

Random spatial growth with paralyzing obstacles

J. van den Berg¹, Y. Peres², V. Sidoravicius³ and M. E. Vares⁴

CWI and VUA, Microsoft and UC Berkeley, IMPA and CBPF. E-mail: J.van.den.Berg@cw.nl; peres@stat.berkeley.edu; vldas@impa.br; eulalia@cbpf.br

Received 1 June 2007; revised 11 December 2007; accepted 18 December 2007

Abstract. We study models of spatial growth processes where initially there are sources of growth (indicated by the colour green) and sources of a growth-stopping (paralyzing) substance (indicated by red). The green sources expand and may merge with others (there is no ‘inter-green’ competition). The red substance remains passive as long as it is isolated. However, when a green cluster comes in touch with the red substance, it is immediately invaded by the latter, stops growing and starts to act as a red substance itself. Our main model space is represented by a graph, of which initially each vertex is randomly green, red or white (vacant), and the growth of the green clusters is similar to that in first-passage percolation. The main issues we investigate are whether the model is well defined on an infinite graph (e.g. the d -dimensional cubic lattice), and what can be said about the distribution of the size of a green cluster just before it is paralyzed. We show that, if the initial density of red vertices is positive, and that of white vertices is sufficiently small, the model is indeed well defined and the above distribution has an exponential tail. In fact, we believe this to be true whenever the initial density of red is positive. This research also led to a relation between invasion percolation and critical Bernoulli percolation which seems to be of independent interest.

Résumé. Nous étudions des modèles de croissance spatiaux qui, au temps initial, ont des sources de croissance (indiquées par la couleur verte) et des sources de substance paralysante arrêtant la croissance (indiquées en rouge). Les sources vertes augmentent et peuvent fusionner avec les autres (il n’y a pas de compétition entre elles). La substance rouge reste passive quand elle est isolée. Cependant, quand un amas vert touche la substance rouge, il est envahi immédiatement par cette dernière, il arrête de grandir et commence à agir comme la substance rouge. Dans notre modèle principal, l’espace est représenté par un graphe dont, à l’instant initial, tous les sommets sont tirés au hasard vert, rouge ou blanc (vide) et la croissance des amas verts est similaire à celle de la percolation de premier-passage. Les problèmes principaux que nous considérons sont les suivants ; est-ce que le modèle est bien défini sur un graphe infini (par exemple le treillis d -dimensionnel) ? Que peut-on dire de la distribution de la taille d’un amas vert juste avant qu’il soit paralysé ? Nous montrons que, si la densité initiale de sommets rouges est positive et que celle des sommets blancs est suffisamment petite, le modèle est en effet bien défini et la distribution ci-dessus mentionnée a une queue exponentielle. Nous conjecturons que ce résultat est vrai dès que la densité initiale des rouges est positive. Ce travail mène également à une relation entre la percolation d’invasion et la percolation de Bernouilli critique qui semble être d’intérêt indépendamment.

MSC: Primary 60K35; secondary 60K37, 82B43

Keywords: Growth process; Percolation; Invasion percolation

¹Research funded in part by the Dutch BSIK/BRICKS project.

²Research supported in part by NSF grant DMS-0605166.

³Partially supported by CNPq, Brazil.

⁴Partially supported by CNPq, Brazil.

1. Introduction

1.1. Description of the model and the main problems

Consider the following model where different ‘objects’ (or ‘populations’) grow simultaneously until they hit a paralyzing substance, in which case they stop growing and become paralyzing themselves: Each vertex of a connected, finite (or countably infinite, locally finite) graph $G = (V, E)$ is initially, independently of the other vertices, white, red or green with probabilities p_w , p_r and p_g respectively. Each edge of G is initially closed. By a green cluster we will mean a maximal connected subgraph of G of which all vertices are green and all edges are open. We denote the green cluster containing v at time t by $C_g(v, t)$. (If v is not green at time t , then $C_g(v, t)$ is empty.) It is clear from the above that initially the only green clusters are single green vertices. These green clusters can grow, merge with other green clusters and finally become paralyzed (red) as follows.

Whenever an edge $e = \langle v, w \rangle$ is closed and has at least one green end-vertex, say v , it becomes open at rate 1. Moreover, immediately after it gets open the following action takes place instantaneously: If exactly one end-vertex, say v , is green and the other, w , is white, w becomes green (and we say, informally, that the green cluster of v grows by absorbing w). If w is red, then each vertex in the green cluster of v becomes red (and we say that the green cluster of v becomes paralyzed). Finally, if both vertices are green, no extra action takes place. (Note that in this case the two vertices may have been in two different green clusters right before the opening of e , but are now in the same green cluster.)

Note that once an edge is open it remains open, that once a vertex is green it never turns white (but may become red), and once a vertex is red it remains red.

Let us first consider the case where the graph G is finite. In that case the above process is clearly well defined and has some obvious properties, which we will state after introducing the following terminology. By a configuration (or ‘site-bond configuration’) we mean an element of $\{0, 1\}^E \times \{\text{green, red, white}\}^V$, where 0 and 1 denote ‘open’ and ‘closed’ respectively. An ‘open-bond cluster’ (with respect to a configuration) is a maximal connected subgraph of G of which all edges are open (for that configuration). We say that it is non-trivial if it has at least one edge. Note that the earlier defined ‘green cluster’ is an open-bond cluster of which each vertex is green. A ‘red cluster’ is defined similarly. We call a configuration admissible if each non-trivial open-bond cluster is either a red cluster or a green cluster. Now we are ready to state the announced simple properties and observations: If G is finite, the process is a Markov chain on the set of admissible configurations. The admissible configurations where no vertices are green or all vertices are green are absorbing, and the chain will with probability 1 end in one of those configurations. In particular, if initially there was at least one red vertex, then every green vertex will eventually become red. Moreover (because initially all edges were closed) at any time, every non-empty red cluster \mathcal{C} contains exactly one vertex v that was originally red. We say that this vertex v is ‘responsible for’ the other vertices in \mathcal{C} becoming red (or, that the vertices in \mathcal{C} became red ‘due to’ v).

If G is *infinite*, for instance the d -dimensional cubic lattice, the situation is much more problematic, due to the fact that the range of the interaction is not bounded: an entire cluster, no matter how large, can change colour instantaneously. The main questions we address in this paper concerning the above process, and some other related models, are:

1. Does the dynamics exist? This is a non-trivial issue for such interacting processes on infinite graphs: See for instance, Aldous’ frozen percolation process [1], which was shown by Benjamini and Schramm ([3], private communication) not to exist in \mathbb{Z}^2 . For related matters on the non-existence of that process, see also Remark (i) in Section 3 of [4] and the example due to Antal Jarai ([11], private communication) which follows it. A crucial difference between Aldous’ model and ours is that in Aldous’ model, clusters freeze only when they are infinite, while we believe that in our model, due to the positive density of initially red vertices, the green clusters do not become infinite (see the next item). A model which has more in common with ours is the forest-fire model studied in [7]. But again there is a major difference: in that model there is a uniform lower bound for the probability that a cluster of interest is ‘destroyed’ before growing further, and this uniform bound is a crucial ingredient in the existence proof in [7]. In our model there seems to be no analog of such a property.
2. Is a green cluster always finite at the moment it becomes red? Does the distribution of its radius (and of its volume) have an exponential tail?

3. Let w be an originally red vertex. Is the set of originally green vertices v with the property that w is responsible for v becoming red, finite? Does the distribution of its volume have an exponential tail?

