ be a family of independent exponentially distributed random variables with parameter 1; T_u may be taken as the length of the period of infection at u, given that infection ever reaches u. After infection, a vertex u attempts to infect its neighbours. Let $(I_{u,v} : u \in \mathbf{Z}^d, v \in \mathcal{E})$ be independent random variables having the exponential distribution with parameter α. We say that u *infects its neighbour* $u + v$ (where $v \in \mathcal{E}$) if $I_{u,v} < T_u$. For each u there is a set $u + N_u$ of neighbours of u which are infected by u. The sets $(N_u : u \in \mathbf{Z}^d)$ are independent and identically distributed according to some probability measure on $\mathcal{E}$, and we denote this measure by μ_α.

Consider the random subgraph Γ of $\vec{\mathbf{L}}^d$ given by: $[u, v\rangle$ is an edge of Γ if and only if u infects v. Then the set C of all vertices which may be reached from 0 along directed paths of Γ has the same distribution as the set of vertices which are ultimately infected in the epidemic system with parameter α and $\beta = 0$.

Proof of Theorem 3.1. Let $\alpha < \alpha_c(0)$ and pick α' such that $\alpha < \alpha' < \alpha_c(0)$. Then $P_{\alpha',0}(0 \to \infty) = 0$. Also, μ_α is strictly dominated by $\mu_{\alpha'}$, whence the claims follow from Theorems 2.1 and 2.2. $\square$

Proof of Theorem 3.2. We have that

$$T \leq \sum_{u \in C} T_u$$

where the T_u are given above. Now,

$$\begin{aligned} P_{\alpha,0}(T > t) &= \sum_{m=1}^{\infty} \sum_{G:|G|=m} P_{\alpha,0}(T > t, C = G) \\ &\leq \sum_{m \leq M} \sum_{G:|G|=m} P_{\alpha,0}\left(\sum_{i=1}^{m} T_i > t\right) + P_{\alpha,0}(|C| \geq M) \end{aligned}$$

where $T_1, T_2, \ldots$ are independent exponentially distributed random variables with parameter 1. Therefore

$$P_{\alpha,0}(T > t) \leq D^M P_{\alpha,0}\left(\sum_{i=1}^{M} T_i > t\right) + e^{-\delta M}$$

where δ is given in Theorem 3.1, and D (> 1) is a constant depending on the dimension d. Using Markov's inequality,

$$P_{\alpha,0}\left(\sum_{i=1}^{M} T_i > t\right) \leq \frac{e^{-\theta t}}{(1-\theta)^M} \quad \text{for } \theta < 1.$$

We may take $\theta = \frac{1}{2}$, and we set $M = \lfloor tr \rfloor$ where $r < 1/(2\log(2D))$, thereby obtaining the required exponential estimate. $\square$

7. Proof of Theorem 4.1

This proof was inspired by ideas in Menshikov and Pelikh (1990) and van den Berg and Ermakov (1996). We do not present all details, but we hope that the following broad account will satisfy most readers. Let $0 < \alpha, \beta < \infty$. In a manner similar to the first part of Section 6, we set out to assign random variables to each vertex u in order to describe the evolution of the process. In fact, our variables will only *partially* describe the process, but this appears to be sufficient for our purpose. To each vertex u, we assign a random element ω_u which, roughly speaking, describes the 'actions' of u between the moment it is infected for the first time (*if* u is ever infected) up to the end of its second infectious period. More precisely, ω_u consists of the following three components. The first component is a marked Poisson point process $\pi_u(1)$ on the positive half-line, where the possible marks are the elements $v \in \mathcal{E}$ (each of which occurs with intensity α) and the special mark 0, which occurs with intensity 1 (the set $\mathcal{E}$ was defined at the beginning of Section 2). This point process, up to the time of the first occurrence of a point with mark 0, describes the infection attempts by u during its first infectious period. The second component of ω_u is an exponentially distributed random variable R_u (with parameter β), denoting the first removal period of u (i.e., the length of the time-interval subsequent to its first infection, during which u is in state 0). Finally, the third component is a point process $\pi_u(2)$ with the same distribution as $\pi_u(1)$. This point process $\pi_u(2)$, at times up to the first occurrence of a point with mark 0, describes the infection attempts by u during its second infectious period.

