THE CONTACT PROCESS ON TREES[1]

BY ROBIN PEMANTLE

Cornell University

The contact process on an infinite homogeneous tree is shown to exhibit at least two phase transitions as the infection parameter λ is varied. For small values of λ a single infection eventually dies out. For larger λ the infection lives forever with positive probability but eventually leaves any finite set. (The survival probability is a continuous function of λ, and the proof of this is much easier than it is for the contact process on d-dimensional integer lattices.) For still larger λ the infection converges in distribution to a nontrivial invariant measure. For any n-ary tree, with n large, the first of these transitions occurs when $\lambda \approx 1/n$ and the second occurs when $1/2\sqrt{n} < \lambda < e/\sqrt{n}$. Nonhomogeneous trees whose vertices have degrees varying between 1 and n behave essentially as homogeneous n-ary trees, provided that vertices of degree n are not too rare. In particular, letting n go to ∞, Galton–Watson trees whose vertices have degree n with probability that does not decrease exponentially with n may have both phase transitions occur together at $\lambda = 0$. The nature of the second phase transition is not yet clear and several problems are mentioned in this regard.

1. Introduction. This paper studies the contact process on trees. A tree is just a connected unoriented graph without cycles. A brief description of the contact process is as follows. The state space, in this case a tree, is any undirected graph. At any instant in time, some vertices of the graph are infected and some are healthy. The infected vertices recover (turn into healthy vertices) at Poisson rate 1, independently. Each infected vertex independently infects each of its healthy neighbors at Poisson rate λ for some fixed parameter λ. More complete descriptions of this are available in [4] and [2], along with a construction (which I will later describe) of the process from a Poisson process called the graphical representation.

There are practical and mathematical reasons to study the contact process on trees. The practical reason is that the process was developed to model the spread of infectious disease. The vertices of the graph represent individuals susceptible to the disease and the edges represent pairs of individuals who may have frequent contact. Traditionally, the graph is taken to be the n-dimensional integer lattice, probably because that makes sense for other interacting

particle systems such as the Ising model for ferromagnetism (although see [5]). While the "real" graph is a large finite graph, a tree is at least as likely as a lattice to serve as a local approximation to the real graph for a population.

One mathematical reason for studying the contact process on trees is to better understand the correspondence between behavior of processes on trees and limits of behaviors on lattices as the dimension of the lattice goes to ∞. For other processes such as random walks and the Ising model, there is a critical dimension beyond which the lattice versions of the process behave in essential ways like the tree version.

Another reason to study trees is the following observation about the contact process on different graphs. If G and H are graphs with a map ϕ from the vertices of G to the vertices of H, such that (x, y) is an edge of G whenever $(\phi x, \phi y)$ is an edge of H, then the contact process may be coupled so that the number of infections on G is always at least the number of infections on H. [Letting infections and recoveries on G proceed as on H but subject to the restriction that there can be at most one infection in any $\phi^{-1}(x)$ gives a process with equal recovery rate but smaller infection rate than the contact process on G. This process has as many infections as the process on H and, by monotonicity, has fewer infections than the contact process on G.] Since any graph has a tree for a universal cover, this relates the contact process on any graph to the contact process on a tree. In particular, among all graphs where every vertex has the same degree $n + 1$, the tree has the best survival probability for the contact process, hence the critical value λ_1 defined below is least for the tree.

The main reason that the contact process on a tree is worth studying is that it exhibits a multiple phase transition. In this respect the behavior differs remarkably from any of the usual interacting particle systems on the lattice. In fact, a major recent result [1] says that no such monkey business can happen with the contact process. That such behavior may be the rule for trees, rather than the exception, is made plausible by the recent work of Grimmett [3] and Newman and Wu [6], which shows a similar phenomenon for percolation on a graph which is the integers crossed with a tree. Since the contact process on a graph G is very similar to oriented percolation on $G \times \mathbf{Z}$, it is not surprising that unoriented percolation on the integers crossed with a tree should behave like the contact process on a tree.

The results in this paper are concerned with the existence and location of the multiple phase transitions. The theory of the contact process on a tree has by no means been tidied up. In particular, I believe, but am far from being able to prove, that there are only two phase transitions. The rest of this introductory section will define the phase transitions and delimit what I know about their existence.

Let $\mathbf{T}$ be an infinite tree with a distinguished vertex ρ called the root. Let $2^{\mathbf{T}}$ be the set of subsets of the vertices of $\mathbf{T}$, and let $\mathscr{A} \subseteq 2^{\mathbf{T}}$ be the finite subsets of vertices. The contact process can be viewed as a random variable taking values in the space of functions ξ with domain $\mathbf{R}^+$ and range $2^{\mathbf{T}}$, where $\xi(t)$ is the set of vertices infected at time t. If the initial configuration $\xi(0)$ is almost

surely in $\mathscr{A}$, then $\xi(t)$ stays in $\mathscr{A}$ almost surely for all time. I will not prove statements on this level, or even formally construct the contact process; the reader may easily do this from [2]. The idea is to construct the graphical representation, which is merely a set of independent Poisson processes. There is a rate 1 Poisson process for each vertex giving the recovery times at that vertex (a vertex ignores recovery times while it is healthy) and there is a rate λ Poisson process for each oriented edge giving the *attempted infection times*, where an attempted infection from an infected node to a healthy node transmits the infection but all other attempted infections are ignored. It is also possible to construct the process directly from its semigroup [4]. Here is some notation.

For any configuration $\eta \in 2^{\mathbf{T}}$, use "$\mathbf{P}_{\eta}(\xi(t)$ has property $X)$" to denote probabilities for the contact process started from the initial condition $\xi(0) = \eta$. Let $\mathbf{P}_{v}$ denote $\mathbf{P}_{\{v\}}$ in this regard. A measure μ is invariant for the contact process if $\xi(t)$ has law μ for any t whenever $\xi(0)$ has law μ. Since the contact process is an attractive system, it is easily seen [4] that there is a stochastically greatest invariant measure $\bar{\mu}$ that can be gotten by taking the limit of the laws of $\xi(t)$ as $t \to \infty$ when $\xi(0)$ is the configuration with every vertex infected. (The contact process, since it is an attractive system, is monotone in the following sense. Two contact processes run from initial states where the infections of one are a subset of the infections of the other may be coupled so as always to remain in this relation. This and its consequences will be used freely in what follows, and sometimes without explicit mention.) The following are all possible definitions of critical values for the infection parameter λ.

DEFINITION 1.1.

$$\lambda_1 = \inf\{\lambda : \mathbf{P}_{\rho}(\xi(t) \neq \varnothing \text{ for all } t) > 0\},$$

$$\lambda_a = \inf\{\lambda : \mathbf{P}_{\rho}(\rho \in \xi(t) \text{ for arbitrarily large } t) > 0\},$$

$$\lambda_b = \inf\left\{\lambda : \limsup_{t \to \infty} \mathbf{P}_{\rho}(\rho \in \xi(t)) > 0\right\},$$

$$\lambda_2 = \inf\left\{\lambda : \liminf_{t \to \infty} \mathbf{P}_{\rho}(\rho \in \xi(t)) > 0\right\},$$

$$\lambda_c = \inf\{\lambda : \xi(t) \text{ conditioned on not being empty } \Rightarrow_{\mathscr{D}} \bar{\mu} \neq \delta_0\},$$

where $\bar{\mu}$ is the upper invariant measure and δ_0 is the point mass at the empty set.

In words: λ_1 is the value at which the process can survive forever with nonzero probability, above λ_a the root becomes reinfected infinitely often, above λ_b the root is reinfected with probability bounded away from 0 at arbitrarily large times, above λ_2 these arbitrarily large times are all times, and above λ_c there is complete convergence to a nontrivial measure (which has been shown for one-dimensional, and more recently any dimensional integer lattice). It is immediate that $\lambda_1 \leq \lambda_a \leq \lambda_b \leq \lambda_2 \leq \lambda_c < \infty$, with the last in-

equality following from the classical results on the contact process on the integers, which are a subgraph of any infinite tree. I believe the following relations hold among the critical values.

CONJECTURE 1. *For any infinite tree, $\lambda_a = \lambda_b = \lambda_2 = \lambda_c$.*

Call a vertex v of an infinite tree $\mathbf{T}$ *essential* if $\mathbf{T} \setminus \{v\}$ has at least two infinite components. Say that a tree has strongly exponential growth if there is some $r \geq 1$ such that for any $v \in \mathbf{T}$ there are at least three essential vertices w whose graph-distance from v is r.

CONJECTURE 2. *For any tree with strongly exponential growth, either all the above critical values are 0, or else $\lambda_1 \neq \lambda_a$.*

The nature of the phase transitions other than λ_1 is evidently unclear, but regarding the first phase transition I believe the following.

CONJECTURE 3. *For any tree, the survival probability starting from a single infection is a continuous function of λ.*

Conjecture 3 is undoubtedly difficult since it includes the case of the one-dimensional integer lattice, a recent and difficult result. Theorem 6.6 says that the survival probability is continuous whenever $\lambda_1 \neq \lambda_a$, thus Conjecture 2 implies Conjecture 3 for trees of strongly exponential growth. The reason some kind of growth condition is needed in Conjecture 2 is that it is easy to make both λ_1 and λ_a equal to the critical value for the integer lattice by making the tree contain increasingly long stretches of vertices of degree 2. The theorems in this paper make some progress toward proving these conjectures, separating λ_1 and λ_a, bounding λ_a and λ_2 near each other and proving continuity of the survival probability, all in special cases. Here is a summary.

Say a tree is homogeneous of degree n if every vertex has $n + 1$ neighbors. For homogeneous trees of degree $n > 2$, Theorem 2.2 states that $0 < \lambda_1 < \lambda_a$. Additionally, it gives reasonable bounds for the location of λ_1, λ_a and λ_2. So there are at least two phase transitions when $n > 2$.

Nonhomogeneous trees come in many shapes. This paper considers two varieties, namely periodic trees and Galton–Watson trees. Since there is no single parameter n measuring the size of such a tree (the mean growth of a Galton–Watson tree being an irrelevant statistic in this setting—Theorem 3.2), bounds on the critical values are harder to get at. The main result here is that $0 = \lambda_1 = \lambda_a = \lambda_b = \lambda_2$ is possible. The last section deals with the continuity of the survival probability as a function of λ. This is linked to the separation of λ_1 and λ_a, and to the question of whether survival is still possible if a single edge is severed. The relation is: $\lambda_1 < \lambda_a \Rightarrow$ survival with a severed edge whenever survival is possible without a severed edge $\Rightarrow$ continuity of the survival probability in λ. Thus again the process behaves nicely on homogeneous trees with $n > 2$.