The organization of the paper is as follows. In Section 1.2 we give a partial answer to the questions listed above. In particular, Theorem 1.1 states that, for $G = \mathbb{Z}^d$ and p_w sufficiently small, the answers to the above questions are positive. Our research also led to a new result for invasion percolation (see Theorem 1.2 and Proposition 1.3). In Section 1.3 we explain the notion of ‘autonomous region’ which plays an important role in this paper. In Section 1.4 we briefly discuss some alternative versions of the model. In Section 2 we give a proof of the main result for the special case where $p_w = 0$. It turns out that that case can be dealt with in a very elegant and transparent way. It serves as an introduction to the proof of the more complicated case where p_w is small but positive, which is given in Section 3. At the end of Section 3 we come briefly back to the alternative versions of the model discussed in Section 1.4.

1.2. Statement of the main results

Let G be a connected, countably infinite graph of bounded degree, and consider the model presented in Section 1.1, with parameters p_w , p_g and p_r . Our main result, Theorem 1.1, states, among other things, that under certain conditions the dynamics is well defined. The formulation of the condition requires some additional notation and terminology: By the distance $d(v, w)$ between two vertices v and w of G we mean the length (i.e. number of edges) of the shortest path from v to w . The diameter of a set of vertices W of G is defined as $\max_{v, w \in W} d(v, w)$, and ∂W will denote the set of all vertices that are not in W but have an edge with some vertex in W . The number of elements of a set W will be denoted by $|W|$. For a finite graph H , denote by $|H|$ the number of vertices in H . Let D denote the maximal degree in G .

For each vertex v of G and $p \in (0, 1)$, let $\xi_v(p)$ denote the expectation of the volume (i.e. number of vertices) of the occupied cluster of v in site percolation on G with parameter p . Further, define

$$\xi(p) = \sup_v \xi_v(p).$$

Recall the definition of $C_g(v, t)$ in Section 1.1. We are now ready to state our main results.

Theorem 1.1. *Suppose that*

$$(D - 1)\xi(p_w) < p_r. \tag{1}$$

We have:

(a) *The dynamics on G is well defined. With probability 1, at any time, each red cluster has a unique initially red vertex.*

(b) *For any originally green vertex v , let $C_g(v) = \bigcup_{t \geq 0} C_g(v, t)$ be the green cluster of v just before it becomes red. Let $|C_g(v)|$ be the number of vertices of $C_g(v)$. Then, with probability 1, $|C_g(v)|$ is finite for each such v . Moreover, the distribution of $|C_g(v)|$ has an exponential tail.*

(c) *If G is a Cayley graph and w is an originally red vertex in G , then the set $D(w)$ consisting of all green vertices that become red due to w is finite; moreover, the diameter of $D(w)$ has an exponential tail. (Here, extending the definition given before in the case of finite G , if v is an originally green vertex and w is the (unique a.s.) originally red vertex in the red clusters that eventually contain v , we say that v becomes red due to w .)*

(d) *If G is the d -dimensional cubic lattice, then the distribution of $|D(w)|$ also has an exponential tail.*

Note that in the case $p_w = 0$, condition (1) of Theorem 1.1 is satisfied for every positive p_r . For this case we have, in addition to Theorem 1.1, considerably stronger results. In particular, the following theorem holds, where we fix $p_w = 0$ and then vary the parameter p_r . In this theorem and its proof, P_p denotes the ordinary (Bernoulli) bond percolation measure with parameter p and P_{cr} stands for P_{p_c} , where p_c denotes the critical probability for this percolation model. By $B(n)$ we denote the set of all vertices at (graph) distance $\leq n$ from some specified vertex O . The event that there is an open path from O to $\partial B(n)$ is denoted by $\{O \leftrightarrow \partial B(n)\}$. Further, the symbol \approx denotes logarithmic equivalence, i.e., we say for two positive functions $g(n)$ and $h(n)$ that $g(n) \approx h(n)$ as $n \rightarrow \infty$, if

$$\frac{\log h(n)}{\log g(n)} \rightarrow 1, \quad n \rightarrow \infty.$$

Let W be a set of vertices in a graph G with a distinguished vertex O . By the *radius* of W we mean the maximal distance from O to a vertex of W . We are now ready to state the following theorem.

Theorem 1.2. *Let $C_g(\cdot)$ be as in part (b) of Theorem 1.1. If G is the square lattice in two dimensions (or the triangular or the hexagonal lattice), and $p_w = 0$, then*

$$P(\text{the radius of } C_g(O) \text{ is at least } n) \uparrow f(n), \quad \text{as } p_r \downarrow 0,$$

where f is a function satisfying

$$f(n) \approx P_{\text{cr}}(O \leftrightarrow \partial B(n)).$$

Theorem 1.2 follows easily from the following proposition concerning invasion percolation on the lattices considered in the theorem. Before we state it, we briefly recall the invasion percolation model (on these lattices) and some of its basic properties. (Invasion percolation was introduced by Wilkinson and Willemsen, see [18]. For a detailed study of this process see [16], or the earlier works [2,6] and [12].) To each edge e we assign, independent of the other edges, a random variable τ_e , uniformly distributed in the interval $(0, 1)$. We construct, recursively, a growing tree. Initially the tree consists only of one vertex, say O . At each step we consider all edges that have exactly one endpoint in the tree that has been created so far. From these edges we select the one with smallest τ value and add it (and its ‘external’ endpoint) to the tree. Let $\tau(n)$ be the τ value of the n th edge invaded by this procedure. For any infinite transitive graph G , it is proved in [10] that

$$\limsup_{n \rightarrow \infty} \tau(n) = p_c, \tag{2}$$

where p_c is the critical probability for bond percolation. Further, note that, if all $\tau(n) < p_c$, then O belongs to an infinite cluster on which all τ values are smaller than p_c . For the graphs in the statement of Theorem 1.2 this latter event has probability 0. (See [8] for this classical result and references.) Hence, for these lattices, (a.s.) there is an n with $\tau(n) > p_c$. This, together with (2), implies that (a.s.) $\tau(n)$ achieves its maximum (and that this maximum is larger than p_c). The following proposition is about the invaded region at the step where this maximum is achieved. Although this and related regions have been under consideration before in the literature (see the subsection ‘Ponds and outlets’ in [17]), this result is, as far as we know, new.

Remark. *The invasion basin of O is defined similarly to the invasion tree, except that at every step, the edge of minimal τ -value among the edges outside the current invasion basin that have at least one endpoint in the basin is added to the basin. The invasion basin is typically not a tree. It is easy to see that each edge e in the invasion tree is in the invasion basin, and the set of sites in the invasion basin immediately before such an edge e is added to it is the same as the set of vertices in the invasion tree immediately before e is added.*

Proposition 1.3. *Consider invasion percolation on the square lattice (or the triangular or the hexagonal lattice) with edge values τ_e . Let \hat{e} be the edge with maximal τ value in the invasion basin (as explained above). Let \hat{R} be the radius of the region that has been invaded up to the step where \hat{e} is invaded. We have:*

$$\begin{aligned} \text{(a)} \quad & P(\hat{R} > n) \geq P_{\text{cr}}(O \leftrightarrow \partial B(n)); \\ \text{(b)} \quad & P(\hat{R} > n) \approx P_{\text{cr}}(O \leftrightarrow \partial B(n)), \quad n \rightarrow \infty. \end{aligned} \tag{3}$$

Remarks.

(a) Proposition 1.3 has triggered further research on the comparison of ponds and critical percolation clusters: see recent refinements and generalizations in [5].

(b) The value \hat{R} above can also be described in the following, somewhat informal, way. Suppose each edge e is closed at time 0 and becomes open at time τ_e (after which it remains open). The open cluster of O grows in time. Up to time p_c it is finite, but at some time larger than p_c it will become infinite (a.s.). The radius of this cluster just before it becomes infinite is \hat{R} .