The ω_u, $u \in \mathbf{Z}^d$, are independent and identically distributed, and for each u we specify that R_u, $\pi_u(1)$, and $\pi_u(2)$ are independent of each other. We continue to write $P_{\alpha,\beta}$ for the appropriate probability measure, and we let $\omega = (\omega_u : u \in \mathbf{Z}^d)$. Any point marked m in a point process is called an m-point.

The realizations of the point processes after the occurrence of the first 0 are without meaning for the time being, but later they (or at least part of them) will be useful in the construction of a certain mapping.

In the following, $T_u(1)$ and $T_u(2)$ will denote the length of the first and second infectious periods of u respectively (i.e., $T_u(1)$ is the position of the first 0 in $\pi_u(1)$ and $T_u(2)$ that of the first 0 in $\pi_u(2)$). For $v \in \mathcal{E}$, we say that u tries to infect $u+v$ in its first (resp. second) infectious period if $\pi_u(1)$ (resp. $\pi_u(2)$) has a v-point in the interval $(0, T_u(1))$ (resp. $(0, T_u(2))$).

Suppose we start with only the origin 0 ill, and all others susceptible. In order to describe the evolution completely, we need of course much more than the above information, but this limited information allows us to construct an appropriate directed graph Γ such that: if there is a path from 0 to a vertex v in this graph, then v is ill at some time.

We perform this construction step by step as follows. First, suppose we are provided with some fixed total order on $\mathbf{Z}^d$. We shall construct a directed graph by means of a sequence of iterations. Let Γ_0 be the directed graph which comprises the single vertex 0 only; we speak of 0 as having been *selected*. At Step 1, we draw an arrow from 0 to all $v \in \mathcal{E}$ for which 0 tries to infect v during its first infectious period, and we call these vertices *children* of 0 (and 0 their *parent*); the consequent directed graph is denoted Γ_1. Note that, although we cannot conclude that each such v is infected directly by 0, we *may* conclude that they are infected by *some* vertex (this is similar to the case $\beta = 0$). Some children v are special: if $\pi_0(1)$ has two v-points in the interval $(0, T_0(1))$ at distance larger than $T_v(1) + R_v$ from each other, (i.e., if the period between the first and the last time during its first period of illness that 0 tries to infect v is larger than the first illness and recovery cycle of v), then a few moments of reflection reveal that v is infected at least twice (although not necessarily by 0). We then call v a *good child* of 0.

More generally, in the construction of Γ_{k+1} at Step $k+1$, we do the following. We select the smallest vertex u (with respect to the pre-defined order) which has not been selected during the contruction of $\Gamma_0, \Gamma_1, \ldots, \Gamma_k$, and which can be reached from 0 along directed paths of Γ_k. Let $v \in \mathcal{E}$ be such that $u+v$ is not a vertex of Γ_k. If u tries to infect $u+v$ during its first infectious period we can, as in the first step, conclude that $u+v$ will be infected at some time. There is another situation in which $u+v$ is infected, namely if u is a good child of its parent, and u tries to infect v during its second infectious period. In either case, we say that $u+v$ is a child of u, and we draw an arrow from u to $u+v$. We denote by Γ_{k+1} the directed graph obtained by augmenting Γ_k with all such arrows and all incident vertices. Finally, as in Step 1, if $\pi_u(1)$ has two v-points in the interval $(0, T_u(1))$ at distance larger than $T_{u+v}(1) + R_{u+v}$ from each other, then we call $u+v$ a good child of u (and conclude that it is infected at least twice).