2. Homogeneous trees. Fix an integer $n \geq 2$. Let **T** be a homogeneous tree where each vertex (also called a node) has $n + 1$ neighbors. Pick any vertex ρ and call that the root; then each vertex other than ρ has one neighbor closer to ρ (the parent) and n neighbors (children) further from ρ. Define the contact process on **T** by letting each infected node recover at rate 1 and each infected node infect each of its neighbors at rate λ, where λ is any positive parameter. Denote by $\xi(t)$ the set of infected nodes at time $t \geq 0$. Thus ξ is viewed as a function from $\mathbf{R}^+$ to $2^{\mathbf{T}}$, or sometimes as a function from $\mathbf{R}^+ \times \mathbf{T}$ to $\{0, 1\}$. Probabilities referring to initial state η at time 0 will be written as $\mathbf{P}_\eta(\xi(t) \cdots)$. When the initial state can be understood, the subscript will be dropped.

THEOREM 2.1. *For all $n > 2$, $0 < \lambda_1 < \lambda_2 < \infty$. For $n = 2$ the same holds except that the middle inequality may not be strict.*

This is a consequence of the following theorem.

THEOREM 2.2.

(i) $$\frac{1}{n} \leq \lambda_1.$$

(ii) $$\lambda_1 \leq \frac{\sqrt{9 + 16/(n - 1)} - 1}{2(n + 1)} < \frac{1}{n - 1}.$$

Note that

$$\frac{\sqrt{9 + 16/(n - 1)} - 1}{2(n + 1)} \approx \frac{2}{3}\left(\frac{1}{n}\right) + \frac{1}{3}\left(\frac{1}{n - 1}\right).$$

(iii) $$\frac{1}{2\sqrt{n}} \leq \lambda_a.$$

(iv) $$\frac{4 + 2/(\sqrt{n} - 1) - \sqrt{8 + 16/(\sqrt{n} - 1) + 4/(\sqrt{n} - 1)^2}}{\cdot \ 2(\sqrt{n} - 1)} \leq \lambda_a.$$

Note that the LHS is asymptotically $(4 - \sqrt{8})/2\sqrt{n}$ and is always larger than (iii).

   (v) *In particular, $0.561722, 0.425516, 0.354248 \leq \lambda_2$ when n is respectively $2, 3, 4$.*

(vi) $$\lambda_2 \leq \min\left(2, \frac{4}{\sqrt{n} - 4}\right) \quad and \quad \limsup_{n \to \infty} \lambda_2\sqrt{n} \leq e = 2.71828\ldots\,.$$

(vii) $$\lambda_c \leq \max\left(\lambda_2, \frac{3 + \sqrt{8n + 1}}{2n - 2}\right),$$

$$where \ asymptotically \quad \frac{3 + \sqrt{8n + 1}}{2n - 2} \approx \frac{\sqrt{2}}{\sqrt{n}}.$$

The upper bound on λ_1 in (ii) is less than the lower bound on λ_a on (iii) except for $n \leq 4$. The reason it is worth improving the relatively easy bound (iii) to the messy bound (iv) is to separate λ_1 and λ_a for $n = 3, 4$, thus showing there are at least two phase transitions except possibly when $n = 2$. The rest of this section is taken up by the proof of Theorem 2.2.

PROOFS. (i) See Liggett [4], page 166. Alternatively, see the remark after the proof of (ii).

(ii) Assign a weight $W(\xi) = ak(\xi) + bc(\xi)$ to each configuration ξ having finitely many infected nodes, where a and b are constants to be determined later, k is the number of infected nodes of ξ, and c is the number of components in the induced subgraph of **T** on ξ, that is, the graph gotten by restricting **T** to ξ. For $\lambda > (\sqrt{9 + 16/(n - 1)} - 1)/2(n + 1)$, it will be shown that $W(\xi)$ is a submartingale, and in fact $\mathbf{E}(W(\xi(t + dt))|\xi(t)) \geq (1 + \varepsilon)W(\xi(t))$ for a suitable $\varepsilon > 0$. Together with the fact that W is a pure jump process with bounded jumps, this implies that the probability of W getting arbitrarily large before hitting 0 is bounded below. (Proof: $1/W$ is a supermartingale when W is greater than an appropriate function of ε and the bound on the jump size, so choose a starting time when W is large enough and stopping when W becomes too small, the optional stopping theorem says the stopping time is reached with probability less than 1.) Here are some preliminary results.

(A) The rate of new infections is precisely $\lambda((n - 1)k + 2c)$. To see this, count all the attempted infections and subtract off the times a node attempts to infect an already infected node. The attempted infection rate is $\lambda(n + 1)k$ and the failure rate is just λ times twice the number of edges having infected nodes on both ends. But the number of such edges is the number of infected nodes minus the number of components they constitute, so the infection rate is $\lambda((n + 1)k - 2(k - c)) = \lambda(n - 1)k + 2c$.

(B) The total recovery rate is precisely k.

(C) The rate at which new components are formed is precisely $k - 2c$, if we agree to count the component lost by the recovery of an isolated infection as -1 new components. With this convention, a new component can only be formed by a recovery, not by a new infection, and conversely, components can only be lost by the joining of previously formed components along new infections. Consider any component of r infected nodes. It contains $r - 1$ connecting edges. A recovery at a node with j edges incident creates $j - 1$ new components (this covers the trivial case $j = 0$ as well). Summing over all nodes gives $\Sigma(j - 1) = (\Sigma j) - r = 2r - 2 - r = r - 2$, where $\Sigma j = 2r - 2$ because each edge is counted exactly twice.

(D) The rate at which components are lost by joining is at most $(n + 1)\lambda(c - 1) < 2\lambda c$. To see this, consider a new, acyclic, bipartite graph, H, whose vertices are the components of the infection together with all *joining vertices*, that is, vertices of the old tree that would join two components by becoming infected. The edges of H are all edges of the tree that connect a

component of the infection to a joining vertex. Now, each joining vertex v, if infected, causes the loss of $e(v) - 1$ components, where $e(v)$ is the number of edges of H incident to v. Since each joining vertex is infected at a rate of at most $(n + 1)\lambda$, summing over all joining vertices gives a total component loss rate of at most $(n + 1)\lambda\sum_{v \in H}(e(v) - 1)$ which is just $n + 1$ times the total number of edges of H minus the total number of joining vertices, since each edge in the bipartite graph H gets counted exactly once. But the total number of edges in the acyclic graph H is just the number of vertices of H minus the number of components, which is at most c plus the number of joining vertices minus 1. The number of joining vertices cancels out and the calculation is done.

Now fix any λ, and calculate whether the infinitesimal expected change in W is always positive:

$$\mathbf{E}(W(\xi(t + dt))|\xi(t)) \geq W(\xi(t)) + \{[((n - 1)\lambda - 1)k + 2\lambda c]a$$
$$+ [k - 2c - ((n + 1)\lambda)c]b\}\,dt.$$

This is linear in c and k, so to check if it is positive for fixed λ, a, b and every c and k, it suffices to check the extreme values for c, namely $c = 1$ and $c = k$. In fact, it is easier to check $c = 0$ instead of $c = 1$. Then the two inequalities to be satisfied are

(1) $$((n - 1)\lambda - 1)a + b > 0,$$

(2) $$((n + 1)\lambda - 1)a - (1 + (n + 1)\lambda)b > 0.$$

By (i), it is safe to assume $(n + 1)\lambda - 1 > 0$ and hence to change (2) into $a > ((n + 1)\lambda + 1)b/((n + 1)\lambda - 1)$. Plugging this into (1) shows there will be a solution when

$$\frac{((n - 1)\lambda - 1)((n + 1)\lambda + 1)}{((n + 1)\lambda - 1)} + 1 > 0,$$

which happens iff $(n + 1)\lambda^2 + \lambda - 2/(n - 1) > 0$, which, for positive λ, will be true when $\lambda > (\sqrt{9 + 16/(n - 1)} - 1)/2(n + 1)$.

Finally, it is easy to see that when this inequality holds, $[\mathbf{E}(W(\xi(t + dt))|\xi(t)) - W(\xi(t))]/dt$ is at least a small constant multiple of bn for each fixed λ, thus at least $\varepsilon W(\xi(t))$ for an appropriate $\varepsilon > 0$ and the argument is done.

REMARK. This method also gives a proof of (i). Instead of the upper bound on component loss in preliminary result (D), use a lower bound of 0. The same calculation now gives $a = b/((n + 1)\lambda - 1)$ and plugging that into (1) with the inequality reversed gives $((n - 1)\lambda - 1)/((n + 1)\lambda - 1) + 1 > 0$ which is satisfiable when $\lambda < 1/n$. Thus for $\lambda < 1/n$, W can be defined to be a supermartingale and the weight must converge to 0.

(iii) This time use a weighting on all the vertices. It will be shown that the total weight of the vertices goes to 0 even if their cardinality does not. Let the weight of a vertex v be $W(v) = n^{-d/2}$, where d is the distance from ρ to v, that is, the number of edges in the shortest path from ρ to v. Let

$$M(t) = \sum_v W(v) I_{v \in \xi(t)} = \text{total weight at time } t.$$

Then

$$\mathbf{E}(M(t + dt)|\xi(t))$$

$$= \mathbf{E}\left(\sum_v I(v \in \xi(t + dt))W(v)|\xi(t)\right)$$

$$\leq \sum_v I(v \in \xi(t + dt))W(v)(-1 + \lambda n^{1/2} + n\lambda n^{-1/2})$$

$$= M(t) + dt M(t)(2\sqrt{n}\,\lambda - 1).$$

So for $\lambda < 1/2\sqrt{n}$, $\mathbf{E}M(t)$ goes to 0 exponentially fast, hence so does $\mathbf{P}(\rho \in \xi(t))$.

(iv) and (v) This time the idea is to take into account the coalescing nature of the infection in order to improve on the bound in (iii). The weight of a vertex will depend on whether its parent is infected and will count less if so, due to the redundancy. Define

$$(3) \qquad\qquad W(v) = r^k(1 - dI(\text{parent of } v \in \xi(t))).$$

where k is the distance from v to the root and r and d are parameters that will be chosen optimally later. The goal is to show that the expected change in the total weight $\sum_\mathbf{T} W(v) I(v \in \xi(t))$ is always negative. Let $\theta_1(v)$ and $\theta_2(v)$ denote respectively the number of infected parents of v (0 or 1) and the number of infected children of v. Define

$$u(v) = -1 + \theta_1(v)d + \theta_2(v)rd + \lambda[r^{-1} - d + (n - \theta_2(v))r(1 - d)].$$

Then $u(v)r^k \partial t$ is an upper bound on the expected increment in the total weight from t to $t + \partial t$ due to the recovery of v or infection of other nodes by v. [This breaks down as follows: recovery loses weight $(1 - \theta_1(v))dr^k$, but gains $\theta_2(v)dr^{k+1}$; infecting the parent gains at most r^{k-1} while losing dr^k, and infecting a child gains $(1 - d)r^{k+1}$ while it may lose something as well.]

One way to show that the expected increment in total weight is negative show that $u(v) < 0$ for each infected node. Trying to count this way is inefficient because d must be kept small so that a node that is completely surrounded by infection does not seem to increase the total weight if it recovers. A solution is to let Y be another parameter to be determined later, and define

$$(4) \qquad\qquad U(v) = u(v) + \theta_1(v)Y/r - \theta_2(v)Y.$$

The sum over all nodes of $U(v)r^k$ is the same as the sum of $u(v)r^k$, because

for each parent and child node that are both infected, a quantity of $r^k Y$ has merely been transferred from parent to child. The U's are calibrated so as no longer to penalize our strategy at nodes with many infected neighbors.