1.3. Description of the model in terms of passage times. Autonomous regions

Consider the description of the dynamics in the beginning of this section, and assume for the moment that the graph is finite. Recall that an open edge remains open and that a closed edge with at least one green end-vertex becomes open at rate 1. This means that if we assign to each edge e an exponentially distributed (mean 1) random variable $\tau(e)$, independent of the other edges (and of the initial colours of the vertices), the time evolution of the process can be completely described in terms of the initial colours of the vertices and the τ -variables of the edges: Each edge e remains closed until the time t at which $L_t(e)$ (defined below) has Lebesgue measure τ_e . (If no such time exists, the edge remains closed forever). Here L_t is defined by

$$L_t(e) = \{s < t: e \text{ has at least one green end-vertex at time } s\}. \tag{4}$$

(Since, once a vertex is green it can change colour only one more time, $L_t(e)$ is clearly an interval or union of two intervals.) When e becomes open and one of its end-vertices is white or red, the appropriate action in the description in Section 1.1 is carried out instantaneously.

In the following this equivalent description of the process turns out to be very convenient. To illustrate it and to emphasize the difference with one of the modified models that will be discussed in Section 1.4, we give the following example.

Example 1.4. Consider the graph with vertices denoted by $\{1, 2, 3, 4, 5\}$ and edges $\langle i, i + 1 \rangle, 1 \leq i \leq 4$. Suppose that the initial colours of the vertices $1, \dots, 5$ are red, green, white, green, red respectively, and that the τ values of the edges $\langle 1, 2 \rangle, \dots, \langle 4, 5 \rangle$ are 6, 3, 4 and 2 respectively. As one can check by following the above description, the initially green vertex 2 becomes red at time 5 due to vertex 5.

Now suppose some finite, but possibly large, graph G is given, together with initial colours $c(v), v \in V$, and ‘opening times’ $\tau(e), e \in E$. Further suppose we are only interested in the time evolution in a small subgraph of G , for instance just one initially green vertex v . Do we need to ‘follow’ the process in the whole graph to reconstruct what happens at v ? Often this is not the case. An instructive example is when v is incident to three edges, e, e' and e'' with the properties that $\tau(e)$ is smaller than $\tau(e')$ and $\tau(e'')$, and that the other end-vertex of e , which we denote by w , is red. In that case we know that v is green until time $\tau(e)$ and from then on is red (which would also happen in the ‘isolated’ graph consisting only of the vertices v and w and the edge e). This holds no matter what the initial colours of the vertices in $V \setminus \{v, w\}$ and the τ -values of the edges in $E \setminus \{e, e', e''\}$ are. Note that this still holds when we extend G to a bigger graph (with c - and τ -variables) as long as we don’t add extra edges to v .

This brings us to the notion of *autonomous set*: Let $H = (V(H), E(H))$ be a finite sub-graph of a graph G , and let \bar{E} be a finite set of external edges of H , i.e. edges in G , which have exactly one vertex in $V(H)$. Assume that we have given an initial colour assignment $c(v)$ to all $v \in V(H)$ and opening times $\tau(e)$ to all $e \in E(H) \cup \bar{E}$. Let \bar{H} be the minimal graph containing H as subgraph and $\bar{E} \subset E(\bar{H})$. We say that (H, \bar{E}) is *autonomous* (with respect to τ and c), if for every finite subgraph G_0 of G which has \bar{H} as a subgraph, the growth process on G_0 starting with a colour pattern and opening times extending the above given c ’s and τ ’s has, restricted to H , always the same time evolution, i.e. the same evolution as it would have with $G_0 = \bar{H}$, and which does not depend on colours at the vertices in \bar{H} not in H . In the simple example considered in the previous paragraph, the graph with vertices v and w , and edge e , together with the set of external edges $\bar{E} = \{e', e''\}$, is autonomous.

Often, when the identity of \bar{E} is obvious and the choice of c - and τ -variables is considered to be known, we simply say that H is autonomous. For this reason we might refer to the autonomous set as “autonomous subgraph.”

Now suppose we have an infinite graph G with given τ - and c -variables. If every vertex (and every edge) is contained in a finite autonomous subgraph of G , the infinite-volume time evolution on G can be defined in an obvious way. The key of the proof of Theorem 1.1 is to show that, under the condition in the theorem, these autonomous subgraphs exist with probability 1. That is, for almost-all initial colour patterns, and almost-all τ -values each vertex and edge is contained in a finite autonomous region.

1.4. Some alternative versions of the model

There are many modifications or generalizations of our model (which we will sometimes call the *basic model* to distinguish it from these modified versions). Below we mention four of them.

(i) In the basic model the τ variables are exponentially distributed. It is easy to see that if the initial colours of the vertices are given, and none of them is white, the time evolution is essentially determined by the order statistics of the τ variables. It is also easy to see that in that case each edge e becomes open at time τ_e or remains closed forever. From such observations it easily follows that, if $p_w = 0$, replacing the exponential distribution of the τ variables by some other continuous distribution, leaves the law of the process unchanged, apart from an obvious time change. This is not true if $p_w > 0$. However, as one can easily see from its proof, Theorem 1.1 remains valid under such replacement of distribution.

(ii) Recall that in our basic model an edge e becomes open at the smallest time t with the property that the subset of times $s < t$ at which e has at least one green end-vertex, has Lebesgue measure τ_e . A natural modification of this rule is the one where $e = \langle v, w \rangle$ becomes open at the smallest time t with the property that v is green throughout the interval $[t - \tau_e, t)$ or w is green throughout the interval $[t - \tau_e, t)$. To illustrate the difference between the rules, consider again the graph with τ values and initial colours in Example 1.4. As can be easily checked, under the modified rule the vertex 2 will no longer become red due to vertex 5 but due to vertex 1 (and at time 6 instead of 5). It turns out that Theorem 1.1 remains valid for this modified model and that its proof only needs some small modifications.

(iii) The third modification is the following model in continuous space. Consider two homogeneous Poisson point processes ζ_G, ζ_R on \mathbb{R}^d , with intensities $\lambda_G = 1, \lambda_R \equiv \lambda \in (0, +\infty)$ respectively. The points of ζ_G (*green*) are interpreted as sources of growth, and those of ζ_R (*red*) as sources of ‘paralyzing poison.’ All other elements of \mathbb{R}^d are uncoloured. From each source in ζ_G at time zero a green Euclidean sphere begins to grow with constant speed 1 (of its radius). When two or more green spheres intersect, they keep growing in the same manner, but we say that they have become connected (are in the same connected green component). If a growing green sphere hits a red region, its *entire* connected green component (note that this is a union of spheres) instantaneously gets red and stops growing. Analogs of the questions for our basic model in Section 1.1, in particular the existence question, arise naturally, but so far we have made very little progress. Although at first sight there is some resemblance with the model studied in [9], the arguments used there seem not to work here.

(iv) Consider the following change of rule of the previous model (model (iii)): When a green sphere hits a red region, *only* the centers of all the spheres of its connected green component become red; the remaining parts of the spheres disappear (become uncoloured). This change makes the model much easier to handle (using an invasion procedure resembling the one we will use in Section 2 for the case $p_w = 0$ of our basic model), but also considerably less natural, and we will not discuss it in more detail.