In this way, we construct an increasing sequence $\Gamma_k = \Gamma_k(\omega)$ of directed graphs, and it is not difficult to show (by induction) that every vertex of every Γ_k is infected at some time. For later purposes, we note the following. Conditional on a vertex u being selected at some given stage, then ω_u is independent of the ω_x for vertices x selected earlier. In particular, $T_u(1)$, $T_u(2)$, and R_u may be regarded as 'fresh', independent exponentially distributed random variables having parameters 1, 1, and β respectively. Also note that if, in the construction above, we were to decide to draw an edge from u to its neighbour v only when u tries to infect $u+v$ during its first infectious period, then we would obtain exactly the model without recovery of Sections 3 and 6. In order to prove our theorem, we have to show that, at each step, with u denoting the vertex selected at that step, the set

$$\begin{aligned} X_u = \big\{v : \ &(u \text{ tries to infect } u+v \text{ during its first infectious period}), \text{ or} \\ &(u \text{ is good and } u \text{ tries to infect } u+v \text{ during its second infectious period})\big\} \end{aligned}$$

is 'uniformly stochastically strictly larger' than the corresponding set we would have in the case $(\alpha, 0)$, i.e., the set

$$Y_u = \big\{v : u \text{ tries to infect } u+v \text{ during its first infectious period}\big\};$$

this will imply (by continuity) that there exists $\gamma > 0$ such that X_u is uniformly stochastically larger than Y_u with α replaced by $\alpha + \gamma$. This implies the claim of the theorem.

In order to achieve this, suppose we observe the evolution of the sequence (Γ_k), but that we do not know the underlying 'fine-structure' ω. Suppose, at Step K, we have obtained that $\mathbf{\Gamma}_K = (\Gamma_0, \Gamma_1, \ldots, \Gamma_{K-1})$ is given by $\mathbf{\Gamma}_K = \boldsymbol{\gamma}_K$, where $\boldsymbol{\gamma}_K = (\gamma_0, \gamma_1, \ldots, \gamma_{K-1})$ is a specified sequence of directed graphs γ_i. Suppose further that $u = u(\mathbf{\Gamma}_K)$ is the selected vertex. Write $A = \{\omega : \mathbf{\Gamma}_K(\omega) = \boldsymbol{\gamma}_K\}$.

Let p denote the parent of u, and let $w = u - p$. Write B be for the event that $\pi_p(1)$ has, in the interval $(0, T_p(1))$, two w-points at distance larger than 1 from each other. Also, let $E_1 = \{T_u(1) < \frac{1}{2}\}$, and $E_2 = \{R_u < \frac{1}{2}\}$. Let $\mathcal{A} \subseteq 2^{\mathcal{E}}$ be increasing and $\mathcal{A} \neq 2^{\mathcal{E}}$, $\mathcal{A} \neq \emptyset$. Finally, let

$$Y'_u = \big\{v : u \text{ tries to infect } u + v \text{ during its second infectious period}\big\}$$

(compare with X_u and Y_u defined above). Clearly, if $B \cap E_1 \cap E_2$ holds, then u is good. Therefore, dropping the subscripts α, β of P for a moment, we have that

$$\begin{aligned}
&P\big(X_u \in \mathcal{A} \,\big|\, A\big) \\
&\quad = P(B \cap E_1 \cap E_2 \,|\, A)\, P\big(Y_u \cup Y'_u \in \mathcal{A} \,\big|\, A \cap B \cap E_1 \cap E_2\big) \\
&\qquad\qquad + P\big((B \cap E_1 \cap E_2)^c \,\big|\, A\big)\, P\big(Y_u \in \mathcal{A} \,\big|\, A \cap (B \cap E_1 \cap E_2)^c\big) \\
&\quad \geq P\big(B \cap E_1 \cap E_2 \,\big|\, A\big) \times \\
&\qquad\qquad\qquad \Big[P\big(Y_u \in \mathcal{A} \,\big|\, A \cap B \cap E_1 \cap E_2\big) + P\big(Y_u \notin \mathcal{A}, Y'_u \in \mathcal{A} \mid A, B, E_1, E_2\big)\Big] \\
&\qquad\qquad + P\big((B \cap E_1 \cap E_2)^c \,\big|\, A\big)\, P\big(Y_u \in \mathcal{A} \,\big|\, A \cap (B \cap E_1 \cap E_2)^c\big) \\
&\quad = P\big(Y_u \in \mathcal{A} \,\big|\, A\big) + P\big(B \cap E_1 \cap E_2 \,\big|\, A\big)\, P\big(Y_u \notin \mathcal{A},\, Y'_u \in \mathcal{A} \,\big|\, A \cap B \cap E_1 \cap E_2\big).
\end{aligned}$$