Now there are $2n + 2$ cases to check in the forward equation, one for each value of $\theta_1(v)$ and $\theta_2(v)$. For each pair of these values, r, d, λ and Y must be chosen to make $U < 0$. Again by linearity, it suffices to consider the cases where $\theta_2(v) = 0$ or n, so the inequalities are

$$-1 + \lambda\left[r^{-1} - d + nr(1 - d)\right] < 0,$$

$$-1 + nrd + \lambda(r^{-1} - d) - nY < 0,$$

$$-1 + d + \lambda nr(1 - d) + Yr^{-1} < 0,$$

$$-1 + d + nrd + Yr^{-1} - nY < 0.$$

These inequalities are not too hard to solve. Given n, r and λ, equality in the first equation gives an optimal choice for d and equality in the third equation gives an optimal choice for Y. It then turns out that equality in both of the remaining equations can be achieved for the greatest λ when $r = 1/\sqrt{n}$, and hence the strict inequalities are satisfied for any smaller λ. Plugging in these values

$$r = \frac{1}{\sqrt{n}}, \qquad d = \frac{2\sqrt{n} - \lambda^{-1}}{1 + \sqrt{n}}, \qquad Y = \left(1 - \sqrt{n}\lambda\right)\left(\frac{1 - \sqrt{n} + \lambda^{-1}}{1 + \sqrt{n}}\right)\Big/\sqrt{n}$$

gives a nasty looking quadratic in λ that simplifies to

$$-\left(1 - \sqrt{n}\right)^2\lambda^2 + \left(4\sqrt{n} - 2\right)\lambda - 2 < 0$$

and has the solution appearing in (iv) (divide through by $1 - \sqrt{n}$ and use the quadratic formula). Values for λ, r, d and Y for $n = 2, 3, 4$ are given below so that skeptics may plug them into the four equations above and check that they actually work. For $n = 2$ use $\lambda = 0.561722$, $r = 0.7071$, $d = 0.434212$ and $Y = 0.082293$; for $n = 3$ use $\lambda = 0.425516$, $r = 0.5773$, $d = 0.407781$ and $Y = 0.089933$; for $n = 4$ use $\lambda = 0.354246$, $r = 0.5012$, $d = 0.391747$ and $Y = 0.088349$.

When all the inequalities are satisfied, it is easy to see, as in the proof of (ii), that in fact $\mathbf{E}(W(\xi(t + dt))|\xi(t)) \leq (1 - \varepsilon)W(\xi(t))$ for a suitable $\varepsilon > 0$. Then again the expected total weight decays exponentially, hence the probability of finding an infection in any finite set decays exponentially and the proof is done.

(vi) The 2 comes from the fact that the critical value for λ on the one-dimensional integer lattice is known to be at most 2 and is an upper bound on λ_2 by monotonicity. For the nontrivial bound, first record two lemmas.

LEMMA 2.3. *Let* M *be a positive integer-valued random variable and pick* $p < \mathbf{E}M$ *(with* $\mathbf{E}M$ *possibly being infinite). For any* $x > 0$, *let* M_x

be a binomial(M, x) random variable. Then there is an $\varepsilon > 0$ for which $\mathbf{P}(M_x \geq 1) > px \wedge \varepsilon$.

PROOF.   Pass to the bounded case as follows: Let L be large enough so that $\mathbf{E}(M \wedge L) > p$; let $M' = M \wedge L$ and $M'_x = \text{binomial}(M', x)$; then

$$\mathbf{P}(M_x \geq 1) \geq \mathbf{P}(M_x \geq 1)$$

and it suffices to prove the lemma for M'. Now

$$\mathbf{P}(M'_x \geq 2) \leq \mathbf{P}(\text{binomial}(L, x) \geq 2) = 1 - (1-x)^L - Lx(1-x)^{L-1}$$
$$= O(x^2).$$

Also, $\mathbf{E}M'_x \leq \mathbf{P}(M'_x = 1) + L\mathbf{P}(M'_x \geq 2)$ and thus

$$\mathbf{P}(M'_x \geq 1) \geq \mathbf{P}(M'_x = 1) \geq \mathbf{E}M'_x - L\mathbf{P}(M'_x \geq 2) > px,$$

for small enough x, which proves the lemma. □

LEMMA 2.4.   Let H be any nondecreasing function on the nonnegative reals with $H(x) \geq x$ on some neighborhood of 0. Suppose f is a function on the nonnegative reals that satisfies (1) $\inf_{0 \leq t \leq L} f(t) > 0$ and (2) $f(t) \geq H(\inf_{0 \leq s \leq t-L} f(s))$ for $t \geq L$, for some $L > 0$. Then $\liminf_{t \to \infty} f(t) > 0$.

PROOF.   For any t_0 and any $\varepsilon > 0$, (2) implies there is a sequence $t_1, t_2, \ldots, t_k$ with $f(t_i) \geq H(f(t_{i-1})) - \varepsilon 2^{-i} \geq f(t_{i-1}) - \varepsilon 2^{-i}$ and $t_k \in [0, L]$. Then $f(t_0) \geq f(t_k) - \varepsilon$, which is bounded away from 0 for small ε by (1). □

Continuing the proof of (vi), let $f(t) = \mathbf{P}_\delta(\rho \in \xi(t))$ in an altered contact process where only n of the root's $n + 1$ neighbors are allowed to infect the root. The object will be to get a recursion $f(t) \geq H(f(t - 2r))$. Here $2r$ is a convenient size for a time increment to be chosen later so as to make H monotone increasing with $H(0) = 0$ and $H'(0) > 1$. Then Lemma 2.4 with $L = 2r$ will imply that $\liminf_{t \to \infty} p(t) > 0$.

To get this recursion, break down $f(t)$ into a product of conditional probabilities as follows. Let

$$H_1(t) = \mathbf{P}(v \in \xi(t-r) \text{ for some child } v \text{ of } \rho)$$

and

$$H_2(t) = \mathbf{P}(\rho \in \xi(t) | v \in \xi(t-r) \text{ for some child } v \text{ of } \rho).$$

Then $f(t) = H_1(t)H_2(t)$. Clearly a lower bound on H_2 that does not depend on t is the probability $p(r) = \mathbf{P}_x(y \in \xi(r))$ in a system with only two neighboring nodes x and y. To calculate this, let $Q(s) = \mathbf{P}_x(y \in \xi(s) | x \in \xi(u)$ for all $u < s)$. Then $Q(0) = 0$ and Q satisfies the differential equation $Q'(s) = -Q(s) + (1 - \lambda)Q(s)$ and solving the equation gives

(5) $$Q(s) = \frac{\lambda}{1+\lambda}(1 - e^{-(1+\lambda)s}).$$

Obtain from this a crude lower bound on $p(r)$ by requiring x to stay infected up to time r:

$$(6) \qquad H_2 \geq p(r) \geq Q(r)\mathbf{P}_x(x \in \xi(s) \text{ for all } s < r)e^{-r}Q(r).$$

A better bound is given by conditioning on the recovery time s of x, requiring y to be infected at time s and remain infected through time r:

$$H_2 \geq p(r)$$

$$\geq Q(r)e^{-r} + \int_0^r \mathbf{P}_x(x \text{ is infected to precisely time } s)Q(s)e^{s-r} \, ds$$

$$= e^{-r}\left[Q(r) + \int_0^r Q(s) \, ds\right]$$

$$= \frac{\lambda}{1+\lambda}e^{-r}\left[1 - e^{-(1+\lambda)r} + r - \frac{1}{1+\lambda}(1 - e^{-(1+\lambda)r})\right].$$

Now for a lower bound on H_1, let M be the number of children of ρ that are infected at time r, not counting the one child is not allowed to infect ρ. The previous calculation says essentially that each child is infected at time r with probability at least $p(r)$, so $\mathbf{E}M \geq np(r)$. If each of these nodes is not allowed to be infected by its parent between time r and time $t - r$, then by definition of f such a node infected at time r is infected at time $t - r$ with probability $f(t - 2r)$. Thus the actual number of children infected at time $t - r$ stochastically dominates a random variable that is binomial($M, f(t - 2r)$). Then by Lemma 2.3, $H_1(t) \geq np(r)f(t - 2r) \wedge \varepsilon$ for some $\varepsilon > 0$ and so $f(t) \geq np(r)^2 f(t - 2r) \wedge \varepsilon$. Now writing $H(x) \geq np(r)^2 x \wedge \varepsilon$, it remains only to show that $H'(0) > 0$, for then Lemma 2.4 will finish the proof.

For the first claim in (vi) use the crude lower bound (6) on $p(r)$ to get

$$\sqrt{H'(0)} = \sqrt{n}\, p(r)$$

$$\geq \sqrt{n}\left(\frac{\lambda}{1+\lambda}\right)e^{-r}(1 - e^{-(1+\lambda)r})$$

$$\geq \sqrt{n}\left(\frac{\lambda}{1+\lambda}\right)e^{-r}(1 - e^{-r}).$$

This holds for any r, and an optimal choice of r is $\ln(2)$, giving $\sqrt{H'(0)} \geq \lambda\sqrt{n}/4(1+\lambda)$, hence $H'(0) > 1$ if $\lambda > 4/(\sqrt{n} - 4)$.

To get the asymptotic result, use the better bound to obtain

$$\sqrt{H'(0)} \geq \sqrt{n}\, p(r)$$

$$\geq \sqrt{n}\left(\frac{\lambda}{1+\lambda}\right)e^{-r}\left(1 - e^{-(1+\lambda)r} + r - \frac{1}{1+\lambda}(1 - e^{-(1+\lambda)r})\right).$$

As $n \to \infty$ the minimum λ satisfying this inequality will have $\sqrt{n}\,\lambda \to e^r r^{-1}$ and choosing $r = 1$ gives $\limsup\sqrt{n}\,\lambda = e$.

(vii) First an outline of the argument. By duality, it suffices to show that when $\lambda > \lambda_2$ and $\lambda > (3 + \sqrt{8n + 1})/(2n - 2)$, a process that lives forever hits any big set. By duality again, this means that the process started from an infection at the root and run for a long time intersects an independent process initially infected on a sufficiently big set and run for a long time. This will be true because each of these processes will at some time infect a particular large configuration if it lives long enough (this uses $\lambda > \lambda_2$) and the configuration is chosen so that any future scenarios started from that configuration intersect with high probability. These two facts become the following two lemmas.