2. Proofs for the case $p_w = 0$

2.1. General properties for the case $p_w = 0$

The case where $p_w = 0$ is considerably easier than the general case and serves as a good introduction to the latter. We start with some deterministic observations and claims. Let us first restrict to a finite graph G , with given τ -values and c -values. We assume that at least one vertex has initial colour red, at least one vertex has initial colour green, and no vertex has initial colour white. Let x be a vertex with initial colour green, and let $t(x)$ denote the time at which x becomes red. Let Π denote the set of all paths of which the starting point is x and the end-vertex has initial colour red. It is easy to see that

$$t(x) \geq \min_{\pi \in \Pi} \max_{e \in \pi} \tau(e). \quad (5)$$

Indeed, for each t smaller than the r.h.s. of (5) there is a ‘cut set’ of edges that are still closed at time t and ‘shield’ x from all initially red vertices. It is also quite easy to see that equality holds in (5). The algorithmic (and inductive) argument below is not the most direct one but has the advantage that it gives more, namely an elegant and suitable construction of an autonomous region. This particular construction will almost immediately lead to a proof of parts (a) and (b) of Theorem 1.1 for the case $p_w = 0$. The ‘algorithm’ is a modification (‘stopped’ version) of the standard invasion percolation procedure (starting at x) described a few lines above Proposition 1.3. At each stage of the procedure we have a tree which is a subgraph of G . Initially this tree consists only of the vertex x . At each step we consider all edges that have exactly one end-vertex in the tree, also called the *external edges* of the tree. Among these edges we select the one with minimal τ -value and add it (and its external end-vertex) to the tree. The procedure is

stopped as soon as an initially red vertex is added to the tree. Let us denote this vertex by R , and the final tree given by this procedure by $T(x)$. Let τ^* be the maximal τ -value on this tree, and e^* the edge where this maximum is attained. Removing this edge from the tree $T(x)$ ‘splits’ the tree in two parts. Let $T_1^*(x)$ denote the part containing x .

Claim 2.1. (i) *The vertex R is responsible for x becoming red.*

(ii) *x becomes red at time τ^* . That is, $t(x) = \tau^*$. Moreover, $C_g(x)$ (defined in part (b) of Theorem 1.1) is the set of vertices of $T_1^*(x)$.*

(iii) *Let \bar{E} denote the set of all edges of which one end-vertex is a vertex of $T(x)$, different from R , and one end-vertex is not in $T(x)$. Let $\widehat{T}(x)$ be the graph with the same vertices as $T(x)$ and with all edges that have both end-vertices in $T(x)$. Then $(\widehat{T}(x), \bar{E})$ is autonomous (with respect to this coloring).*

Proof. The proof of the claim is by induction on the number of steps in the above invasion procedure. If the number of steps is 1 we are in the situation that the edge incident to x with minimal τ -value has a red end-vertex, and the above claim follows easily. (Note that this case corresponds with the example in the second paragraph of Section 1.3.) Now suppose the number of steps is larger than 1. Consider the edge e^* defined above. Let E^* denote the set of external edges, except e^* itself, at the stage of the procedure immediately before e^* was added. From the definition of invasion percolation, all edges in E^* have τ -value larger than τ^* . On the other hand, all edges that were added after that step have, by definition, τ -value smaller than τ^* . Therefore the edges in E^* were never added to the tree. Hence, since R was added after e^* (and was the first red point added to the tree), it follows that every path in G from x to a red point contains e^* or an edge in E^* . Therefore, by (5) we get that

$$t(x) \geq \tau^*.$$

To get the reversed inequality, note the following. Let y denote the external end-vertex of e^* when e^* was added to the tree. We already remarked that removing e^* from $T(x)$ ‘splits’ $T(x)$ in two separate trees, and we denoted the part containing x by $T_1^*(x)$. Let $T_2^*(x)$ denote the other part. It follows from the above that $T_2^*(x)$ contains y and R . We will assume that the initial colour of y is green (otherwise the claim follows easily). It is easy to see from the above that a similar invasion procedure as before, but now starting at y instead of x , has as its final tree the tree $T_2^*(x)$. By the induction hypothesis we have that y becomes red at the time which is equal to the maximal edge value in $T_2^*(x)$ and hence before time τ^* , and that R is responsible for y becoming red. Also note that, from the earlier observations, just before time τ^* there is an open path from x to the end-vertex $\neq y$ of e^* . Since e^* becomes open at time τ^* it follows that x becomes red at time τ^* . Moreover, since R is responsible for y becoming red, it is also responsible for x becoming red. This (and the earlier made observation that all external edges $\neq e^*$ of $T_1^*(x)$ have τ value larger than τ^*) completes part (i) and (ii) of the proof of Claim 2.1. Similar arguments show part (iii). \square

Now we are ready to handle the case where G is infinite. If G is infinite and $p_r > 0$, it is not a priori clear that the process described in Section 1.1 is well defined. However, the above invasion procedure and the corresponding claim now give us the instrument to define it and to give a proof of Theorem 1.1 in this particular case.

2.2. Proof of Theorem 1.1 for the case $p_w = 0$

For each green vertex x simply run the invasion procedure starting from x . Since the initial colours and the τ variables are independent, we have, at each step in the invasion from x , probability p_r of hitting a red vertex (independently of the previous steps in this invasion). Hence the invasion procedure starting at x stops with probability 1, and (by part (iii) of Claim 2.1) yields an autonomous region containing x . Since the graph has countably many vertices, this yields a construction of the process on G and completes the proof of part (a) of the theorem. Moreover it shows that Claim 2.1 also holds (a.s.) for G . Further, the number of steps in the invasion procedure from an initially green vertex clearly has a geometric distribution: the probability that it is larger than n is $(1 - p_r)^n$. Since (by part (ii) of Claim 2.1) $|C_g(v)|$ is at most the number of steps in the invasion procedure, part (b) of the theorem follows.

Proof of part (c). For each pair of vertices x, y , let $I(x, y)$ denote the event that x is initially green and that y is initially red and responsible for x becoming red. It follows immediately from the above that for all vertices x and all m

$$\sum_{y: d(x,y) \geq m} P(I(x, y)) = P(d(x, R(x)) \geq m) \leq (1 - p_r)^m. \tag{6}$$

Further, using that G is a Cayley graph, the ‘mass transport principle’ (see e.g. Section 7.1 in [15] or [10]) gives:

$$P(D(w) \text{ has radius } \geq m) \leq \sum_{v: d(v,w) \geq m} P(I(v, w)) = \sum_{v: d(v,w) \geq m} P(I(w, v)),$$

which by (6) is at most $(1 - p_r)^m$. This completes the proof of part (c) of the theorem. □

Proof of part (d). As we will see, this follows from earlier observations, together with a block argument which is quite similar to one in percolation theory, due to Kesten (see [13]). Below we denote the d -dimensional cubic lattice simply by \mathbb{Z}^d .

Let, as before, $T(x)$ denote the tree produced by the invasion procedure starting at x . We want to prove exponential decay for $P(|D(v)| > n)$, where v is an initially red point. Without loss of generality we take $v = \mathbf{0}$. We say that a finite set W of vertices containing $\mathbf{0}$ is a lattice animal (abbreviated as l.a.) if for all $w \in W$ there is a path in \mathbb{Z}^d from $\mathbf{0}$ to w of which every vertex is in W . From the definitions (and since, as we saw in (c), $D(\mathbf{0})$ is a.s. finite), it is clear that $D(\mathbf{0})$ is a lattice animal. Let L be an even integer and consider the partition of \mathbb{Z}^d into cubes $Q_L(x) := [-L/2, L/2)^d + Lx$, $x \in \mathbb{Z}^d$. We say that $x \in \mathbb{Z}^d$ is *fine* if $Q_L(x) \cap D(\mathbf{0}) \neq \emptyset$. Let V_F denote the set of all vertices that are fine. Since $D(\mathbf{0})$ is a lattice animal, V_F is also a lattice animal. Further, we say that $x \in \mathbb{Z}^d$ is *proper* if $Q_L(x)$ contains a vertex y with $|T(y)| > L/4$, and write $I(x \text{ is proper})$ for the indicator function of the corresponding event. (Here $T(\cdot)$ is as defined in the invasion procedure earlier in this section.) Finally, a subset of \mathbb{Z}^d is *proper* if every element in the set is proper. It is clear that for every $x \neq \mathbf{0}$, if x is fine, then x is proper. It is also clear that if $D(\mathbf{0})$ contains vertices outside $Q_L(\mathbf{0})$, then $\mathbf{0}$ is also proper. Recall from Claim 2.1(iii) that for each tree T in \mathbb{Z}^d and each vertex y , the event $\{T(y) = T\}$ depends only on the c -values of the vertices of T and the τ -values of the edges that have at least one end-vertex in T . From this it easily follows that the process $(I(x \text{ is proper}), x \in \mathbb{Z}^d)$ is 2-dependent (see e.g. [8] for this notion). Let $\varepsilon = \varepsilon(L) = \varepsilon(L, d)$ be the probability that a given vertex is proper. Since, for each y , the distribution of $|T(y)|$ is geometric (and $|Q_L(y)|$ is polynomially bounded in L) it is clear that for fixed d

