1

Basic concepts and results

Percolation theory was founded by Broadbent and Hammersley [1957], in order to model the flow of fluid in a porous medium with randomly blocked channels. Interpreted narrowly, it is the study of the component structure of random subgraphs of graphs. Usually, the underlying graph is a lattice or a lattice-like graph, which may or may not be oriented, and to obtain our random subgraph we select vertices or edges independently with the same probability p. In the quintessential examples, the underlying graph is $\mathbb{Z}^d$.

The aim of this chapter is to introduce the basic concepts of percolation theory, and some easy fundamental results concerning them.

We shall use the definitions and notation of graph theory in a standard way, as in Bollobás [1998], for example. In particular, if Λ is a graph, then $V(\Lambda)$ and $E(\Lambda)$ denote the sets of vertices and edges of Λ, respectively. We write $x \in \Lambda$ for $x \in V(\Lambda)$. We also use standard notation for the limiting behaviour of functions: for $f = f(n)$ and $g = g(n)$, we write $f = o(g)$ if $f/g \to 0$ as $n \to \infty$, $f = O(g)$ if f/g is bounded, $f = \Omega(g)$ for $g = O(f)$, and $f = \Theta(g)$ if $f = O(g)$ and $g = O(f)$.

The standard terminology of percolation theory differs from that of graph theory: vertices and edges are called *sites* and *bonds*, and components are called *clusters*. When our random subgraph is obtained by selecting vertices, we speak of *site percolation*; when we select edges, *bond percolation*. In either case, the sites or bonds selected are called *open* and those not selected are called *closed*; the *state* of a site or bond is open if it is selected, and closed otherwise. (In some of the early papers, the term 'atom' is used instead of 'site', and 'dammed' and 'undammed' for 'closed' and 'open'.) In site percolation, the *open subgraph* is the subgraph induced by the open sites; in bond percolation, the *open subgraph* is formed by the open edges and all vertices; see Figure 1.

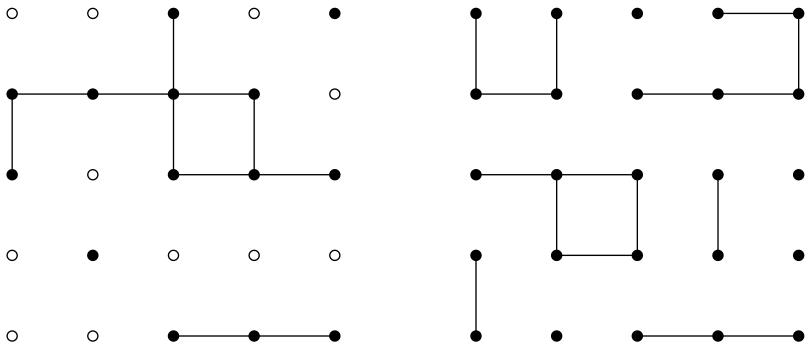

Figure 1. Parts of the open subgraphs in site percolation (left) and bond percolation (right) on the square lattice $\mathbb{Z}^2$. On the left, the filled circles are the open sites; the open subgraph is the subgraph of $\mathbb{Z}^2$ induced by these. For bond percolation, the open subgraph is the spanning subgraph containing all the open bonds.

To streamline the discussion, we shall concentrate on unoriented percolation, i.e., on (bond and site) percolation on an unoriented graph Λ. We assume that Λ is connected, infinite, and locally finite (i.e., every vertex has finite degree). In general, Λ is a *multi-graph*, so multiple edges between the same pair of vertices are allowed, but not loops. Most of the interesting examples will be simple graphs.

Often, we shall choose bonds or sites to be open with the same probability p, independently of each other. This gives us a probability measure on the set of subgraphs of Λ; in bond percolation we write $\mathbb{P}^{\mathrm{b}}_{\Lambda,p}$ for this measure, and in site percolation $\mathbb{P}^{\mathrm{s}}_{\Lambda,p}$. More often than not, we shall suppress the dependence of these measures on some or all parameters, and write simply $\mathbb{P}$ or $\mathbb{P}_p$. Similarly, Λ^{b}_p is the open subgraph in bond percolation, and Λ^{s}_p in site percolation.