LEMMA 2.5. *Suppose $\lambda > \lambda_2$. Then for any finite set S and any initial configuration η, $\mathbf{P}_\eta(S \subseteq \xi(t)$ for arbitrarily large $t) = \mathbf{P}_\eta(infection\ lives\ for-ever)$.*

PROOF. Assume without loss of generality that η is finite; an easy limit argument handles the infinite case. Fix $\lambda > \lambda_2$. The first observation is that for any node v, $\mathbf{P}_v(v \in \xi(t)$ for arbitrarily large $t) = p > 0$, where $\mathbf{P}_v$ denotes $\mathbf{P}_{\{v\}}$, the contact process run from only v infected initially. This is immediate from the definition of λ_2, since if $\mathbf{P}_v(\sup\{t: v \in \xi(t)\} < \infty) = 1$, then $\mathbf{P}_v(v \in \xi(t)$ for some $t \geq T) \to 0$ as $T \to \infty$, contradicting $\limsup_T \mathbf{P}_v(v \in \xi(T)) > 0$. Next observe that $\mathbf{P}_\eta(S \subseteq \xi(t)$ for arbitrarily large $t | v \in \xi(t)$ for arbitrarily large $t) = 1$. This follows from the Markov property and the fact that $\mathbf{P}(S \subseteq \xi(t + 1) | v \in \xi(t))$ is bounded away from 0 for a fixed finite S, fixed $v \in \mathbf{T}$ and any initial state of the contact process.

The last step in proving the lemma is to show that if the process lives forever, then with probability 1 there is some v that is infected at an unbounded set of times. Assume the contrary. Then there is some configuration η' such that conditional upon the process being in configuration η' at time t, the probability is greater than $1 - p$ that the process lives forever but no node is infected at unbounded times. This contradicts the definition of p above, since η' must have at least one infected node, which is then infected at arbitrarily large future times with probability at least p. □

LEMMA 2.6. *Define a version of the contact process starting with a single infection at the root, where no vertex may infect its parent. Let $\xi(t)$ and $\xi'(t)$ be independent such processes, started simultaneously. Then $\mathbf{P}(\xi(t) \cap \xi'(t) \neq \varnothing$ for all $t) > 0$, provided*

$$\lambda > \frac{3 + \sqrt{8n + 1}}{2n - 2}.$$

PROOF. The argument will show that the set of nodes simultaneously infected in ξ and ξ' contains a supercritical branching process. This process Z

will be coupled to ξ and ξ' and for each $v \in \mathbf{T}$ the set $\{t: v \in Z(t)\}$ will be a possibly empty interval called the lifetime of v. The set of nodes that are alive is a branching process, oriented by the orientation of the rooted tree. Specifically, let ρ be the single node of Z alive at time 0. Let each node v of Z die whenever v recovers in ξ or ξ'. Thus lifetimes in the Z process end at Poisson rate 2. Dead nodes may not come back to life. A node v in the Z process is born at time t when the following conditions are met. The parent of v must be alive in the Z process, since some time $t_0 < t$. The node v must be infected in both ξ and ξ', having just received an attempted infection from its parent in one of them at time t (it is acceptable if the node was already infected when it received the attempt). Furthermore, if ξ' is the process in which the infection occurs at time t, then the most recent infection of v in the ξ process must have been after time t_0; similarly if the infection at time t occurs in the ξ process, then the most recent infection of v in the ξ' process must have been after time t_0.

Whether a node is ever alive in the Z process depends on: whether the parent was ever alive; the parent's lifetime, $t_1 - t_0$; and the attempted infections and recoveries at that node in the ξ and ξ' processes during the interval $[t_0, t_1)$. Since the distribution of attempted infections in an interval depends on the interval only through its length, and since lifetimes are i.i.d., Z is a Galton–Watson process. The mean number of offspring is easy to calculate exactly. The probability that v is ever alive given that the parent is born at time t is just the probability that v receives an attempted infection from its parent in both ξ and ξ' before its parent recovers. This is just $2\lambda/(2\lambda + 2)$ times $\lambda/(\lambda + 2)$, the first of these being the probability that even one attempt arrives in time, and the second being the race between the other attempt and the recovery of the parent. The branching process is supercritical if n times this quantity is greater than 1, which happens when $n\lambda^2 > \lambda^2 + 3\lambda + 2$, and solving the inequality gives the condition in the statement of the lemma. Finally, note that $Z(t) \subseteq \xi(t) \cap \xi'(t)$, hence $\xi(t) \cap \xi'(t) \neq \varnothing$ for all t when Z lives forever. $\square$

LEMMA 2.7. *Let η be the configuration where the infected nodes are those at distance precisely r from the root, for some $r > 0$. Let λ be as large as in Lemma 2.6, and let ξ and ξ' be two independent versions of the contact process run from initial state η. Then for any $\varepsilon > 0$, r can be chosen large enough so that $\inf_{s,t} \mathbf{P}(\xi(s) \cap \xi'(t) \neq \varnothing) > 1 - \varepsilon$.*

PROOF. By duality, $\mathbf{P}(\xi(s) \cap \xi'(t) \neq \varnothing) = \mathbf{P}(\xi((s + t)/2) \cap \xi'((s + t)/2) \neq \varnothing)$, so it suffices to consider the case $s = t$. By monotonicity, the processes dominate altered versions where no node may infect its parent. These processes proceed independently on the subtrees below each of the $(n + 1)n^{r-1}$ nodes in η, and on each subtree ξ and ξ' are versions of the process described in Lemma 2.6. By independence and the conclusion of Lemma 2.6, r can be chosen large enough so that ξ and ξ' intersect for all time on at least one of the subtrees with probability $\geq 1 - \varepsilon$. $\square$

R. PEMANTLE

PROOF OF (vii). Convergence in distribution is determined by finite-dimensional marginals, which by inclusion–exclusion are determined by the probabilities of infection existing in each finite set. So it suffices to show for finite sets $S \subseteq \mathbf{T}$ that $\mathbf{P}_\eta(S \cap \xi(t) \neq \varnothing) \to \bar\mu(S \cap \xi \neq \varnothing)$ times the survival probability of η. By duality, this is just the product of the survival probabilities of S and η.

Let $\varepsilon > 0$ be arbitrary and choose r as in Lemma 2.7. Let ξ and ξ' be independent versions of the contact process run from initial states S and η respectively. Let

$$\tau = \inf\{t \colon v \in \xi(t) \text{ for all } v \text{ with } d(\rho, v) = r\},$$

$$\sigma = \inf\{t \colon v \in \xi'(t) \text{ for all } v \text{ with } d(\rho, v) = r\}.$$

Let ξ'' and ξ''' be independent contact processes run from the initial configuration in Lemma 2.7, coupled to ξ and ξ' respectively so that $\xi''(t) \subseteq \xi(t + \tau)$ whenever $\tau < \infty$ and $\xi'''(t) \subseteq \xi'(t + \sigma)$ whenever $\sigma < \infty$. Then $\mathbf{P}(S \cap \xi(T) \neq \varnothing) = \mathbf{P}(\xi'(T/2) \cap \xi(T/2) \neq \varnothing)$ by duality, which

$$= \iint_{\infty \geq s,\, t \geq 0} \mathbf{P}(\sigma \in ds)\mathbf{P}(\tau \in dt)\mathbf{P}(\xi'(T/2) \cap \xi(T/2) \neq \varnothing | \sigma = s, t = t)$$

$$\geq \iint_{T/2 \geq s,\, t \geq 0} \mathbf{P}(\sigma \in ds)\mathbf{P}(\tau \in dt)\mathbf{P}(\xi'(T/2) \cap \xi(T/2) \neq \varnothing | \sigma = s, \tau = t)$$

$$\geq \iint_{T/2 \geq s,\, t \geq 0} \mathbf{P}(\sigma \in ds)\mathbf{P}(\tau \in dt)\mathbf{P}(\xi'''(T/2 - \sigma) \cap \xi''(T/2 - \tau) \neq \varnothing)$$

$$\geq \iint_{T/2 \geq s,\, t \geq 0} \mathbf{P}(\sigma \in ds)\mathbf{P}(\tau \in dt)(1 - \varepsilon)$$

$$= (1 - \varepsilon)\mathbf{P}(\sigma, \tau \leq T/2).$$

The last expression converges to $(1 - \varepsilon)$ times the product of the survival probabilities from S and η according to Lemma 2.5, since σ and τ are independent. Since ε is arbitrary, this shows $\liminf_t \mathbf{P}_\eta(S \cap \xi(t) \neq \varnothing)$ is at least the product of the survival probabilities. Duality gives trivially that $\limsup_t \mathbf{P}_\eta(S \cap \xi(t) \neq \varnothing)$ is at most the product of the survival probabilities, and the theorem is proved. $\square$

3. Nonhomogeneous trees.

The upper and lower bounds Theorem 2.2 gives for the critical values of λ_2 on homogeneous trees are within a constant factor. Since the constant may be quite large, it was reasonable to get some more accurate values for small n. In the case of nonhomogeneous trees, more work is needed just to get the bounds within a constant asymptotic factor, as shown by the example in the following paragraph. This section is therefore devoted to improving the bounds so that they are again within a constant factor of each other. This section does not discuss bounds on λ_1 per se, but when the trees are nonhomogeneous enough, the better upper bounds on λ_2 get pushed down by more than a constant factor below the upper bounds on λ_1 that can be obtained using the methods of the previous section, which means

that the correct value of λ_1 in these cases is also more than a constant factor lower than can be gotten by the methods of the previous section. In fact, the last example of this section shows that λ_1 and λ_2 may both be 0 on a Galton–Watson tree whose generating function has all moments. For ease of exposition, large constants are chosen with reckless abandon, leaving the interested reader the exercise of determining better constants.

To see that there is indeed a problem, try to apply the methods of the previous section to the following tree. The root has n children, each of which has only one child. Each of the n grandchildren has n children, and in general the number of children alternates between 1 and n along any direct lineage.

For an upper bound on λ_2, adapt the method used to prove (iv). Let $f(t) = \mathbf{P}(\rho \in \xi(t))$ as before, and let $g(t) = \mathbf{P}_{\mathbf{T}'}(\rho \in \xi(t))$, where $\mathbf{T}'$ is the subtree rooted at a child of the root of $\mathbf{T}$. Then the same argument as before gives $f(t) \geq H(g(t - 2r))$ and $g(t) \geq J(f(t - 2r))$, where H is the same as in the proof of (iv) and J is similar but based on only one child. Then monotonicity of H gives $f(t) \geq (H \circ J)(f(t - 4r))$. Now since $(H \circ J)'(0) = H'(0)J'(0)$, Lemma 2.4, together with (6), implies that $\lambda > \lambda_2$ whenever the quantity $Q(r)$ in (5) satisfies $e^{-4r}Q(r)^4 \cdot 1 \cdot n > 1$. This means λ must be of order $n^{-1/4}$ as n gets large. Thus if G is the geometric mean of n and 1, an upper bound on the second critical value for λ is

$$(7) \qquad\qquad \lambda_2 \leq cn^{-1/4} = c/\sqrt{G}.$$

For trees with longer periodicity, the same argument yields $\lambda_2 \leq c/\sqrt{G}$, where G is the geometric mean of the whole sequence of family sizes.