$$\varepsilon(L, d) \rightarrow 0 \quad \text{as } L \rightarrow \infty.$$

The above mentioned 2-dependence gives that there is a constant $C_1 = C_1(d)$ such that for every set $W \subset \mathbb{Z}^d$

$$P(W \text{ is proper}) \leq \varepsilon^{|W|/C_1}. \tag{7}$$

Finally, we use that there is a constant $C_2 = C_2(d)$ such that the number of lattice animals of size m is at most C_2^m , see [8]. Together, the above gives that (noting that each l.a. of size $\geq m$ contains a l.a. of size m) for n large enough (depending on L),

$$\begin{aligned} P(|D(\mathbf{0})| > n) &\leq P\left(\exists \text{ a proper l.a. of size } \left\lceil \frac{n}{|Q_L|} \right\rceil\right) \\ &\leq C_2^{n/|Q_L|+1} \varepsilon(L)^{n/(|Q_L|C_1)} \\ &= C_2 \left[(C_2 \varepsilon(L))^{1/C_1} \right]^{1/|Q_L| n}. \end{aligned} \tag{8}$$

Taking L so large that $C_2 \varepsilon(L)^{(1/C_1)} < 1$ completes the proof of part (d). □

This completes the proof of Theorem 1.1 for the special case where $p_w = 0$.

2.3. Proofs of Proposition 1.3 and Theorem 1.2

We first prove Proposition 1.3. We say that an edge is p -open if $\tau_e < p$. Define p -open paths and p -open clusters in the obvious way. To prove the proposition we will derive suitable lower and upper bounds for the l.h.s. of (3) in terms of an expression of the form of its r.h.s.

The lower bound is very easy: Since $\tau_{\hat{e}} > p_c$ (see the paragraph below (2)), it follows immediately that (a.s.) the region which is already invaded at the step where \hat{e} is invaded, contains all the vertices of the p_c -open cluster of O . Hence the l.h.s. of (3) is larger than or equal to the r.h.s.

The upper bound is more complicated. We use the standard percolation notation $\theta(p)$ for the probability that O is in an infinite p -open cluster.

Define, for each p and n , the following two events:

$$A_{n,p} = \{\exists \text{ a } p\text{-closed circuit with diameter } \geq n \text{ in the dual lattice that contains } O \text{ in its interior}\}.$$

$$D_p = \{O \text{ belongs to an infinite } p\text{-open cluster}\}.$$

Note that $P(D_p) = \theta(p)$ and that if $p_1 < p_2$, then $D_{p_1} \subset D_{p_2}$ and $A_{n,p_2} \subset A_{n,p_1}$.

Let $\hat{\tau} = \tau_{\hat{e}}$. Let p' be some number between p_c and 1. The following observation is straightforward.

Observation.

- (a) If $\hat{\tau} > p'$ and $\hat{R} \geq n$, then there is a $p > p'$ such that the event $A_{n,p}$ occurs.
- (b) If $\hat{\tau} < p'$, then there is a $p < p'$ such that D_p occurs.

Let, for $p > p_c$, $L(p)$ be the correlation length ($=L(p, \varepsilon_0)$) as defined in Section 1 in the paper by Kesten [14] on scaling relations. (See [14].) That is, $L(p)$ is the smallest n such that the probability that there is a p -open horizontal crossing of a given $n \times n$ box is larger than $1 - \varepsilon_0$. Here ε_0 is an appropriately (sufficiently small) chosen positive constant. (From this definition it is clear that $L(p)$ is non-increasing in p on the interval $(p_c, 1]$.) It is well known (see (2.25) in [14] and the references preceding that equation) that there are constants $C_1 > 0$ and $C_2 > 0$ such that for all $p > p_c$ and all n ,

$$P_p(A_{n,p}) \leq C_1 \exp\left(-\frac{C_2 n}{L(p)}\right). \tag{9}$$

Further, Theorem 2 in [14] says that there is a constant $C_3 > 0$ such that, for all $p > p_c$,

$$\theta(p) \leq C_3 P_{\text{cr}}(O \leftrightarrow \partial B(L(p))). \tag{10}$$

Now take, for p' , the supremum of those p for which $L(p) > n/(C_4 \log n)$, where C_4 is a positive constant that will be appropriately chosen later. Obviously,

$$P(\hat{R} \geq n) \leq P(\hat{R} \geq n, \hat{\tau} > p') + P(\hat{\tau} < p'). \tag{11}$$

The first term in the r.h.s. of (11) is, by Observation (a) and the ‘nesting’ property of the events $A_{n,p}$ (stated in the sentence below the definition of these events), smaller than or equal to

$$\lim_{p \downarrow p'} P(A_{n,p}) \leq \limsup_{p \downarrow p'} C_1 \exp\left(-\frac{C_2 n}{L(p)}\right) \leq C_1 \exp(-C_2 C_4 \log n), \tag{12}$$

where the first inequality follows from (9) and the second inequality from the definition of p' .

The second term of (11) is, by Observation (b) and the ‘nesting’ property of the events D_p , smaller than or equal to

$$\lim_{p \uparrow p'} \theta(p) \leq \limsup_{p \uparrow p'} C_3 P_{\text{cr}}(O \leftrightarrow \partial B(L(p))) \leq C_3 P_{\text{cr}}\left(O \leftrightarrow \partial B\left(\frac{n}{C_4 \log n}\right)\right), \quad (13)$$

where the first inequality follows from (10) and the second follows by (again) using the definition of p' . Putting (11), (12) and (13) together we have

$$P(\hat{R} \geq n) \leq C_3 P_{\text{cr}}\left(O \leftrightarrow \partial B\left(\frac{n}{C_4 \log n}\right)\right) + C_1 \exp(-C_2 C_4 \log n). \quad (14)$$

It is believed that $P_{\text{cr}}(O \leftrightarrow \partial B(n))$ has a power law behaviour. This has only been proved for site percolation on the triangular lattice. However, for the percolation models we are considering, we do know that this function of n has power-law lower and upper bounds. As a consequence we can choose C_4 so large that the second term in the r.h.s. of (14) is (for all large enough n) smaller than the first term. Finally, it follows quite easily from RSW arguments (see e.g. Sections 11.7 and 11.8 in [8]) that $P_{\text{cr}}(O \leftrightarrow \partial B(n/C_4 \log n)) \approx P_{\text{cr}}(O \leftrightarrow \partial B(n))$. This completes the proof of Proposition 1.3.

Now we are ready to prove Theorem 1.2. The invasion procedure in Section 2.1, which was used in the proof of Theorem 1.1, differs from the ‘ordinary’ invasion percolation model (described in the paragraphs preceding Proposition 1.3), in that it stops as soon as the growing tree ‘hits’ a red vertex. There is strictly speaking another difference: the τ -values in the former case were exponentially distributed and those in the latter case were uniformly distributed on $(0, 1)$. However, that difference clearly does not matter, and in the rest of this proof we assume the τ variables to be uniformly distributed on $(0, 1)$. Let us call the former procedure a ‘stopped’ invasion procedure (with parameter p_r), and the latter an ordinary invasion procedure. All these procedures (the stopped procedures with p_r varying between 0 and 1, and the ordinary procedure) can be coupled in the following natural way: Assign to each vertex v , independent of the others, (and of the τ variables) a random variable $\rho(v)$, uniformly distributed on the interval $(0, 1)$. When we now do invasion percolation (w.r.t. the τ variables) and stop when we hit a vertex with ρ value smaller than p_r , this corresponds exactly with the above mentioned stopped invasion with parameter p_r . In this coupled setting, the set $C_g(O)$ for the stopped model with parameter p_r is clearly non-increasing in p_r , and the union of these sets over all the values $p_r > 0$ is exactly the region mentioned in Proposition 1.3. Theorem 1.2 now follows from this proposition.