Formally, given a graph Λ with edge-set E, a (bond) *configuration* is a function $\omega : E \to \{0,1\}$, $e \mapsto \omega_e$; we write $\Omega = \{0,1\}^E$ for the set of all (bond) configurations. A bond e is open in the configuration ω if and only if $\omega_e = 1$, so configurations correspond to open subgraphs. Let Σ be the σ-field on Ω generated by the cylindrical sets

$$C(F,\sigma) = \{\omega \in \Omega : \omega_f = \sigma_f \text{ for } f \in F\},$$

where F is a finite subset of E and $\sigma \in \{0,1\}^F$. Let $\mathbf{p} = (p_e)_{e \in E}$, with $0 \le p_e \le 1$ for every bond e. We denote by $\mathbb{P}^{\mathrm{b}}_{\Lambda,\mathbf{p}}$ the probability measure on (Ω, Σ) induced by

$$\mathbb{P}^{\mathrm{b}}_{\Lambda,\mathbf{p}}\big(C(F,\sigma)\big) = \prod_{\substack{f\in F\\ \sigma_f=1}} p_f \prod_{\substack{f\in F\\ \sigma_f=0}} (1-p_f). \tag{1}$$

When $p_e = p$ for every edge e, as before, we write $\mathbb{P}^{\mathrm{b}}_{\Lambda,p}$ for $\mathbb{P}^{\mathrm{b}}_{\Lambda,\mathbf{p}}$.

In the measure $\mathbb{P}^{\mathrm{b}}_{\Lambda,\mathbf{p}}$, the states of the bonds are independent, with the probability that e is open equal to p_e; thus, for two disjoint sets F_0 and F_1 of bonds,

$$\mathbb{P}^{\mathrm{b}}_{\Lambda,\mathbf{p}}\big(\text{the bonds in } F_1 \text{ are open and those in } F_0 \text{ are closed}\big) = \prod_{f\in F_1} p_f \prod_{f\in F_0} (1-p_f).$$

We call $\mathbb{P}^{\mathrm{b}}_{\Lambda,\mathbf{p}}$ an *independent bond percolation measure* on Λ. The special case where $p_e = p$ for every bond e is exactly the measure $\mathbb{P}^{\mathrm{b}}_{\Lambda,p}$ defined informally above. The formal definitions for *independent site percolation* are similar.

Let us remark that site percolation is more general, in the sense that bond percolation on a graph Λ is equivalent to site percolation on $L(\Lambda)$, the *line graph* of Λ. This is the graph whose vertices are the edges of Λ; two vertices of $L(\Lambda)$ are adjacent if the corresponding edges of Λ share a vertex; see Figure 2.

Although in this chapter we shall make some remarks about general infinite graphs, the main applications are always to 'lattice-like' graphs. These graphs have a finite number of 'types' of vertices and of edges. Occasionally, we may select vertices or edges of different types with different probabilities.

For a fixed underlying graph Λ, there is a natural coupling of the measures $\mathbb{P}^{\mathrm{b}}_{\Lambda,p}$, $0 \le p \le 1$: take independent random variables X_e for each bond e of Λ, with X_e uniformly distributed on $[0,1]$. We may realize Λ^{b}_p as the spanning subgraph of Λ containing all bonds e with $X_e \le p$. In this coupling, if $p_1 < p_2$, then $\Lambda^{\mathrm{b}}_{p_1}$ is a subgraph of $\Lambda^{\mathrm{b}}_{p_2}$. A similar coupling is possible for site percolation.

An *open path* is a path (i.e., a self-avoiding walk) in the open subgraph. For sites x and y, we write '$x \to y$' or $\{x \to y\}$ for the event that there is an open path from x to y, and $\mathbb{P}(x \to y)$ for the probability of this event in the measure under consideration. We also write '$x \to \infty$' for the event that there is an infinite open path starting at x.