It is clear that the occurrence of A (resp. B) is completely determined by the set of ω_x for which the vertex x was selected before step K, while, as noted before, ω_u is independent of the family of such ω_x. Therefore, the first term in the final expression equals $P_{\alpha,\beta}(Y_u \in \mathcal{A})$, and the second term equals

$$P_{\alpha,\beta}(B \mid A) P_{\alpha,\beta}(E_1) P_{\alpha,\beta}(E_2) P_{\alpha,\beta}(Y'_u \in \mathcal{A}) P_{\alpha,\beta}(Y_u \notin \mathcal{A} \mid E_1),$$

which in turn equals

$$P_{\alpha,\beta}\big(B \,\big|\, A\big) P_{\alpha,\beta}\big(E_2\big) P_{\alpha,\beta}\big(Y'_u \in \mathcal{A}\big) P_{\alpha,\beta}\big(Y_u \notin \mathcal{A},\ E_1\big).$$

Since the last three factors are strictly positive and do not depend on A, it suffices to prove that $P_{\alpha,\beta}(B \mid A)$ is uniformly (over all events A of the form given above) bounded away from 0. This final statement is proved using the following 'modification' argument.

Let A_w be the subset of A containing all ω which satisfy the following four properties.

(i) $\pi_p(1)$ has a w-point in the interval $(T_p(1), T_p(1) + 1)$ and in the interval $(T_p(1) + 2, T_p(1) + 3)$.

(ii) If p is not a good child (for the configuration ω), then $\pi_p(1)$ has no points in the interval $(T_p(1), T_p(1) + 3)$ with marks other than w.

(iii) If p is a good child, then: for each $y \in \mathcal{E} \setminus \{w\}$ such that $\pi_p(2)$ has a y-point in the interval $(0, T_p(2))$ and $\pi_p(1)$ has no y-point in the interval $(0, T_p(1))$, then $\pi_p(1)$

has exactly one y-point in the interval $(T_p(1), T_p(1)+3)$, and besides w these are the only marks in $\pi_p(1)$ occurring in the interval $(T_p(1), T_p(1)+3)$.

(iv) $\pi_p(1)$ has no point at position $T_p(1)+3$, and its only point at position $T_p(1)$ is the one with mark 0.

Remark: The fourth property may appear superfluous since it has probability 1, but we include it in order to make precise our argument involving the map H below. Note that (ii)–(iv) imply that $\pi_p(1)$ has no 0-point in the interval $(T_p(1), T_p(1)+3]$.

It may be seen that $P_{\alpha,\beta}(A_w) \geq gP_{\alpha,\beta}(A)$, where $g = g(\alpha)$ is given by

$$g = (1-e^{-\alpha})^2 e^{-3} \min_{0\leq k\leq 2d-1} \Big\{ e^{-3k\alpha} \big[e^{-3\alpha} 3\alpha \big]^{2d-1-k} \Big\}.$$

[The first factor comes from (i), the second from the required absence of 0-points in the interval $(T_p(1), T_p(1)+3]$, and the third factor from (ii) and (iii).]