3. Proof for the case $p_w > 0$

In this section we prove Theorem 1.1 for the case $p_w > 0$. Recall that in the special case where there are no white vertices (see Section 2) there was an elegant invasion procedure which produced, with probability 1, a finite autonomous set containing a given vertex or edge. This is much more complicated in the general case, when there are white vertices. We still have a procedure which, if it stops, gives an autonomous set containing, say, a given vertex x . This algorithm starts as before, with one invasion tree, which initially consists only of the vertex x , and which grows by invading the edge with minimal τ value. However, when we hit a ‘fresh’ white vertex y we have to investigate the ‘space–time paths from outside’ that have possibly influenced y . This is done by starting new invasion trees in the green vertices on the boundary of the white cluster of y . As before, an invasion tree stops when it invades a red vertex. In the situation in the previous section this also marked the end of the algorithm. But in the current situation it only marks the end of one invasion tree, while the others keep growing and creating new invasion trees. In this way the algorithm might go on forever. However, we show that under the condition in Theorem 1.1 the algorithm, which is described more precisely below, does end.

The input is a connected graph $G = (V, E)$, the initial colours $c(v)$, $v \in V$, and the opening times $\tau(e)$, $e \in E$, and the vertex x or edge e for which we want to find an autonomous region. Here we only handle the case concerning a vertex x and we assume that x is green; the other cases can be done in a very similar way. For the moment it suffices to restrict to finite graphs. The algorithm will produce an autonomous subgraph H and, for some vertices v of H , non-negative numbers $t_g(v)$ and $t_r(v)$, and for some edges e of H a positive number $t(e)$. Here $t_g(v)$ and $t_r(v)$ will denote the time at which v becomes green and red, respectively. The value $t(e)$ will be the time when e becomes open.

It will be clear from the description below that, at each stage of the algorithm the edges to which a t -value has been assigned form a collection of disjoint trees. Each tree in this collection has one of two labels: ‘active’ or ‘paralyzing.’ How these labels are assigned is described in Section 3.1. The collection of active trees is denoted by \mathcal{T}_a and the collection of paralyzing trees by \mathcal{T}_p . As we will see, new active or paralyzing trees are ‘created’ during the algorithm, and active trees can merge with each other or with a paralyzing tree. In the former case the new tree is active, in the latter case it is paralyzing.

The set of edges which have at least one end-vertex in an active tree (and not both end-vertices in the same active tree) is denoted by \mathcal{E} . With some abuse of terminology we say that a vertex is in \mathcal{T}_a if it is a vertex of some tree in \mathcal{T}_a . A similar remark holds w.r.t. \mathcal{T}_p .

Apart from the above, we need the following auxiliary variables and structures, which will be assigned during the algorithm.

The first auxiliary structure we mention here is a set S , which can be interpreted as the set of all initially white vertices that ‘have been seen until the current stage’ in the algorithm. We say that a vertex ‘is registered’ if it is in \mathcal{T}_p , \mathcal{T}_a or S . Further, to each edge $e \in \mathcal{E}$ (as introduced above) a value $t_1(e)$ will be assigned, which can be interpreted as a tentative, possible value for $t(e)$.

Finally, the following definition will be important: The *white cluster* $C_w(v)$ of a vertex v is defined as the maximal connected subset of G of which all vertices y have initial colour $c(y) = \text{white}$. Note that this notion, in contrast with the notion of green clusters (defined in Section 1) does not involve the state (open/closed) of the edges. The boundary of the white cluster of v , denoted by $\partial C_w(v)$, is the set of all vertices that are not in $C_w(v)$ but have an edge to some vertex in $C_w(v)$. If $c(v)$ is not white, then $C_w(v)$ and $\partial C_w(v)$ are empty.

3.1. Description of the algorithm

Using the notions above we are now ready to describe the algorithm. It starts with action 1 below, followed by an iteration of (some of) the other actions. Recall that $c(x)$ is green.

1. Initialization of some of the variables and structures.

Set $\mathcal{T}_p = \emptyset$, $\mathcal{T}_a = \{\{x\}\}$ and $S = \emptyset$.

Set $t_g(x) = 0$, \mathcal{E} as the set of all edges incident to x , and $t_1(e) = \tau(e)$ for all edges $e \in \mathcal{E}$.

2. Selection of minimal external edge.

Remove from \mathcal{E} all edges of which both endpoints are in the same tree of \mathcal{T}_a .

Comment: such edges can have resulted from some of the actions below.

If $\mathcal{E} = \emptyset$, stop. Otherwise, let e be the edge in \mathcal{E} with minimal t_1 -value.

Write $e = \langle v, y \rangle$ with v in \mathcal{T}_a . (This way of writing is of course not unique if both end-vertices of e are in \mathcal{T}_a but that doesn’t matter.) Let T denote the tree in \mathcal{T}_a of which v is a vertex.

If y is not in \mathcal{T}_a , \mathcal{T}_p or S (that is, y is ‘fresh’) go to 2a, else go to 2b.

2a. Fresh vertex.

Determine $c(y)$.

If $c(y) = \text{red}$, set $t(e) = t_1(e)$ and go to 3a.

If $c(y) = \text{green}$, set $t(e) = t_1(e)$ and go to 4.

If $c(y) = \text{white}$, go to 6.

2b. Registered vertex.

Set $t(e) = t_1(e)$.

If y is in \mathcal{T}_p go to 3b.

If y is in \mathcal{T}_a go to 5.

Else go to 7.

3a. Fresh red.

Comment: This case can be handled in almost the same way as 3b below and therefore, with an ‘administrative trick,’ we simply turn it into the latter case:

Set $t_r(y) = 0$. Add to \mathcal{T}_p the tree which consists only of the vertex y .

Go to 3b.

3b. Active tree T becomes paralyzed.

Set $t_r(z) = t(e)$ for all vertices z of T .

Remove from \mathcal{E} all edges of which one end-vertex is in T and the other end-vertex is not in \mathcal{T}_a . Let T' be the tree in \mathcal{T}_p of which y is a vertex. Replace, in \mathcal{T}_p , the tree T' by that obtained from ‘gluing together’ T and T' via the edge e . Remove T from \mathcal{T}_a .

Go to 2.

4. Fresh green.

Set $t_g(y) = 0$. For each edge e' incident to y that was not yet in \mathcal{E} : add e' to \mathcal{E} and set $t_1(e') = \tau(e')$. Replace, in \mathcal{T}_a , the tree T by a new tree obtained from gluing y to T by the edge e .

Go to 2.

5. Two active trees join.

Let $T' \in \mathcal{T}_a$ be the active tree of which y is a vertex. Replace, in \mathcal{T}_a , the trees T and T' by a new tree obtained from ‘gluing together’ T and T' with the edge e .

Go to 2.

6. Fresh white.

Add every vertex of $C_w(y)$ to S .

For each vertex z in $\partial C_w(y)$ that has $c(z) = \text{green}$ and is not in \mathcal{T}_a or \mathcal{T}_p , do the following:

Set $t_g(z) = 0$; add the tree $\{z\}$ to \mathcal{T}_a ; add to \mathcal{E} each edge e' incident to z that is not yet in \mathcal{E} , and set $t_1(e') = \tau(e')$.