An *open cluster* is a component of the open subgraph. As the graphs

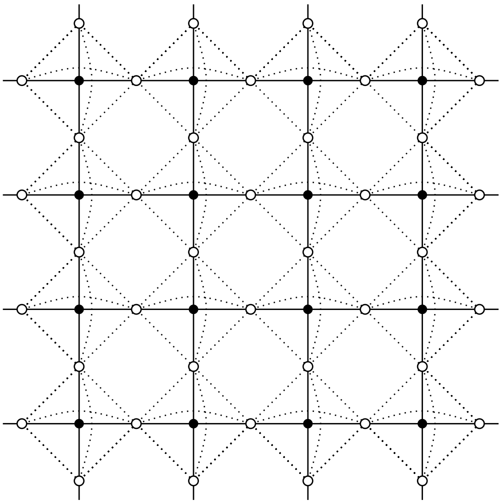

Figure 2. Part of the square lattice $\mathbb{Z}^2$ (solid circles and lines) and its line graph $L(\mathbb{Z}^2)$ (hollow circles and dotted lines). Note that $L(\mathbb{Z}^2)$ is isomorphic to the non-planar graph obtained from $\mathbb{Z}^2$ by adding both diagonals to every other face.

we consider are locally finite, an open cluster is infinite if and only if, for every site x in the cluster, the event $\{x \to \infty\}$ holds. Given a site x, we write C_x for the open cluster containing x, if there is one; otherwise, we take C_x to be empty. Thus $C_x = \{y \in \Lambda : x \to y\}$ is the set of sites y for which there is an open x–y path. Clearly, in bond percolation, C_x always contains x, and in site percolation, $C_x = \emptyset$ if and only if x is closed.

Let $\theta_x(p)$ be the probability that C_x is infinite, so $\theta_x(p) = \mathbb{P}_p(x \to \infty)$. Needless to say, $\theta_x(p)$ depends on the underlying graph Λ, and whether we take bond or site percolation. More formally, for bond percolation, for example,

$$\theta_x(p) = \theta_x(\Lambda_p^{\mathrm{b}}) = \theta_x^{\mathrm{b}}(\Lambda; p) = \mathbb{P}_{\Lambda,p}^{\mathrm{b}}(|C_x| = \infty),$$

where $|C_x| = |V(C_x)|$ is the number of sites in C_x. We shall use whichever form of the notation is clearest in any given context. In future, we shall introduce such self-explanatory variants of our notation without further comment; we believe that this will not lead to confusion. Two sites x and y of a graph Λ are *equivalent* if there is an automorphism of Λ mapping x to y. When all sites are equivalent (i.e., the symmetry group of the graph Λ acts transitively on the vertices), we write $\theta(p)$ for $\theta_x(p)$ for any site x. The quantity $\theta(p)$, or $\theta_x(p)$, is sometimes known as the *percolation probability*.

Clearly, if x and y are sites at distance d, then $\theta_x(p) \ge p^d\theta_y(p)$, so either $\theta_x(p) = 0$ for every site x, or $\theta_x(p) > 0$ for every x. Trivially, from the coupling described above, $\theta_x(p)$ is an increasing function of p. Thus there is a *critical probability* p_{H}, $0 \le p_{\mathrm{H}} \le 1$, such that if $p < p_{\mathrm{H}}$, then $\theta_x(p) = 0$ for every site x, and if $p > p_{\mathrm{H}}$, then $\theta_x(p) > 0$ for every x. The notation p_{H} is in honour of Hammersley. When the model under consideration is not clear from the context, we write $p_{\mathrm{H}}^{\mathrm{s}}(\Lambda)$ for site percolation on Λ and $p_{\mathrm{H}}^{\mathrm{b}}(\Lambda)$ for bond percolation.