On the other hand, the somewhat ad hoc methods in (iii) and (iii') can also be adapted to this tree to give lower bounds on λ_2. The reader is invited to check that the best simple weighting scheme along the lines of (iii) for trees whose generation sizes alternate between a and b has the weight of a vertex with a children equal to $\sqrt{a}$ times the weight of each of its children, and the weight of a vertex with b children equal to $\sqrt{b}$ times the weight of each of its children (this scheme is not optimal when the periodicity of the generations is 3 or more). Under this weighting scheme, the total weight goes to 0 almost surely whenever $\lambda < 1/(\sqrt{a} + \sqrt{b})$. When $a = 1$ and $b = n$, this gives the lower bound

$$(8) \qquad\qquad 1/(1 + \sqrt{n}) \leq \lambda_a \leq \lambda_2.$$

Thus the bounds in (7) and (8) have different asymptotics as $n \to \infty$, one decreasing as $n^{-1/4}$ and the other as $n^{-1/2}$, in contrast to the case where the tree was homogeneous and the bounds were always within a contrant factor.

Which of these asymptotics for λ_2 is sharp, if either? The somewhat surprising answer is that (8) is sharp, even though (7) is clearly a better representative for the growth rate of the tree. In fact, for reasonably regular nonhomogeneous trees, the critical value of λ_2 is determined by $M =$ the (essential) supremum of the number of children of a vertex and is at most a constant times $rM^{-1/2}$, where r is a logarithmic measure of how far apart vertices with M children are from each other.

There are several kinds of nonhomogeneous trees one might wish to study. The periodic trees mentioned earlier in this section are one such class and Galton–Watson trees are another. Theorem 4.1 below will be used to obtain results about both of these cases. Since the proofs for Galton–Watson trees are harder, these will be given first and a brief version of the other will follow.

Let $\mathbf{T}$ be a Galton–Watson tree with generating function $f = \sum_{n \geq 0} a_n x^n$, that is, there is a root ρ and each vertex v has a random number $C(v)$ of children, where $C(v)$ are independent and each is equal to n with probability a_n. Assume $f'(1) > 1$ so that the tree is infinite with probability $p > 0$. To run the contact process on a Galton–Watson tree that begins with only the root infected, choose $\mathbf{T}(\omega)$ according to the Galton–Watson distribution and let $\xi(t, \omega | \mathbf{T}(\omega))$ be the contact process run on the tree $\mathbf{T}(\omega)$ starting with only the root infected.

PROPOSITION 3.1. *For a given offspring generating function f, let*

$$q(\mathbf{T}) = \liminf_{t \to \infty} \mathbf{P}(\rho \in \xi(t) | \mathbf{T}(\omega) = \mathbf{T}).$$

Then either q is almost surely 0 or else $\mathbf{P}(q > 0 |\, |\mathbf{T}| = \infty) = 1$.

PROOF. Let $p_1 = \mathbf{P}(q = 0)$. If the root has n children with subtrees $\mathbf{T}_1, \ldots, \mathbf{T}_n$, then the $\mathbf{T}_i$ are independent, identically distributed with distribution independent of n. Furthermore, if $q(\mathbf{T}_i) > 0$ for any i, then $q(\mathbf{T}) > 0$. To see this use the same argument as in (vi) of the last section for any fixed r to get

$$\mathbf{P}(\rho \in \xi(t)) \geq e^{-2r} Q^2 \mathbf{P}_{\mathbf{T}_i}(\rho \in \xi(t - 2r)).$$

Taking lim inf's over t gives $q(\mathbf{T}) \geq$ a constant times $q(\mathbf{T}_i)$ for each i. Now by independence of the subtrees $\mathbf{T}_i$, $p_1 = \mathbf{P}(q/(\mathbf{T}) \neq 0) \leq \sum_n a_n (p_1)^n = f(p_1)$. That means that either $p_1 = 1$ or p_1 is bounded above by the unique fixed point of f in $(0, 1)$. In the first case q is almost surely 0 and in the second case $\mathbf{P}(q = 0)$ is at most $\mathbf{P}(|\mathbf{T}| < \infty)$; since it is also at least $\mathbf{P}(|\mathbf{T}| < \infty)$, these two must be equal and the proposition is proved. □

In light of this, it is natural to define $\lambda_2(f)$ to be the inf of λ such that q is almost surely nonzero for the contact process with parameter λ run on a Galton–Watson tree with generating function f conditioned to be infinite. The following upper bound on λ_2 can now be obtained.

THEOREM 3.2. *Let r be the maximum of 2 and $c_2 \ln(1/na_n)/\ln(f'(1))$. There are constants c_2 and c_3 such that for any $n > 1$,*

$$\lambda_2 \leq c_3 \sqrt{r \ln(r) \ln(n)/n}\,.$$

The class of periodic trees may be defined in any of a number of ways so that a corresponding theorem is true. Here is one possible course broad

enough to include the example at the beginning of this section. Let G be any finite graph with a root, ρ, a second distinguished vertex, σ, and another set of distinguished vertices $v_1, \ldots, v_m$ distinct from each other. Suppose that for some j_1 and j_2 a path from σ to each v_i of length at most j_2 exists that is disjoint from a fixed path of length j_1 connecting σ and ρ. Let n be 1 less than the number of neighbours of σ and let $j = j_1 + j_2$. Construct a graph that will be an m-ary tree of G's as follows.

Let $\mathbf{T}^{(1)}(G) = G$. For each $i \geq 1$ let $\mathbf{T}^{(i+1)}(G)$ be a copy of G together with m copies of $\mathbf{T}^{(i)}(G)$ but with ρ in the ith copy of $\mathbf{T}^{(i)}$ identified with v_i in the copy of G, so there remains only a single root in $\mathbf{T}^{(i+1)}$, namely the one from the copy of G. The increasing limit of these trees (as seen from the root) is the desired graph $\mathbf{T}(G)$. If G is a tree, then so is $\mathbf{T}(G)$.

THEOREM 3.3. *For G, n, j, m and $\mathbf{T}(G)$ as above, let r be the maximum of 2 and $\lceil j/\ln(m) \rceil$. Then there is a constant c_4 such that $\lambda_2 \leq c_4 \sqrt{r \ln(r) \ln(n)/n}$.*

The proofs of Theorems 3.2 and 3.3 will be given in Section 5 since they rely on the results from Section 4 on finite stars.

4. Finite stars. Let $\mathbf{T}$ be a star of size n, that is, $\mathbf{T}$ consists of a root ρ and n other vertices each connected only to ρ. Although any finite system is eventually trapped in the state of zero infection, this star turns out to be capable of storing the infection for a long time when λ is an appropriate constant times $1/\sqrt{n}$. Since the application of these results is to get the two asymptotics is (7) and (8) to agree within a constant factor, there is no loss of generality in assuming n is greater than 64, say, so that $n^{-1/2}$ and $n^{-1/4}$ are actually distinguishable. It will also be convenient to restrict λ to being less than 1; again, nothing is lost, since all the critical values are less than 1 for $n \geq 64$.

Before stating the theorems precisely, here is a heuristic explanation of the results. Suppose n is large and let $\lambda = a/\sqrt{n}$ for some constant a. Let x be the number of infected neighbors of ρ divided by $\sqrt{n}$ and pretend that x is a continuous rather than a discrete parameter. When the root is not infected, x decreases deterministically at a rate x (i.e., $x = e^{-t}x_0$) and when the root is infected, x approaches a at rate $a - x$ [i.e., $x - a = e^{-t}(x_0 - a)$]. The root, when infected, recovers at rate 1, and when healthy, is infected at rate $(a/\sqrt{n})(x\sqrt{n}) = ax$. In other words, the state of infection is a Markov process on $\{0, 1\} \times [0, a]$, where the second coordinate is x and the first coordinate is 0 if the root is healthy and 1 if the root is infected. Its natural scale can be calculated exactly, giving a function $W: \{0, 1\} \times [0, a] \to \mathbf{R}$ for which $W(t)$ is a martingale. Writing $f(x) = W(1, x)$ and $g(x) = W(0, x)$ yields

$$(a - x)f' + (g - f) = 0,$$

$$-xg' + ax(f - g) = 0,$$

which have the unique (up to affine transformation) solution

$$(9) \quad f(x) = \int_0^x e^{-as}(a - s)^{-2}\,ds, \qquad g(x) = f(x) - e^{-ax}(a - x)^{-1}.$$

From this it is easy to read off the behavior of $x(t)$ and, in particular, to see that for $k < l < m$, the chance of hitting $x = a/k$ before hitting $x = a/m$, starting from $x = a/l$, is of the order $e^{-a^2(1/l - 1/m)}$.

Returning to the actual situation, where the number of infections is integral, there are roundoff errors in the statement of an analogous result due to the greatest integer function, and roundoff errors in the proof due to the fact that the functions f and g are no longer very nice. To kill the roundoff error, I require n to be larger than some N and choose sequentially a bunch of constants satisfying certain relations. The constants involved are named c_5 through c_{13} and are fixed throughout once they are chosen. These can be evaluated easily by tracing the argument, but the value of N that this yields is far from optimal.

THEOREM 4.1. *Let* c_9, c_{10}, c_{11} *be constants satisfying* $1/4 > 1/c_{10} + 1/c_{11} > 1/c_{10} > 1/c_9 > 0$ *and let* c_5 *satisfy* $1/c_5 < 1/c_{10} - 1/c_9$. *For* n *a positive integer and* a *a positive real number, run the contact process on a star of size* n *with* $\lambda = a/\sqrt{n}$ *and initially* y *nodes infected, where* $y = \lfloor a\sqrt{n}/c_{10} \rfloor$. *Then the conditions*

$$n > (1/c_{10} - 1/c_9 - 1/c_5)^{-1} \quad and \quad 4 \le a \le \sqrt{n}$$

imply that the probability of the number of infected nodes dropping below $a\sqrt{n}/c_9$ *before reaching* $y + a\sqrt{n}/c_{11}$ *is at most* e^{-a^2/c_5}.

COROLLARY 4.2. *If* n *is large enough, then beginning with only* ρ *infected at time* 0, *the probability of the number of infected nodes staying above* $a\sqrt{n}/c_9$ *during the time interval* $1 \le t \le e^{a^2/c_5}/2c_{11}$ *is at least* $e^{-1}/5$.

REMARKS. (1) The important fact here is that the infection lives for a time exponential in a^2 with a probability bounded away from 0; this is what creates the order $\sqrt{\ln(n)/n}$ bound in Theorems 3.2 and 3.3, via (11) in the next section. (2) As observed earlier, requiring n to be large and $\lambda = a/\sqrt{n}$ to be less than 1 gives away nothing.

PROOF OF COROLLARY 4.2. The probability that ρ remains infected until at least time 1 is e^{-1}. Remember this value: $p_1 = e^{-1}$. Conditioned on that, the probability of finding at least y nodes infected at time 1 is at least the probability of a binomial with parameters n and $Q(1)$ being at least y, where $Q(1) = (a/(a + \sqrt{n}))(1 - e^{-(1 + a/\sqrt{n})}) > (1 - e^{-1})a/2\sqrt{n}$ is as in (5). The mean of this binomial is at least $a/\sqrt{n}/4$ and the variance is at most $nQ(1) \le an/(a + \sqrt{n})$ so if n is large enough, the binomial is at least $y = a\sqrt{n}/c_{10}$ with probability arbitrarily close to 1, say $4/5$. Remember this too: $p_2 = 4/5$.