For each vertex z in $\partial C_w(y)$ that has $c(z) = \text{red}$ and is not in \mathcal{T}_p , set $t_r(z) = 0$ and add the tree $\{z\}$ to \mathcal{T}_p .

Go to 2.

7. Registered white.

Set $t_g(y) = t(e)$. Replace, in \mathcal{T}_a , the tree T by the tree obtained from T by ‘gluing’ the vertex y to it by the edge e .

For each edge $e' = \langle y, z \rangle$ of y that is not in \mathcal{E} , add it to \mathcal{E} and set $t_1(e')$ as follows:

If z is in \mathcal{T}_p but $c(z) \neq \text{red}$, set

$$t_1(e') = t(e) + \tau(e') - (t_r(z) - t_g(z)), \quad (15)$$

else set

$$t_1(e') = t(e) + \tau(e').$$

Comment: The subtracted term in (15) accounts for the time that e' already had a green end-vertex. See also the Remark at the end of Section 3.2.

Go to 2.

Remark. Note that initially there is only one active tree and that new active trees are only formed in part 6 of the algorithm. Also note that initially there are no paralyzing trees; these can be formed in part 6 and in part 3a. Moreover, 3a always leads, via 3b, to the elimination of an active tree. Now consider the case that G has no vertices with initial colour white. Then the algorithm never enters part 6 (neither part 7) so that throughout the algorithm there is one active tree until a red vertex is ‘hit.’ From such considerations it is easily seen that in this case the algorithm reduces to the one described in Section 2.

3.2. Correctness of the algorithm

If G is finite the above algorithm will clearly stop. Moreover, we claim that if G has at least one vertex with initial colour red, we have the following situation at the end of the algorithm: The set of active trees \mathcal{T}_a is empty. The set \mathcal{T}_p contains one or more trees, and the vertex x is in one of them. Each of these trees has exactly one vertex with initial colour red, and this vertex is ‘responsible’ for the other vertices in that tree to become red. The following pair, (H, \bar{E}) , is autonomous: The vertices of H are the vertices in \mathcal{T}_p together with all vertices in S . The edges of H are all edges of which both end-vertices are in the above set. The set \bar{E} is the set of all edges of which one end-vertex is a vertex v of H with $c(v) \neq \text{red}$, and the other end-vertex is not in H . Further, each initially green vertex v of H becomes red at time $t_r(v)$.

The ‘correctness’ of the above algorithm (that is, the above claim) can, in principle, be proved by induction, e.g. on the number of edges. Instead of giving a full proof (which would be extremely tedious) we present the key ideas/observations ((a)–(d) below) to be used in such proof.

(a) As in many induction proofs it is useful, or even necessary, (for carrying out the induction step) to generalize the statement one wants to prove. In the current situation this generalization is as follows: In the above algorithm, information is stored in the administration when the vertices involved are ‘encountered’ by the algorithm. In particular, in action 6 a white cluster and its boundary are ‘stored’ because a vertex of the white cluster had been encountered (as endpoint of the edge selected in action 2). The same algorithm still works if at one or more stages of the algorithm such information about a white cluster (and its boundary) is stored ‘spontaneously’ (that is, without this cluster having been encountered in the sense above).

(b) The main observation for doing induction on the number of edges is the following: Let, among all edges with at least one initially green endpoint, \hat{e} be the one with minimal τ -value. Let \hat{x} and \hat{y} denote its endpoints. We may assume that \hat{x} is initially green. It is clear that the first thing that happens in the ‘real’ growth process is the opening of \hat{e} (namely, at time $\tau(\hat{e})$). It is also clear that from that moment on the growth process behaves as if starting on a graph with one vertex less, namely the graph obtained by ‘identifying’ (or gluing together) \hat{x} and \hat{y} (with an obviously assigned colour: green if $c(y)$ is white or green; red if $c(y)$ is red).

(c) To carry out the induction step it has to be shown that the algorithm has a property analogous to that for the real process described in (b) above. That this is indeed the case can be seen as follows: As long as \hat{x} and \hat{y} are not ‘registered’ in the algorithm, the algorithm behaves the same as it would behave for the graph obtained after the identification described in (b). Moreover, one can easily see from the description of the algorithm that immediately after one of these vertices is registered, the other one also is, and that they are immediately ‘attached to each other’ (by the edge \hat{e}) in the same tree.

(d) The following side remark must be added to (c) above: Suppose that $\hat{y} \in C_w(y)$ in action 6 at some stage of the algorithm. This cluster $C_w(y)$ could be larger than that in the graph obtained by identifying \hat{y} and \hat{x} . This means that in that step ‘more information is collected’ than in the situation where \hat{x} and \hat{y} would be identified from the beginning. It is exactly for this issue that the generalized algorithm (and claim) in (a) was given.

3.3. Proof of Theorem 1.1

Proof of Theorem 1.1. It follows, in the same way as in the case $p_w = 0$, that on an infinite graph the dynamics is well defined provided the algorithm stops with probability 1. We will show that, under the condition (1) in the statement of the theorem, the algorithm indeed stops. In fact, the arguments we use will give something stronger, namely Proposition 3.1, from which not only part (a) of Theorem 1.1 follows, but which we will also use to prove parts (b), (c) and (d).

Proposition 3.1. *Under the condition of Theorem 1.1, we have that, for each x , the above mentioned algorithm stops, and, moreover, the distributions of the volume and the diameter of the graph H defined above have an exponential tail.*

Proof. By the k th step of the algorithm we mean everything done by the algorithm between the k th and $(k + 1)$ th time the algorithm ‘enters’ part 2a in the description in Section 3.1. Recall that we say that a vertex is registered if it is in \mathcal{T}_a , \mathcal{T}_p or S . Let v_k be the number of registered vertices at the beginning of step k . (In particular, $v_1 = 1$.) If the algorithm is already terminated during step j for some $j < k$, we set v_k equal to the number of registered vertices at the moment of termination. Further, let y_k denote the ‘fresh’ vertex (i.e. the vertex y in part 2a of the description in Section 3.1) treated in step k of the algorithm. (In particular, y_1 is the end-vertex of the edge incident to x with minimal τ value.) Let $\eta_k = v_{k+1} - v_k$. Further, let α_k denote the net increase of the number of active trees during step k of the algorithm. If the algorithm is terminated during step k , we set $\alpha_k = -1$. (This choice is somewhat arbitrary; it is simply a suitable choice to ensure that certain statements below hold for all k .)

Note that the initial colours of the vertices are independent random variables, each being white, red or green with probability p_w , p_r and p_g respectively. It is clear from the algorithm that we may consider the colour of a vertex as ‘hidden’ until the moment the vertex becomes registered. Let \mathcal{F}_k be all information obtained by the algorithm until the beginning of step k (including the identity but not the colour of y_k).

Let $N = \min\{n: 1 + \sum_{k=1}^n \alpha_k = 0\}$. It is easy to see that if N is finite the algorithm stops during or before step N , and the number of vertices in the above defined graph H is

$$1 + \sum_{k=1}^N \eta_k. \tag{16}$$

Note that if $c(y_k)$ is white, the procedure is sent to part 6, and the newly registered vertices in step k of the algorithm are exactly the vertices of $C_w(y_k)$ and the not yet registered vertices of $\partial C_w(y_k)$; moreover, $|\mathcal{T}_a|$ increases during this step by at most the number of green vertices in $\partial C_w(y_k)$. We write *at most*, because during the remainder of step k no new active trees are created but already present active trees may disappear (which happens if the algorithm enters part 3b before it enters part 2a again).

Similarly, if $c(y_k)$ is red or green, then the only newly registered vertex is y_k itself; moreover, in the former case $|\mathcal{T}_a|$ goes down during step k by at least 1, while in the latter case it goes down or doesn't change.