Now, let H be the mapping which assigns to each $\omega \in A_w$ the element $H(\omega)$ obtained by shifting the first 0-point in $\pi_p(1)$ distance 3 to the right. Using the note in the remark above, this map $H : A_w \to H(A_w)$ is one–one and measure-preserving, in the sense that $P_{\alpha,\beta}(H(A_w)) = P_{\alpha,\beta}(A_w)$. Therefore we have that

$$P_{\alpha,\beta}(H(A_w)) = P_{\alpha,\beta}(A_w) \geq gP_{\alpha,\beta}(A).$$

Moreover, it may be checked that the properties (i)–(iv) of A_w guarantee that $H(A_w) \subseteq A$ (in particular, property (iii) guarantees that, although the possible goodness of p may be disturbed by the map, this has no serious consequences), and that $H(A_w) \subseteq B$. It follows that

$$P_{\alpha,\beta}\big(B \,\big|\, A\big) \geq P_{\alpha,\beta}\big(H(A_w) \,\big|\, A\big) = \frac{P_{\alpha,\beta}(H(A_w))}{P_{\alpha,\beta}(A)} \geq g(\alpha),$$

which completes the proof.

8. Proof of Theorem 4.4

Throughout this section, we think of the epidemic process as being generated by its graphical representation. That is to say, we are given appropriate families of Poisson processes which may be used to couple together the different epidemic processes corresponding to different initial conditions. Such constructions are standard, and may be found in Bezuidenhout and Grimmett (1990), Griffeath (1979), Harris (1974), Liggett (1985), and elsewhere. We shall continue to use the notation $P_{\alpha,\beta}$ to denote the relevant probability measure; this notation is not entirely appropriate, since the initial configuration (3.1) is not germane to the following discussion.

Let A be a finite subset of $\mathbf{Z}^d$. We shall show that the following holds for small positive β. There exists an a.s. finite (random) time T_A such that the space–time region $A \times [T_A, \infty)$ $(\subseteq \mathbf{Z}^d \times [0,\infty))$ contains no 2's. This statement implies the claim of the theorem. We prove this under the assumption that $d = 2$, in order to avoid more cumbersome notation; the case $d = 1$ is simpler, and no essentially new difficulty emerges when $d \geq 3$.

We define two space–time regions:

$$\mathcal{A} = [-2L, 2L]^2 \times [0, 2T], \qquad \mathcal{B} = [-L, L]^2 \times [T, 2T]$$

where L and T are integers to be chosen later. Define $\mathcal{C}$ to be the part of the 'boundary' of the box $\mathcal{A}$:

$$\mathcal{C} = \Big\{(m, n, t) \in \mathcal{A} : |m| = 2L \text{ or } |n| = 2L \text{ or } t = 0\Big\}$$

We will compare the process ξ_t to a certain dependent percolation process on the set $\mathcal{L} = \mathbf{Z}^2 \times \mathbf{Z}_+$, where $\mathbf{Z}_+ = \{0, 1, 2, \ldots\}$. We say that the site (k, m, n) in $\mathcal{L}$ is *wet* if there exist no 2's in the box $(kL, mL, nT) + \mathcal{B}$ *regardless of the states of sites in the boundary* $(kL, mL, nT) + \mathcal{C}$. Note that the event $\{(k, m, n) \text{ is wet}\}$ depends only on the existence (or not) of paths of infection within $\mathcal{A}$. We shall require this uniformity on the states of the boundary in order to ensure that the percolation process in $\mathcal{L}$, although dependent, has an interaction with only finite range. Sites which are not wet are called *dry*.

Let $\alpha_c(\infty) < \alpha < \alpha_c(0)$, and $\epsilon > 0$. We will now show that there exist L, T, and $\beta_c > 0$ (depending on α, ϵ) such that:

$$P_{\alpha,\beta}\big((k, m, n) \text{ is dry}\big) \leq \epsilon \quad \text{if } \beta < \beta_c.$$

We start by showing the above property when $\beta = 0$. Then, using a continuity argument, we shall deduce that the inequality remains true for small positive β. By translation-invariance, it suffices to consider the site $(0, 0, 0) \in \mathcal{L}$.