Next, conditional on all that, Theorem 4.1 implies that the number of times the number of neighbors of ρ goes from y to $y + a\sqrt{n}/c_{11}$ before becoming less than $a\sqrt{n}/c_9$ dominates a geometric random variable with mean at least e^{a^2/c_5}. Finally, if n is large enough to apply Theorem 4.1, then the waiting times for each new infection are independent from everything else and are at least exponential with mean $1/a\sqrt{n}$ (the infection rate summed over all neighbors of the root is at most $a\sqrt{n}$). The mean sum of $a\sqrt{n}/c_{11}$ of these waiting times e^{a^2/c_5} times over is thus at least $e^{a^2/c_5}/c_{11}$. Any sum of a geometric random number of i.i.d. exponential random variables is at least half its mean with probability at least $1/4$ if the geometric has large enough mean. So the probability that the sum of the waiting times is at least $e^{a^2/c_5}/2c_{11}$ conditional on there being at least e^{a^2/c_5} transitions from y to $y + a\sqrt{n}/c_{11}$ is at least $1/4$. Call this p_3. Now multiply p_1, p_2 and p_3 to give the result. □

PROOF OF THEOREM 4.1. For each state ξ, let $I(\xi)$ be 1 if the root is infected and 0 otherwise. Assign each state ξ the weight

$$W(\xi) = e^{-ax/10\sqrt{n}}\left(1 - I(\xi)\frac{\sqrt{n}(e^{a/10\sqrt{n}} - 1)}{a}\right),$$

where x is the number of infected neighbors of ρ. This W is not a martingale like the heuristic one was, but reducing the exponent from $ax/\sqrt{n}$ to $ax/10\sqrt{n}$ makes it a supermartingale by, as will be evident, a comfortable margin. Note that by the size assumption on a, $a/10\sqrt{n} \leq 0.1$ so $e^{a/10\sqrt{n}} - 1 \leq 1.2a/10\sqrt{n}$, hence $\sqrt{n}(e^{a/10\sqrt{n}} - 1)/a \leq 0.12$ and W is always positive. Let τ be the first time that ρ has at least $y + a\sqrt{n}/c_{11} < a\sqrt{n}/4$ infected neighbors. The state $\xi(t)$ is a continuous-time Markov chain on the space $\{0, 1\} \times \{0, 1, \ldots, n\}$, where ξ is represented as the pair $(I(\xi), \text{number of infections})$. The transition rates are

$$(0, x) \to (0, x - 1) \quad \text{at rate } x,$$

$$(0, x) \to (1, x) \quad \text{at rate } \frac{ax}{\sqrt{n}},$$

$$(1, x) \to (0, x) \quad \text{at rate } 1,$$

$$(1, x) \to (1, x - 1) \quad \text{at rate } x,$$

$$(1, x) \to (1, x + 1) \quad \text{at rate } \frac{a}{\sqrt{n}}(n - x).$$

For $W(\xi(t \wedge \tau))$ to be a supermartingale, two inequalities must be satisfied. The transitions from $(0, x)$ yield the inequality

$$\frac{ax}{\sqrt{n}}\frac{\sqrt{n}}{a}(e^{a/10\sqrt{n}} - 1) \geq x(e^{a/10\sqrt{n}} - 1),$$

which is satisfied with equality. Setting $k = (\sqrt{n}/a)(e^{a/10\sqrt{n}} - 1)$, the transitions from $(1, x)$ yield

$$k + x(1 - k)(e^{a/10\sqrt{n}} - 1) \leq \frac{a}{\sqrt{n}}(n - x)(1 - k)e^{-a/10\sqrt{n}}(e^{a/10\sqrt{n}} - 1)$$

and dividing out by $1 - k$ and by $e^{a/10\sqrt{n}} - 1$ gives

$$\frac{k}{(1 - k)(e^{a/10\sqrt{n}} - 1)} + x \leq \frac{a}{\sqrt{n}}(n - x)e^{-a/10\sqrt{n}},$$

$$x + \frac{\sqrt{n}}{a - \sqrt{n}(e^{a/10\sqrt{n}} - 1)} \leq \frac{a}{\sqrt{n}}(n - x)e^{-a/10\sqrt{n}},$$

$$e^{a/10\sqrt{n}}\sqrt{n}\,x + n\frac{e^{a/10\sqrt{n}}}{a - \sqrt{n}(e^{a/10\sqrt{n}} - 1)} \leq an - ax,$$

$$(a + \sqrt{n}\,e^{a/10\sqrt{n}})x \leq n\left(a - \frac{e^{a/10\sqrt{n}}}{a - \sqrt{n}(e^{a/10\sqrt{n}} - 1)}\right).$$

Since $a \leq \sqrt{n}$, the left-hand side is at most $(1 + e^{1/10})\sqrt{n}\,x$, and since $a \geq 4$, the right-hand side is at least $an(1 - e^{1/10}/(4 - \sqrt{n})(e^{2/5\sqrt{n}-1})) > an/2$ (the denominator on the right-hand side has a local minimum at either endpoint of $[4, \sqrt{n}]$ with the absolute minimum at 4). Thus the inequalities are satisfied as long as $1.11\sqrt{n}\,x < an/2$, which is always satisfied when $x < a/4\sqrt{n}$, which is always satisfied up to the stopping time. This establishes that $W(\xi(t \wedge \tau))$ is a supermartingale.

To finish proving the theorem, suppose that ρ has y infected neighbors at time t, where y is as in the statement of the theorem. There is no loss of generality in the worst-case assumption that ρ is not infected at time t. Then $W(\xi(t)) = e^{-ay/10\sqrt{n}}$ and by definition of y, this is between $e^{-a^2/10c_{10}}$ and $e^{-a^2/10c_{10}+a/10\sqrt{n}}$. But $W(\xi) \geq e^{-a^2/10c_9}$ the first time the number of infected nodes falls below $a\sqrt{n}/c_9$. Use the fact that W is always positive together with the optional stopping theorem to see that W can reach the value $e^{-a^2/10c_9}$ with probability at most $\exp(-a^2/10c_{10} + a/10\sqrt{n} + a^2/10c_9)$. This gives the desired conclusion as long as $a^2/10c_{10} - a/10\sqrt{n} - a^2/10c_9 > a^2/c_5$. Factoring out $a^2/10$ yields the equivalent condition $1/c_{10} - 1/c_9 - 1/c_5 > 1/a\sqrt{n}$ and since a is at least 1, the condition $n > (1/c_{10} - 1/c_9 - 1/c_5)^{-1}$ finishes the proof. □

One more lemma will be needed for the next section, so we record it now.

LEMMA 4.3. *Let n, a and λ be as in Theorem 4.1 and run the contact process on any graph containing the finite tree* **T**, *where* **T** *consists of a star of size n with the root infected at time 0, to which has been added a single chain*

of length r of descendants of some child v_1 of ρ, denoted $v_2, \ldots, v_r$. Then

(10)
$$\mathbf{P}\big(v_r \text{ is infected before time } e^{a^2/c_5}/2c_{11}\big)$$
$$> e^{-1}/5 - (1 - c^r\lambda^r)^{(e^{a2/c_5})/2c_{11}r - 2},$$

where c is a constant independent of everything else.

PROOF. By monotonicity, it is sufficient to prove the statement for a modified contact process in which v_2 cannot infect v_1, so the process restricted to the star of size n is unaffected by what happens at v_1 and below. Also by monotonicity assume without loss of generality that the graph is actually equal to **T**. Let $\mathscr{B}_t$ be the event that the number of infected nodes of the star at time t is at least $a\sqrt{n}/c_9$, $\mathscr{B} = \cap\{\mathscr{B}_t: 1 \le t \le e^{a^2/c_5}/1c_{11}\}$, and let $\mathscr{F}_t$ be the σ-algebra of events up to time t. Then $\mathbf{P}(v_r$ is infected at some time $s \in (t, t + r)|\mathscr{F}_t, \mathscr{B}_t)$ is bounded below by the probability that ρ becomes infected before time $t + 1$ given $\mathscr{F}_t, \mathscr{B}_t$ times the product of the probabilities that v_i infects $v_i + 1$ before becoming uninfected and this happens within time 1. Working out these probabilities gives a lower bound of $(c\lambda)^r$ for a universal constant c.

Now let $L = \lfloor(1 + e^{a^2/c_5})/2c_{11}r\rfloor - 1$ and consider a sequence of times $\{t_i\} = \{1, 1 + r, 1 + 2r, \ldots, 1 + Lr\}$. For each such time t_i, $\mathbf{P}(v_r \notin \xi(t)$ for $t_i < t < t_i + r|\mathscr{F}_{t_i}, \mathscr{B}_{t_i}) \le 1 - c^r\lambda^r$, and so $\mathbf{P}(v_r \notin \xi(t)$ for $1 \le t \le 1 + Lr|\mathscr{B}) \le (1 - c^r\lambda^r)^L$, since each $\mathscr{B}_{t_i} \in \mathscr{F}_{t_i}$ and $\mathscr{B} \in \cap \mathscr{B}_{t_i}$. Corollary 4.2 gives $\mathbf{P}(\mathscr{B}) > e^{-1}/5$ and since L is at least as big as the exponent in the right-hand side of (10), the lemma is proved. □

5. Proofs of Theorems 3.2 and 3.3 and examples.

PROOF OF THEOREM 3.2. Assume without loss of generality that $a_n > 0$ (or else r is infinite). Also, there is no loss of generality in assuming n to be large enough to apply Corollary 4.2, since c_2 and c_3 may be taken large enough so that $c_3\sqrt{r \ln(r) \ln(n)}/n$ is at least 2 for any n that is too small, and 2 is an upper bound on all critical values on trees. By Proposition 3.1, is suffices to prove that $\liminf_{t \to \infty} \mathbf{P}(\rho \in \xi(t)) > 0$ without conditioning on **T**. Let

$$g(t) = \mathbf{P}(\rho \in \xi(t)) \ge a_n\mathbf{P}(\rho \in \xi(t)|\rho \text{ has at least } n \text{ children}).$$

To calculate $g(t)$, condition on the root having at least n children until further notice. Let c be the constant in Lemma 4.3. Treat r as a parameter to be determined later and choose $a = a(r)$ so that

(11)
$$e^{a^2/c_5}/2c_{11}r > 2 + 10\big(c^r(a/\sqrt{n})^r\big)^{-1}.$$

Using the inequality $(1 - x)^{1/x} < e^{-1}$, observe that (11) forces the right-hand side of (10) to be at least $c_8 = e^{-1}/4 - e^{-10}$ when λ is at least $a/\sqrt{n}$. To satisfy (11), it suffices to choose $a > \sqrt{c_5 \ln\big(2c_{11}r \cdot 10 \cdot c_1^{-r}\sqrt{n}^r\big)}$, which is at most $c_3\sqrt{r \ln r \ln n}$ for an appropriate c_3 as long as $r \ge 2$. Set $\lambda = a/\sqrt{n}$. If

this is sufficient to force $\mathbf{P}(q(\mathbf{T}) > 0 | \, |\mathbf{T}| = \infty) = 1$, then $\lambda_2(f) \leq c_3\sqrt{r \ln(r)\ln(n)/n}$ as desired. Now the root has at least n children; select n of them and ignore the rest (by monotonicity). Among all the vertices at distance r from ρ, look for those having at least n children; call these $v_r^1, \ldots, v_r^j$, where j is the random cardinality of the set of such descendants. From the facts that ρ has n children, that each descendant has an expected $f'(1)$ children and that a node has n children with probability a_n, it follows that

$$\mathbf{E}j \geq na_n(f'(1))^{r-1}.$$

Now let M be the random number of vertices among $v_r^1, \ldots, v_r^j$ that are infected before $e^{a^2/c_5}/2c_{11}$. By the choice of a, Lemma 4.3 and the above expression for $\mathbf{E}j$,

$$\mathbf{E}M \geq c_8na_n(f'(1))^{r-1}.$$

Now the argument proceeds as in (iv) of the previous section. Fix t and let x be the minimum value of $g(s)$ for $0 \leq s \leq t - e^{a^2/c_5}/2c_{11}$. Recursively calculate a lower bound on $g(t)$ as follows.