The component structure of the open subgraph undergoes a dramatic change as p increases past p_{H}: if $p < p_{\mathrm{H}}$ then the probability of the event E that there is an infinite open cluster is 0, while for $p > p_{\mathrm{H}}$ this probability is 1. To see this, note that the event E is independent of the states of any finite set of bonds or sites, so Kolmogorov's 0-1 law (see Theorem 1 in Chapter 2) implies that $\mathbb{P}_p(E)$ is either 0 or 1. If $p < p_{\mathrm{H}}$, so that $\theta_x(p) = 0$ for every x, then

$$\mathbb{P}_p(E) \le \sum_x \theta_x(p) = 0,$$

and if $p > p_{\mathrm{H}}$, then $\mathbb{P}_p(E) \ge \theta_x(p) > 0$ for some site x (and so for all sites), implying that $\mathbb{P}_p(E) = 1$. One says that *percolation occurs* in a certain model if $\theta_x(p) > 0$, so $\mathbb{P}_p(E) = 1$. With a slight abuse of terminology, we use the same word both for this particular event and for the measures studied; this is not ideal, but, as in so many subjects, the historical terminology is now entrenched.

To start with, the theory of percolation was concerned mostly with the study of critical probabilities, i.e., with the question of when percolation occurs. Now, however, it encompasses the study of much more detailed properties of the random graphs arising from percolation measures. In fact, great efforts are made to describe the structure of these random graphs at or near the critical probability, even when we cannot pin down the critical probability itself. In Chapter 7, we shall get a glimpse of the huge amount of work done in this area, although in a setting in which the critical probability is known.

The theory of percolation deals with infinite graphs, and many of the basic events studied (such as the occurrence of percolation) involve the states of infinitely many bonds. Nevertheless, it always suffices to consider events in *finite* probability spaces, since, for example,

$\theta_x(p) = \lim_{n\to\infty} \mathbb{P}_p(|C_x| \geq n)$. In this book, almost all the time, even the definition of the infinite product measure will be irrelevant.

For $p < p_{\mathrm{H}}$, the open cluster C_x is finite with probability 1, but its expected size need not be finite. This leads us to another critical probability, p_{T}, named in honour of Temperley. Again, we write $p_{\mathrm{T}}^{\mathrm{s}}(\Lambda)$ or $p_{\mathrm{T}}^{\mathrm{b}}(\Lambda)$ for site or bond percolation on Λ. For a site x, set

$$\chi_x(p) = \mathbb{E}_p(|C_x|),$$

where $\mathbb{E}_p$ is the expectation associated to $\mathbb{P}_p$. If all sites are equivalent, we write simply $\chi(p)$. Trivially, $\chi_x(p)$ is increasing with p, and, as before, $\chi_x(p)$ is finite for some site x if and only if it is finite for all sites. Hence there is a critical probability

$$p_{\mathrm{T}} = \sup\{p : \chi_x(p) < \infty\} = \inf\{p : \chi_x(p) = \infty\},$$

which does not depend on x. By definition, $p_{\mathrm{T}} \leq p_{\mathrm{H}}$. One of our aims will be to prove that $p_{\mathrm{T}} = p_{\mathrm{H}}$ for many of the most interesting ground graphs, including the lattices $\mathbb{Z}^d$, $d \geq 2$.

There are very few cases in which p_{H} and p_{T} are easy to calculate. The prime example is the *d-regular infinite tree*, otherwise known as the *Bethe lattice* (see Figure 3). For the purposes of calculation, it is more

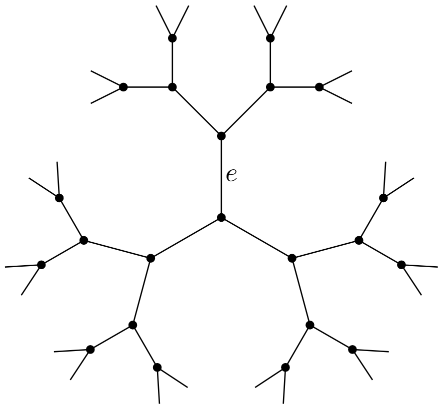

Figure 3. The 3-regular tree, for which $p_{\mathrm{T}}^{\mathrm{b}} = p_{\mathrm{H}}^{\mathrm{b}} = p_{\mathrm{T}}^{\mathrm{s}} = p_{\mathrm{H}}^{\mathrm{s}} = 1/2$. Deleting an edge ($e$, for example), this tree falls into two components, each of which is a 2-branching tree.

convenient to consider the *k-branching tree* T_k. This is the rooted tree in which each vertex has k children, so all sites but one have degree $k+1$. Writing v_0 for the root of T_k, let $T_{k,n}$ be the section of this tree up to