Suppose that $(x, y, t) \in \mathcal{B}$ is such that $\xi_t(x, y) = 2$. Then there must exist some point $(x', y', t') \in \mathcal{C}$ such that (a) $\xi_{t'}(x', y') = 2$, and (b) there exists a 'chain' of infection from (x', y', t') to (x, y, t) lying entirely within $\mathcal{A}$. Such (x', y', t') must lie either on the 'bottom' of $\mathcal{C}$ (i.e., have $t' = 0$) or on one of its 'sides' (i.e., have $t' \neq 0$). In the former case, the infection originating at (x', y', t') must have survived at least time T before it reaches $\mathcal{B}$, whilst in the latter case it must have radius at least L. We propose to use Theorems 3.1 and 3.2 to control the probabilities of these possibilities. However, these theorems were proved under the assumption that infection originated in *a single vertex* having state 2, surrounded by an ocean of 1's. The effect of augmenting the original configuration by adding extra infected vertices is to diminish the set of the points in space–time reached by the infection starting at (x', y', t'); this holds when $\beta = 0$ since any extra initial infections may cause the subsequent removal of points which might otherwise have assisted the spread of infection from (x', y', t').

We may therefore apply Theorems 3.1 and 3.2 to deduce that there exist $\gamma, \lambda > 0$ such that

$$P_{\alpha,0}\big((0, 0, 0) \text{ is dry}\big) \leq 8T(4L + 1)e^{-\gamma L} + (4L + 1)^2 e^{-\lambda T}.$$

We may take $L = T$ sufficiently large such that

$$P_{\alpha,0}\big((0,0,0) \text{ is dry}\big) \le \tfrac{1}{2}\epsilon.$$

As observed above, the state of any site in $\mathcal{L}$ depends only on the graphical representation within the appropriate copy of $\mathcal{A}$. Since this region is bounded, the density of wet sites is a continuous function of α and β. Therefore there exists $\beta_c = \beta_c(\alpha, \epsilon) > 0$ such that

$$P_{\alpha,\beta}\big((0,0,0) \text{ is dry}\big) \le \epsilon \quad \text{if } \beta < \beta_c. \tag{8.1}$$

We choose β_c accordingly, noting that we may take $\beta_c(\alpha, \epsilon)$ to be a strictly positive and continuous function of α on $(0, \alpha_c(0))$.

We now position oriented edges between sites in $\mathcal{L}$ in order to obtain a percolation model. Let $\mathcal{A}(k,m,n) = (kL, mL, nT) + \mathcal{A}$. For each pair $(k,m,n), (x,y,z) \in \mathcal{L}$, we draw an oriented edge from (k,m,n) to (x,y,z) if and only if $n \le z$ and $\mathcal{A}(k,m,n) \cap \mathcal{A}(x,y,z) \ne \emptyset$. The wet sites on the ensuing directed graph constitute a (dependent) percolation model. There exists an absolute constant K, depending only on the number d of dimensions (here, $d = 2$), such that any set of sites of $\mathcal{L}$ have independent states whenever the graph-theoretic distance between any pair of such sites exceeds K (this distance is to be measured on the undirected graph obtained from $\mathcal{L}$ by removing the orientations). Furthermore, there exist positive finite constants δ, ν such that following two statements hold. First, the number of self-avoiding oriented paths on $\mathcal{L}$, having length r and any given endpoint, is no larger than δ^r. Secondly, any self-avoiding path of length r contains at least νr sites, the distance between any pair of which exceeds K.

Let $x \in \mathbf{Z}^2$ and let T_x be the supremum of all times t such that $\xi_t(x) = 2$. We claim that T_x is a.s. finite if β is sufficiently small. The theorem will follow from this statement, since the T_A given there satisfy $T_A = \max\{T_x : x \in A\}$. It will suffice to prove that T_0 is a.s. finite, since the argument is 'translation-invariant'.

Suppose that $T_0 > TM$. Then there exists m $(\ge M-1)$ with the property that $(0,0,m)$ is the endpoint of an oriented dry path of $\mathcal{L}$ whose other endpoint has the form $(x,y,0)$ for some $x, y \in \mathbf{Z}^2$. By the above remarks,

$$P_{\alpha,\beta}(T_0 > TM) \le \sum_{m \ge M-1} \sum_{r \ge m} \delta^r \eta^{\nu r}, \tag{8.2}$$

where $\eta = P_{\alpha,\beta}\big((0,0,0) \text{ is dry}\big)$. By (8.1), we may choose $\beta_c = \beta_c(\alpha)$ such that the right side of (8.2) is finite whenever $\beta < \beta_c$ and $M \ge 2$. When this holds, the right side approaches 0 as $M \to \infty$, implying that T_0 is a.s. finite as required.