Ignoring all but the first infection of each v_r^i by its parent, any of these that are infected at some times $s < t - e^{a^2/c_5}/2c_{11}$ will evolve independently from time s to time $t - e^{a^2/c_5}/2c_{11}$ and then be infected with probability at least $x = \inf\{g(s): s \leq t - e^{a^2/c_5}/2c_{11}\}$. Thus the distribution of the number of such nodes infected at time $t - e^{a^2/c_5}/2c_{11}$ given M dominates a random variable M_x that is binomial(M, x). Thus for small enough x, $\mathbf{P}(M_x \geq 1) \geq x(\mathbf{E}M - \varepsilon)$ for arbitrarily small ε. Finally, if some v_r^i is infected at time $t - e^{a^2/c_5}/2c_{11}$, then the probability of finding ρ infected at time t is bounded below by the probability p_1 that the contact process starting with only v_r^i infected at time $t - e^{a^2/c_5}/2c_{11}$ infects ρ at some time s with $t - e^{a^2/c_5}/2c_{11} \leq s \leq t$ times the probability p_2 that ρ is infected at time t given infection of ρ at some such time s. Then p_1 is at least c_8 by Lemma 4.3 and (11) with the star of size n centered at v_r^i this time. Meanwhile, p_2 is at least $(e^{-1}/5)e^{-1}a^2/(c_9 + a^2)$, which will be called c_{13}. [See this by applying Corollary 4.2 together with the fact that the probability of ρ becoming reinfected between times $t - 1$ and t given at least $a/\sqrt{n}/c_9$ infections during this time, and then staying reinfected, is at least $a^2/(c_9 + a^2)$.]

Thus as $x \to 0$, $g(t)$ is bounded below by $c_8c_{13}x\mathbf{E}(M) \geq c_8^2c_{13}xna_n(f'(1))^{r-1}$. Now choose c_2 so that putting $r = c_2 \ln(1/na_n)/\ln(f'(1))$ forces $c_8^2c_{13}2na_n(f'(1))^{r-1}$ to be greater than 1. Applying Lemma 2.4 then gives $\inf_t g(t) > 0$. $\square$

PROOF OF THEOREM 3.3. Passing to a subtree if necessary, assume without loss of generality that the parent of σ is on the unique path of length j_1 connecting ρ to σ and that there is a unique path connecting σ to each v_i and it passes through a child of σ, not through the parent of σ.

For a positive integer q to be determined later, consider $\mathbf{T}(G)$ as being put together from units of size $\mathbf{T}^{(q)}$ instead of from units of size G. Then ρ has at

least n neighbors and furthermore it has at least m^q descendants at distance at most qj, all of which continue recursively the property of having the right number of neighbors and well-endowed descendants.

Now use the argument for the previous theorem with m^q instead of $a_n(f'(1))^{r-1}$. Then a must be at least $c_3\sqrt{qj \ln(qj)}\ln(n)/n$ as before, and the condition needed for Lemma 2.4 is that $c_8^2 c_{13} m^q > 1$. This means that $q = c'/\ln(m)$ suffices, and setting $r = qj/c'$ proves the theorem. $\square$

A couple of examples conclude this section. The first shows that the upper bound for λ_2 in Theorem 3.3 can be asymptotically sharp. In fact, for this example it is shown that λ_1 is within a constant factor of the upper bound for λ_2 as $n \to \infty$, so Theorem 3.3 provides asymptotically tight bounds on both critical values.

EXAMPLE. This example calculates λ_1 for a particular tree in order to show that the order $\sqrt{\ln(n)/n}$ bound given by Theorem 3.3 may be sharp to within a constant factor. Let $\mathbf{T}$ be a homogeneous binary tree to every vertex of which has been added n children that have no children themselves. Thus each vertex has $n + 2$ children, only 2 of which have descendants. To get a lower bound on λ_1, define the weight of a configuration as the sum over all nodes that have children of 1 if the node or any of its childless children are infected and 0 otherwise. The following partial converse to Corollary 4.2 will be needed.

PROPOSITION 5.1. *Let $\mathbf{T}$ be a star of size n with $\lambda = a/\sqrt{n} < 1$ as usual. There is an N for which $n > N$ implies that the probability starting from configuration η of the infection dying out before time $\ln(n) + 2(\ln|\eta| \vee 4)$ is at least $e^{-4a^2}/10$.*

PROOF. First observe that from any configuration η, $\mathbf{P}_\eta(|\xi(2(\ln|\eta| \vee 4))| < 2a\sqrt{n}) > 3/4$. To see this, couple the process on $\mathbf{T}$ to one on another star of size n, $\mathbf{T}'$, where the root of $\mathbf{T}'$ always remains infected, but none of the children of ρ' in $\mathbf{T}'$ starts out infected. The coupling has corresponding children of ρ and ρ' recover at the same times and has an infection occurring in $\mathbf{T}'$ whenever one occurs in $\mathbf{T}$ but not necessarily vice versa. When every child of ρ has recovered once, the number of infected children of ρ is at most the number of infected children of ρ'. This happens by time $2(\ln|\eta| \vee 4)$ with probability $(1 - e^{-2(\ln|\eta| \vee 4)})^{|\eta|} > 7/8$. But the number of infected children of ρ' at any time is a random variable converging upwards to a stationary distribution that has mean approximately $a\sqrt{n}$ and for large enough n has probability less than $1/8$ of being above $2a\sqrt{n}$. Thus the coupling establishes the observation. Now from a configuration where at most $2a\sqrt{n}$ nodes are infected, it remains to show that the probability of the infection dying out within time $\ln(n)$ is at least $e^{-4a^2}/7.5$. The probability that no reinfection ever occurs is at least $(1/(1 + \lambda))^{4a\sqrt{n}} > (1 - a/\sqrt{n})^{4a\sqrt{n}} \approx e^{-4a^2}$. If no reinfection

occurs, the infection dies out within time $\ln(n)$ with probability at least $(1 - e^{-\ln(n)})^{2a\sqrt{n}} = (1 - 1/n)^{2a\sqrt{n}} \geq e^{-2}$ since $a \leq \sqrt{n}$ by assumption. Multiplying together the conditional probabilities gives $e - 4a^2(3/4)/e^2 \geq e^{-4a^2}/10$.

$\square$

What is needed below is actually an altered version of this proposition saying that even if ρ is given three additional neighbors that are always infected, the probability of the infection dying out except on the additional three neighbors in the allotted time is still at least $e^{-4a^2}/10$. The reader may fill in the details to beef up the proof.

Going back now to the example, consider any node v that contributes 1 to the weight of the configuration either by having an infected child or by being a node with children that is itself infected. There are only three neighbors of v that may increase the weight of the configuration if they get infected by v, so the expected contribution of v to increasing the weight of the configuration in a time period of length $3\ln(n)$ is at most

$$9\lambda \ln(n) = 9a \ln(n)/\sqrt{n}.$$

The probability of v decreasing the weight during this time because the infection dies out on v and its childless children is at least $e^{-4a^2}/10$, by the beefed-up version of the previous proposition (since $|\eta|$ is necessarily at most n). Now pick $a = \sqrt{\ln(n)}/10$. This makes

$$e^{-4a^2}/10 = n^{0.1}/10\sqrt{n}/2 > 9\sqrt{\ln(n)} \ln(n)/10\sqrt{n} = 9a \ln(n)/\sqrt{n}$$

for large enough n. Thus any node contributes an expected weight that is less than 1 before it and its childless children all recover. The usual supermartingale argument then implies that the infection dies out almost surely and so

$$(12) \qquad\qquad\qquad \lambda_1 \geq \sqrt{\ln(n)}/10\sqrt{n}$$

for large enough n.

On the other hand, $\mathbf{T}$ is a periodic tree: $\mathbf{T} = \mathbf{T}(G)$ for the graph G with root $\rho = \sigma$ and $n + 2$ other vertices all connected only to the root, two of which are labeled v_1 and v_2. In the notation of Section 3, $m = 2$, $j = 1$ and what is called n in Theorem 3.3 is $n + 2$. Then Theorem 3.3 gives

$$(13) \qquad\qquad\qquad \lambda_2 \leq c_4\sqrt{\ln n}/\sqrt{n}.$$

Thus for this class of trees, both (12) and (13) are sharp to within a constant factor.

EXAMPLE. This example is thrown in to show that even medium-sized tails on the generating function of a Galton–Watson tree can force the critical values down to 0. For any positive $\gamma < 1$, let

$$f(x) = \sum_n ce^{-n^\gamma}x^n$$

be the generating function for the number of children, where c is a normalizing constant. Then $\ln(1/na_n) \approx n^\gamma$, so Theorem 3.2 gives $\lambda_2 \leq c\sqrt{n^{\gamma-1}\ln(n)}$, which goes to 0 as $n \to \infty$. Thus λ_2 is less than any positive number, hence $\lambda_1 = \lambda_2 = 0$.

6. Continuity of the survival probability. Continuity from the right of the survival probability as a function of λ is immediate from the monotonicity, since versions of the process for each λ can all be coupled so that the event of survival is continuous from the right in λ. At the moment, I can only prove continuity from the left when the phase transitions at λ_1 and at λ_a are distinct. It is also necessary to assume that the tree is homogeneous or periodic. In particular, the survival probability is a continuous function of λ for homogeneous trees with $n > 2$. I believe that survival probability is continuous for arbitrary trees.

I will present the argument that $\lambda_1 < \lambda_a$ implies continuity of the survival probability only for homogeneous trees; the generalization to periodic trees is straightforward but tedious. In the foregoing discussions, let $p(\lambda)$ denote the survival probability for the contact process on an n-ary homogeneous tree with only the root infected initially. Begin with a counting lemma.

DEFINITION 6.1. For a finite set $S \subseteq \mathbf{T}$, say a vertex $v \in S$ is surrounded in S iff each of the $n + 1$ components of $\mathbf{T} \setminus \{v\}$ intersects S.