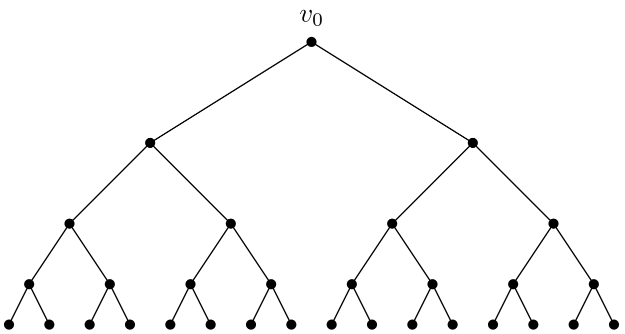

Figure 4. The tree $T_{2,4}$ with root v_0.

height (or, following the common mathematical convention of planting trees with the root at the top, depth) n, as in Figure 4. Taking the bonds to be open independently with probability p, let $\pi_n = \pi_{k,n}(p)$ be the probability that $T_{k,n}$ contains an open path of length n from the root to a leaf. Since such a path exists if and only if, for some child v_1 of v_0, the bond v_0v_1 is open and there is an open path of length $n-1$ from v_1 to a leaf, we have

$$\pi_n = 1 - (1 - p\pi_{n-1})^k = f_{k,p}(\pi_{n-1}). \tag{2}$$

On the interval $[0,1]$, the function $f_{k,p}(x)$ is increasing and concave, with $f_{k,p}(0) = 0$ and $f_{k,p}(1) < 1$, so $f_{k,p}(x_0) = x_0$ for some $0 < x_0 < 1$ if and only if $f'_{k,p}(0) = kp > 1$; furthermore, the fixed point x_0 is unique when it exists (see Figure 5). Thus, if $p > 1/k$, then, appealing to (2) we see

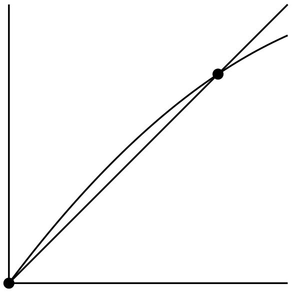

Figure 5. For $k = 2$ and $p = 2/3$, the increasing concave function $f(x) = f_{k,p}(x) = \frac{4}{3}x - \frac{4}{9}x^2$ satisfies $f' = \frac{4}{3} - \frac{8}{9}x$, $f(0) = 0$, $f(3/4) = 3/4$ and $f(1) = 8/9$.

that $\pi_{n-1} \geq x_0$ implies $\pi_n \geq x_0$. Since $\pi_0 = 1$, it follows that $\pi_n \geq x_0$ for every n, so $\theta^{\mathrm{b}}_{v_0}(T_k; p) \geq x_0 > 0$, implying that $p^{\mathrm{b}}_{\mathrm{H}}(T_k) \leq 1/k$. Also, if $p \leq 1/k$, then π_n converges to 0, the unique fixed point of $f_{k,p}(x)$, and so $\theta^{\mathrm{b}}_{v_0}(T_k; p) = 0$. Hence, the critical probability $p^{\mathrm{b}}_{\mathrm{H}}(T_k)$ is equal to $1/k$.

Turning to p_{T}, note that the probability that a site y at graph distance ℓ from the root v_0 belongs to C_{v_0} is exactly p^ℓ. Thus,

$$\chi^{\mathrm{b}}_{v_0}(T_k; p) = \mathbb{E}(|C_{v_0}|) = \sum_{y \in T_k} \mathbb{P}(y \in C_{v_0}) = \sum_{\ell=0}^{\infty} k^\ell p^\ell,$$

which is finite for $p < 1/k$ and infinite for $p \geq 1/k$. Thus the critical probability $p^{\mathrm{b}}_{\mathrm{T}}(T_k)$ is also equal to $1/k$.