Since the estimates presented above depend only on the graphical representation, and not on the initial configuration, we may deduce also that $\theta(\alpha, \beta) = 0$ under the conditions of the theorem.

The above proof is related in part to Theorem 2 of Durrett and Schinazi (1993). Using further arguments based on the shape theorem for oriented percolation, one may obtain linearly growing estimates for cones in space–time which are devoid of vertices having state 2. Similar arguments have been used by Durrett (1992) in the context of the Greenberg–Hastings and cyclic colour models. □

Acknowledgements. G.R.G. and R.B.S. began this project at IMPA (Rio de Janeiro) in 1995 while attending the Coloquio Brasileiro de Matematica. They thank Pablo Ferrari for hospitality and CNPq for financial support.

References

Aizenman, M. and Barsky, D. J. (1987). Sharpness of the phase transition in percolation models. *Communications in Mathematical Physics* **108**, 489–526.

Aizenman, M. and Grimmett, G. (1991). Strict monotonicity for critical points in percolation and ferromagnetic models. *Journal of Statistical Physics* **63**, 817–835.

Andjel, E. and Schinazi, R. (1996). A complete convergence theorem for an epidemic model. *Journal of Applied Probability* **33**, 741–748.

Berg, J. van den (1997). Some reflections on disjoint occurrences of events. In preparation.

Berg, J. van den and Ermakov, A. (1996). A new lower bound for the critical probability of site percolation on the square lattice, *Random Structures and Algorithms* **8**, 199–212.

Berg, J. van den and Kesten, H. (1985). Inequalities with applications to percolation and reliability. *Journal of Applied Probability* **22**, 556–569.

Bezuidenhout, C. E. and Grimmett, G. R. (1990). The critical contact process dies out. *Annals of Probability* **18**, 1462–1482.

Durrett, R. (1992). Multicolor particle systems with large threshold and range. *Journal of Theoretical Probability* **5**, 127–152.

Durrett, R. and Neuhauser, C. (1991). Epidemics with recovery in $D = 2$. *Annals of Applied Probability* **1**, 189–206.

Durrett, R. and Schinazi, R. (1993). Asymptotic critical value for a competition model *Annals of Applied Probability* **3**, 1047–1066.

Griffeath, D. (1979). *Additive and Cancellative Interacting Particle Systems.* Lecture Notes in Mathematics no. 724, Springer-Verlag, Berlin.

Grimmett, G. R. (1989). *Percolation.* Springer-Verlag, Berlin.

Grimmett, G. R. (1996). Percolation and disordered systems. St Flour Lecture Notes, to appear.

Harris, T. E. (1974). Contact interactions on a lattice. *Annals of Probability* **2**, 969–988.

Kesten, H. (1981). Analyticity properties and power law estimates in percolation theory. *Journal of Statistical Physics* **25**, 717–756.

Kuulasmaa, K. (1982). The spatial general epidemic and locally dependent random graphs. *Journal of Applied Probability* **19**, 745–758.

Liggett, T. M. (1985). *Interacting Particle Systems.* Springer-Verlag, New York.

Menshikov, M. V. (1986). Coincidence of critical points in percolation problems. *Soviet Mathematics Doklady* **33**, 856–859.

Menshikov, M. V. and Pelikh, K. D. (1990). Percolation with several types of defects. Estimates of the critical probability on the square lattice. [In Russian] *Mathematical Notes* **46**, 38–47; [English translation] *Math. Notes Acad. Sci. USSR* **46**, 778–785.

Reimer, D. (1997). Proof of the van den Berg–Kesten inequality. (Revised version of 1995 paper entitled ‘Butterflies’.) To appear.