LEMMA 6.2. *Let S be a finite set of vertices of a homogeneous tree, $\mathbf{T}$. Then the number of surrounded vertices of S is less than $|S|/n$.*

PROOF. The proof proceeds by repeated reductio ad absurdum. Suppose for a contradiction that the result is false. Let S be a counterexample with as few vertices as possible, and among such counterexamples, let S have as small a value as possible for the sum $\sum_{v,w \in S} d(v, w)$. The *induced subgraph* on S is just the vertices of S together with whichever edges of $\mathbf{T}$ connect two vertices in S.

I claim that the induced subgraph on S must be connected. Suppose not. Let $v \in \mathbf{T} \setminus s$ be a vertex for which $\mathbf{T} \setminus v$ has at least two components intersecting S (the existence of such a v is equivalent to S failing to be connected). There are two cases. First, suppose more than two components of $\mathbf{T} \setminus v$ intersect S. Writing $S = \cup S_i$, where each S_i is the intersection of S with a single component of $\mathbf{T} \setminus v$, deletion of any single one of the S_i from S does not affect which of the remaining vertices are surrounded. Then by deleting the S_i with the smallest fraction of surrounded nodes, a smaller counterexample is obtained, which is a contradiction. The other case is that exactly two components of $\mathbf{T} \setminus v$ intersect S. Let S_1 and S_2 be the intersection of S with these two components and let v_1 and v_2 be the neighbours of v in the two components. Let ϕ be any graph-automorphism of $\mathbf{T}$ sending v_2 to v and v to v_1. Let $S' = S_1 \cup \phi[S_2]$. Then vertices in S_1 are surrounded in S

iff they are surrounded in S', and vertices in S_2 are surrounded in S iff their images under ϕ are surrounded in S'. So S' is a counterexample with as many vertices as S but a smaller value $\Sigma_{v,\,w\,\in\,S}\,d(v, w)$, which is a contradiction.

Now that it is established that the induced subgraph on S is connected, count the edges of S in two different ways. First, since the induced subgraph is a tree, there is one fewer edge than there are vertices. Counting oriented edges doubles this, so there are $2|S| - 2$ oriented edges. On the other hand, there are $n + 1$ oriented edges leading out of every surrounded vertex and at least one edge leading out from every vertex, so the number of oriented edges is at least $|S| + n$(number of surrounded vertices), thus $2|S| - 2 \geq |S| + n$(number of surrounded vertices), and there are less than $|S|/n$ surrounded vertices. This is not a counterexample, thus no counterexample exists. □

The rest of the discussion deals with the contact process on a tree with a severed edge. For each positive integer r, let $\xi^{(r)}$ be a contact process run from an initial infection at the root with an edge severed at distance r from the root, that is, no infection passes over that edge. Write $\xi^{(r;\lambda)}$ or $\xi^{(;\lambda)}$ when it is necessary to emphasize the choice of λ. An adjacent edge counts as being at distance 1. The main lemma on continuity from the left is the following.

LEMMA 6.3. *Let* **T** *be a homogeneous tree. For a fixed λ, let $\xi^{(r)}$ be as above, and suppose the survival probability for the infection in $\xi^{(r)}$ is positive for some r. Then the survival probability for the unaltered contact process is continuous from the left for that value of λ.*

REMARK. There is a corresponding lemma in percolation theory stating that the percolation probability is continuous provided that percolation in $\mathbf{Z}^d$ implies percolation in a half-space.

PROOF. First note that r can be taken to be 1, since the initial infection in the $\xi^{(1)}$ process has a nonzero chance of moving to a distance r from the severed edge before it dies, thus dominating a $\xi^{(r)}$ process a positive fraction of the time.

The next claim is that if $\xi^{(1)}$ survives with positive probability at a particular value of λ, then it survives with positive probability for some smaller value λ'. To find such a λ', begin by setting $p = p(\lambda)$ equal to the survival probability for $\xi^{(1;\lambda)}$ and pick t large enough so that

$$\mathbf{P}\big(|\xi^{(1;\lambda)}(t)| \geq 10(1 + \lambda)/\lambda p\big) \geq 3p/4.$$

Such a t exists, since $|\xi^{(1)}(t)| \to \infty$ almost surely when the infection lives forever. Pick $\lambda' \in [\lambda/2, \lambda)$ large enough so that

$$\mathbf{P}\big(|\xi^{(1;\lambda')}(t)| \geq 5(1 + \lambda')/\lambda' p\big) \geq p/2.$$

This is possible because $p/2 < 3p/4$, $5(1 + \lambda')/\lambda' p < 10(1 + \lambda)/\lambda p$ and the distribution of $\xi^{(1)}(t)$ is continuous in λ for fixed t.

For this paragraph, let the infection rate be understood to be λ'. Let $S \subseteq \xi^{(1;\lambda')}(t)$ be the random subset of the infected nodes at time t that are not surrounded in the set of infected nodes at time t. Since $1 - 1/n \geq 1/2$, Lemma 6.2 shows that S contains at least half the nodes infected at time t. Thus

(14) $$\mathbf{P}(|S| \geq 5(1 + \lambda')/2\lambda'p) \geq p/2.$$

For each $v \in S$ there is by definition of S some neighbor, say $g(v)$, for which the component of $\mathbf{T} \setminus v$ containing $g(v)$ does not intersect $\xi^{(1;\lambda')}(t)$. For each such v, the probability is $\lambda'/(1 + \lambda')$ that v will pass the infection on to $g(v)$ before recovering. Let A be the random set of such v. Forbid any infection of nodes in S from time t onward. Then the process proceeds independently on each component of $\mathbf{T} \setminus S$ containing some $x \in A$, since communication has been severed. Furthermore, the starting configuration on each such component is precisely a single infection and the process evolves as $\xi^{(1;\lambda')}$, since an edge next to the initial edge has been severed. Then the $\xi^{(1)}$ process dominates a set of independent $\xi^{(1)}$ processes, started from each vertex in A. As in the theory of Galton–Watson processes, the survival probability is greater than the fixed point of a generating function for the probabilities of the various cardinalities of A. This must be positive if $\mathbf{E}|A| > 1$. Since

$$E(|A|\,|\,|S|) > 5(1 + \lambda')/2\lambda p) > 5/2p,$$

the desired result follows from (14).

Finally, to show continuity from the left, recall $p(\lambda')$ is the survival probability for a severed edge and parameter λ', and let q be the survival probability for the unaltered contact process with parameter λ. Let $\varepsilon > 0$ be arbitrary and let M be large enough so that $(1 - p(\lambda'))^M < \varepsilon/4$. Also pick M large enough so that any binomial random variable with mean at least M is at least $M/2$ with probability at least $1 - \varepsilon/4$. Let t be large enough so that

$$\mathbf{P}(|\xi^{(;\lambda)}(t)| \geq 4(1 + \lambda')M/\lambda') \geq q - \varepsilon/4,$$

for the parameter value λ, and pick λ'' with $\lambda' \leq \lambda'' < \lambda$ large enough so that

$$\mathbf{P}(|\xi^{(\lambda'')}(t)| \geq 2(1 + \lambda'')M/\lambda'') \geq q - \varepsilon/2.$$

Rerun the previous argument but for infection rate λ''. Let S be the unsurrounded infected nodes in $\xi^{(;\lambda'')}(t)$, so each $v \in S$ has a neighbor $g(v)$ in a component of $\mathbf{T} \setminus S$ with no infection. There is a set A of nodes $g(v)$ for $v \in S$ that are infected before v recovers. Formally, A is a set of nodes $g(v)$ such that $v \in S$ and there exists a $t' > t$ for which v is infected throughout the time interval $[t, t']$ and such that v infects $g(v)$ at time t'. The binomial hypothesis on M guarantees that A has cardinality at least M with probability at least $q - 3\varepsilon/4$, since each $g(v)$ is independently in A with probability $\lambda''/(1 + \lambda'')$. Each node in A defines an independent process which is a restarted version of $\xi^{(1)}$ with parameter $\lambda'' \geq \lambda'$, so the hypothesis on $p(\lambda')$ and M gives a survival probability of at least $q - \varepsilon$ for the whole process. Since ε was arbitrary, the lemma is proved. $\square$

It remains to show that the hypothesis of this lemma is satisfied when the phase transitions at λ_1 and λ_a are distinct. This is done in two steps.

LEMMA 6.4.   *Suppose the contact process is run on a homogeneous tree for some $\lambda < \lambda_a$. Then $\mathbf{P}(v$ is ever infected$) \to 0$ as $d(\rho, v) \to \infty$.*

PROOF.   Suppose not. Then by monotonicity it decreases to some bound $b > 0$. Since $\lambda < \lambda_a$ the root is infected finitely often almost surely, so there is some time t and configuration η such that $\mathbf{P}($survival without infecting the root any more $| \xi(t) = \eta) > 1 - b$. Since any node alive at time t infects the root with probability at least b, this is a contradiction. $\square$

LEMMA 6.5.   *Suppose $\lambda < \lambda_a$. Then the survival probability for $\xi^{(r)}$ on a homogeneous tree approaches the survival probability for the unaltered contact process as $r \to \infty$.*

PROOF.   Couple all the $\xi^{(r)}$ with the unaltered process, so that they evolve identically until an infection in the unaltered process crosses an edge that is severed in the other process. By Lemma 6.4 the probability of the processes ever becoming uncoupled goes to 0 as r goes to ∞, and the result follows. $\square$

Putting all these lemmas together gives the following theorem.

THEOREM 6.6.   *If $\lambda_1 < \lambda_a$, then the survival probability for the contact process from a single initial infection on a homogeneous tree is continuous in λ.*

PROOF.   When the survival probability is 0, there is nothing to prove. When it is positive and $\lambda < \lambda_a$, then some $\xi^{(r)}$ also has positive survival probability by Lemma 6.5, and continuity follows from Lemma 6.3. For $\lambda \geq \lambda_a$, the existence of a λ between λ_1 and λ_a implies the survival with positive probability of a severed-edge contact process for that and hence all higher λ, and Lemma 6.3 again applies. $\square$

REFERENCES

[1] BEZUIDENHOUT, C. and GRIMMETT, G. (1990). The critical contact process dies out. *Ann. Probab.* **18** 1462–1482.
[2] DURRETT, R. (1988). *Lecture Notes on Particle Systems.* Wadsworth, New York.
[3] GRIMMETT, G. and NEWMAN, C. (1990). Percolation in $\infty + 1$ dimensions. In *Disorder in Physical Systems* (G. Grimmett and D. Welsh, eds.). Clarendon, Oxford.
[4] LIGGETT, T. (1985). *Interacting Particle Systems.* Springer, New York.
[5] LYONS, R. (1989). The Ising model and percolation on trees and tree-like graphs. *Comm. Math. Phys.* **125** 337–353.
[6] NEWMAN, C. and WU, C. (1991). Markov fields on branching planes. *Probab. Theory and Related Fields* **85** 539–552.

DEPARTMENT OF MATHEMATICS
UNIVERSITY OF WISCONSIN–MADISON
VAN VLECK HALL
480 LINCOLN DRIVE
MADISON, WISCONSIN 53706