For any infinite tree, after conditioning on the root x being open, the open clusters containing x in site and bond percolation have exactly the same distribution. Indeed, each child of a site in the open cluster lies in the open cluster with probability p. Thus, for the k-branching tree T_k, we have $p^{\mathrm{s}}_{\mathrm{H}} = p^{\mathrm{b}}_{\mathrm{H}} = p^{\mathrm{s}}_{\mathrm{T}} = p^{\mathrm{b}}_{\mathrm{T}} = 1/k$. It is easy to show similarly, or indeed to deduce, that the four critical probabilities associated to the $(k+1)$-regular tree are also equal to $1/k$.

The argument above amounts to a comparison between percolation on T_k and a certain branching process; we shall give a slightly less trivial example of such a comparison shortly. If Λ is any graph with maximum degree Δ, then a 'one-way' comparison with a branching process shows that all critical probabilities associated to Λ are at least $1/(\Delta - 1)$. To see this more easily, note that for every $y \in C_x$ there is at least one open path in Λ from x to y. Thus $\chi_x(p) = \mathbb{E}_p(|C_x|)$ is at most the expected number of open (finite) paths in Λ starting at x. There are at most $\Delta(\Delta - 1)^{\ell-1}$ paths in Λ of length ℓ starting at x, so

$$\chi^{\mathrm{b}}_x(p) \leq 1 + \sum_{\ell \geq 1} \Delta(\Delta - 1)^{\ell-1} p^\ell$$

and

$$\chi^{\mathrm{s}}_x(p) \leq p + \sum_{\ell \geq 1} \Delta(\Delta - 1)^{\ell-1} p^{\ell+1},$$

for bond and site percolation respectively. Both sums converge for any $p < 1/(\Delta - 1)$, so $p^{\mathrm{b}}_{\mathrm{T}}(\Lambda)$, $p^{\mathrm{s}}_{\mathrm{T}}(\Lambda) \geq 1/(\Delta - 1)$. As $p_{\mathrm{H}} \geq p_{\mathrm{T}}$, the corresponding inequalities for p_{H} follow. This shows that among all graphs with maximum degree Δ, the Δ-regular tree has the lowest critical probabilities.

There are various trivial changes we can make to a graph whose effect on the critical probability is easy to calculate. For example, if Λ is any graph and $\Lambda^{(\ell)}$ is obtained from Λ by subdividing each edge $\ell-1$ times, then $p_c^b(\Lambda^{(\ell)}) = p_c^b(\Lambda)^{1/\ell}$, where p_c^b is p_H^b or p_T^b. Also, if $\Lambda^{[k]}$ is obtained from Λ by replacing each edge by k parallel edges, then $1-p_c^b(\Lambda^{[k]}) = (1-p_c^b(\Lambda))^{1/k}$, where p_c^b is p_H^b or p_T^b. Of course, $p_c^s(\Lambda^{[k]}) = p_c^s(\Lambda)$. Combining these operations, we may replace each bond of a graph by k independent paths of length ℓ to obtain a new graph. For bond percolation, the critical probabilities p_{old} and p_{new} satisfy

$$1-(1-p_{\mathrm{new}}^{\ell})^k = p_{\mathrm{old}}.$$

In this way, by a trivial operation on the graph, a critical probability in the interval $(0,1)$ can be moved very close to any point of $(0,1)$.

If we know the critical probability for a graph Λ, then we know instantly the critical probabilities for a family of graphs Λ' obtained by sequences of trivial operations from Λ, as in Figure 6.

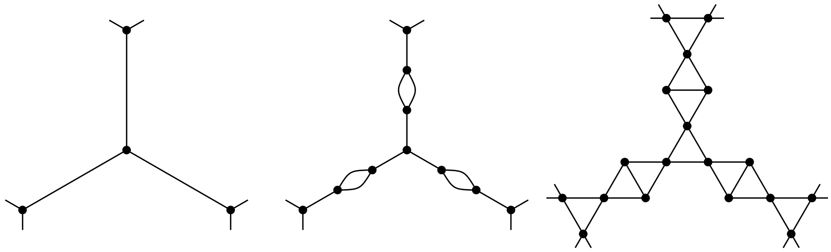

Figure 6. Transforming one bond percolation model into another, and then into a site percolation model. If the first (the hexagonal lattice) has critical probability p, then the second has critical probability r satisfying $r^3(2-r) = p$, which is also the critical probability for site percolation on the third graph.