For every connected set W of vertices with $|W| \geq 2$, the number of vertices in the boundary of W is at most $(D - 1)|W|$; hence, we have

$$\eta_k \leq D|C_w(y_k)| + \mathbb{I}_{\{c(y_k) \text{ not white}\}}, \tag{17}$$

$$\alpha_k \leq (D - 1)|C_w(y_k)| - \mathbb{I}_{\{c(y_k) \text{ is red}\}}. \tag{18}$$

Note that (since y_k is 'fresh') the conditional probability that $c(y_k)$ is red, white or green, given \mathcal{F}_k , is p_r , p_w and p_g respectively. Also note that, by the condition in the theorem, $p_w < 1/(D - 1)$ and hence (as is well known and easy to check) there is a $q < 1$ such that for all n and all vertices v ,

$$P(|C_w(v)| \geq n) \leq q^n. \tag{19}$$

Moreover, it is easy to see that conditioned on \mathcal{F}_k , which includes the information that y_k is a specific vertex, say y , the cluster size $|C_w(y_k)|$ is stochastically smaller than $|C_w(y)|$. Hence the bound (19) also holds (a.s.) if we replace its l.h.s. by $P(|C_w(y_k)| \geq n | \mathcal{F}_k)$. This, combined with (17) immediately gives that there is a $\gamma < 1$ such that for all k and n ,

$$P(\eta_k \geq n | \mathcal{F}_k) \leq \gamma^n. \tag{20}$$

As to the α 's, define (compare (18)), for every vertex v ,

$$\alpha(v) = (D - 1)|C_w(v)| - \mathbb{I}_{\{c(v) \text{ is red}\}}. \tag{21}$$

Let $\alpha'(v)$, $v \in V$, be independent copies of the $\alpha(v)$, $v \in V$. By a similar stochastic domination argument that led to (20), we have for all vertices v , and all positive integers k and n ,

$$P(\alpha_k \geq n | \mathcal{F}_k, y_k = v) \leq P(\alpha(v) \geq n) = P(\alpha'(v) \geq n). \tag{22}$$

And, again by (19), there is a $\lambda < 1$ such that for all n and v

$$P(\alpha'(v) \geq n) = P(\alpha(v) \geq n) \leq \lambda^n. \tag{23}$$

Further note that, for each vertex v , we have $E(|C_w(v)|) = \xi_v(p_w)$. Hence, condition (1) in Theorem 1.1 says that there is an $\varepsilon > 0$ such that for all vertices v we have

$$E(\alpha'(v)) = E(\alpha(v)) < -\varepsilon. \tag{24}$$

From (22) and the definition of the random variables $\alpha'(v)$, $v \in V$, it follows (from stochastic domination) that, for all positive integers K ,

$$P\left(\sum_{k=1}^K \alpha_k \geq 0\right) \leq \sup^* P\left(\sum_{k=1}^K \alpha'(v_k) \geq 0\right), \tag{25}$$

where we use “ \ast ” to indicate that the supremum is taken over all tuples of K distinct vertices v_1, v_2, \dots, v_K . From (23) and (24) it follows (by standard large-deviation upper bounds for independent random variables) that there is a $\beta < 1$ such that for all K and all distinct vertices v_1, v_2, \dots, v_K ,

$$P\left(\sum_{k=1}^K \alpha'(v_k) \geq 0\right) \leq \beta^K.$$

From this and (25) it follows that the distribution of N has an exponential tail.

Putting this together with (20) and (16) we get that the number of vertices in H has an exponential tail. Indeed the event that $1 + \sum_{k=1}^N \eta_k \geq n$ is contained in the union of the events $N \geq an$ and $\sum_{k=1}^{an} \eta_k \geq n$; the probabilities of these events decay exponentially in n for suitable a .

This completes the proof of Proposition 3.1. (Note that the diameter of H is at most its volume, since H is a connected graph.) \square

Parts (a) and (b) of Theorem 1.1 follow immediately from Proposition 3.1 (noting that the vertices of $C_g(x)$ belong to H).

Using Proposition 3.1, *Parts (c) and (d) of Theorem 1.1 can now be derived in the same way as in the special case $p_w = 0$ in Section 2. This completes the proof of Theorem 1.1.* \square

Remark. *For the alternative model (i) in Section 1.4, the proof of Theorem 1.1 is exactly the same. Note that the proof doesn't use that the τ 's are exponentially distributed, it applies in the same manner to any continuous distribution.*

For the alternative model (ii) the algorithm in Section 3.1 needs a few small adaptations. Apart from this the proof remains practically the same.

Acknowledgments

Two of the authors (V. S. and M. E. V.) learned about the continuum model from E. J. Neves. We thank Antal Járai for comments on Proposition 1.3 and Chuck Newman for drawing our attention to the article [17]. We also thank Ron Peled and the referees for corrections in the first manuscript.

References

- [1] D. J. Aldous. The percolation process on a tree where infinite clusters are frozen. *Proc. Camb. Phil. Soc.* **128** (2000) 465–477. MR1744108
- [2] K. S. Alexander. Percolation and minimal spanning forests in infinite graphs. *Ann. Probab.* **23** (1995) 87–104. MR1330762
- [3] I. Benjamini and O. Schramm. Private communication, 1999.
- [4] J. van den Berg and B. Tóth. A signal-recovery system: asymptotic properties, and construction of an infinite-volume process. *Stochastic Process. Appl.* **96** (2001) 177–190. MR1865354
- [5] J. van den Berg, A. Járai and B. Vágvolgyi. The size of a pond in 2D invasion percolation. *Electron. Comm. Probab.* **12** (2007) 411–420. MR2350578
- [6] J. T. Chayes, L. Chayes and C. M. Newman. Bernoulli percolation above threshold: an invasion percolation analysis. *Ann. Probab.* **15** (1987) 1272–1287. MR0905331
- [7] M. Dürre. Existence of multi-dimensional infinite volume self-organized critical forest-fire models. *Electron J. Probab.* **11** n. 21, (2006) 513–539 (electronic). MR2242654
- [8] G. R. Grimmett. *Percolation*, 2nd edition. Springer, 1999. MR1707339
- [9] O. Häggström and R. Meester. Nearest neighbor and hard sphere models in continuum percolation. *Random Structures Algorithms* **9** (1996) 295–315. MR1606845
- [10] O. Häggström, Y. Peres and R. H. Schonmann. Percolation on transitive graphs as a coalescent process: relentless merging followed by simultaneous uniqueness. In *Perplexing Problems in Probability* (M. Bramson and R. Durrett, Eds) **44** 69–90. Birkhäuser, Boston, 1999. MR1703125
- [11] A. Járai. Private communication, 1999.
- [12] A. Járai. Invasion percolation and the incipient infinite cluster in 2D. *Comm. Math. Phys.* **236** (2003) 311–334. MR1981994
- [13] H. Kesten. Analyticity properties and power law estimates in percolation theory. *J. Statist. Phys.* **25** (1981) 717–756. MR0633715
- [14] H. Kesten. Scaling relations for 2D percolation. *Comm. Math. Phys.* **109** (1987) 109–156. MR0879034
- [15] R. Lyons, Y. Peres. Probability on trees and networks. Available at <http://mypage.iu.edu/~rdlyons/>.
- [16] R. Lyons, Y. Peres and O. Schramm. Minimal spanning forests. *Ann. Probab.* **34** (2006) 1665–1692. MR2271476
- [17] D. L. Stein and C. M. Newman. Broken ergodicity and the geometry of rugged landscapes. *Phys. Rev. E* **51** (1995) 5228–5238.
- [18] D. Wilkinson and J. F. Willemsen. Invasion percolation: a new form of percolation theory. *J. Phys. A* **16** (1983) 3365–3376. MR0725616