It is easy to show that any $0 < \pi < 1$ is the critical probability for some graph, indeed, for some tree. Let T be a finite rooted tree with height (depth) h, with ℓ leaves. Let $T^1 = T$, and let T^n be the rooted tree of height hn formed from T^{n-1} by identifying each leaf with the root of a copy of T. For example, if T is a star with k edges, then T^n is the tree $T_{k,n}$ defined above. Let T^∞ be the 'limit' of the trees T^n, defined in the obvious way.

Taking the bonds of T to be open independently with probability p, the number of leaves of T joined to the root by open paths has a certain distribution X with expectation $p^h\ell$. Now suppose that the bonds of T^∞ are open independently with probability p, and let X_n be the number of sites of T^∞ at distance hn from the root joined to the root by open paths. Then the sequence $(X_0, X_1, X_2, \dots)$ is a branching process: we have $X_0 = 1$, and each X_n is the sum of X_{n-1} independent copies of the distribution X. As X is bounded, excluding the trivial case $p = \ell = 1$, it is easy to show (arguing as above for T_k) that percolation occurs if and only if $\mathbb{E}(X) > 1$, i.e., if and only if $p^h\ell > 1$; this is a special case of the fundamental result of the theory of branching processes. In fact, one obtains

$$p_{\mathrm{T}}^{\mathrm{s}}(T^\infty) = p_{\mathrm{H}}^{\mathrm{s}}(T^\infty) = p_{\mathrm{T}}^{\mathrm{b}}(T^\infty) = p_{\mathrm{H}}^{\mathrm{b}}(T^\infty) = \ell^{-1/h}. \qquad (3)$$

Suppose now that $k \ge 1$ and $1/(k+1) < \pi < 1/k$. Define $0 < \alpha < 1$ by $(k+1)^\alpha k^{1-\alpha} = 1/\pi$. Let $\mathbf{a} = (a_i)_{i=1}^\infty$ be the 0-1 sequence with density α constructed as follows: whenever 2^{j-1} divides i but 2^j does not, set $a_i = 1$ if and only if the jth bit in the binary expansion of α is 1. Let $T_{\mathbf{a}}$ be the rooted tree in which each site at distance i from the root has $k + a_{i+1}$ children. It is easy to check that, for each n, we can find trees T' and T'' of height $\ell = 2^n$ such that $(T')^\infty \subset T_{\mathbf{a}} \subset (T'')^\infty$, where T'' has $(k+1)/k$ times as many leaves as T'. Using (3), one can easily deduce that $p_{\mathrm{c}}(T_{\mathbf{a}}) = \pi$, where p_{c} denotes any of the four critical probabilities we have defined.

Alternatively, let T be the random rooted tree in which each site has $k+1$ children with probability r and k children with probability $1-r$, with the choices made independently for each site. It is easy to show that with probability 1 this random tree has $p_{\mathrm{c}}(T) = 1/(k+r)$.

In general, it is easy to calculate the various critical probabilities for a graph that is 'sufficiently tree-like'. For example, for $\ell \ge k \ge 3$, let $C_{k,\ell}$ be the *cactus* shown in Figure 7. This graph is formed by replacing each vertex of the k-regular tree T_k by a complete graph on ℓ vertices, and joining each pair of complete graphs corresponding to an edge of T_k by identifying a vertex of one with a vertex of the other, using no vertex in more than one identification. We call the vertices resulting from these identifications *attachment vertices*. Although $C_{k,\ell}$ contains many cycles, it still has the global structure of a tree, and percolation on $C_{k,\ell}$ may again be compared with a branching process.

Indeed, let K_ℓ be a complete graph with k distinguished (attachment) vertices $v_1, \dots, v_k$. Taking the edges of K_ℓ to be open independently with probability p, let X_p be the random number of vertices among $v_2, \dots, v_k$ that may be reached from v_1 by open paths. Let us explore the open cluster of a given initial site x of $C_{k,\ell}$ by working outwards from x. Except at the first step, from each attachment vertex that we