

BALANCED LOADS IN INFINITE NETWORKS¹

BY BRUCE HAJEK

University of Illinois at Urbana-Champaign

A set of nodes and a set of consumers are given, and to each consumer there corresponds a subset of the nodes. Each consumer has a demand, which is a load to be distributed among the nodes corresponding to the consumer. The load at a node is the sum of the loads placed on the node by all consumers. The load is balanced if no single consumer can shift some load from one node to another to reduce the absolute difference between the total loads at the two nodes. The model provides a setting to study the performance of load balancing as an allocation strategy in large systems.

The set of possible balanced load vectors is examined for infinite networks with deterministic or random demands. The balanced load vector is shown to be unique for rectangular lattice networks, and a method for computing the load distribution is explored for tree networks. An FKG-type inequality is proved. The concept of load percolation is introduced and is shown to be associated with infinite sets of nodes with identical load.

1. Introduction. Two simple examples illustrate the concept of balanced loads in infinite networks. Suppose that a unit load is associated with each edge of an infinite tree graph, in which each node has $d + 1$ neighbors, as pictured in Figure 1. Further, suppose that the load from each edge is allocated to the two endpoints of the edge and that the total load at each node is the sum of the loads assigned by the $d + 1$ incident edges. The resulting load is said to be balanced if, for each edge, the absolute difference of the total loads at the endpoints cannot be reduced by shifting some load from one endpoint to the other. For example, suppose each edge allocates all of its load to that one of its endpoints which is further from some designated reference node Δ . This results in one unit of load at each node except Δ , and no load at Δ . This load vector is not balanced. For example, a half unit of load could be shifted from one of the neighbors of node Δ to node Δ , and then the absolute difference between the load at the two nodes would be reduced from 1 to 0. One might continue shifting loads in an effort to obtain a balanced load vector. That process is *load balancing* and is an interesting topic in its own right, but the topic of this paper is *balanced loads*, which are loads that might be thought of as the output of a load balancing algorithm.


To continue with the example, suppose the original allocation is changed by shifting all the load for each edge along an infinite chain of edges leading

Received March 1994; revised June 1995.

¹This research was supported by the Joint Services Electronics Program Grant N00014-90-J-1270 and by NSF Grant NCR-93-14253.

AMS 1991 subject classification. Primary 60K35; secondary 60C05, 60G60, 60K30.

Key words and phrases. Load balancing, percolation, assignment problem.

FIG. 1. Portion of an infinite tree graph for $d = 2$.

from Δ , as indicated in Figure 2. This produces one unit of load at every node. Clearly, this is a balanced load vector. Is it unique? If $d \geq 2$ the answer is no. Another load vector can be obtained by modifying the assignment just described by shifting the load of every edge from one endpoint to the other. This yields the load vector identically equal to d , which if $d \geq 2$ is a distinct balanced load vector.

Let $x(v)$ denote the total load at a node v . We have shown that $x(v) \equiv 1$ and $x(v) \equiv d$ are each balanced load vectors (with two different corresponding assignments). By taking a convex combination of the two assignments, one observes that $x(v) \equiv \alpha$ is also a possible balanced load vector for any $\alpha \in [1, d]$. Intuitively speaking, the nonuniqueness of the balanced load is a result of the effect of the boundary condition at ∞ .

Is it possible that the load at a node can exceed d for a balanced load vector? The answer is no. In fact, let us indicate why for any balanced load vector the load $x(v)$ at any node v is in the interval $[1, d]$. In particular, the balanced load vector is unique if $d = 1$ and is not unique if $d \geq 2$. To begin, let F be a finite, connected subset of nodes. Then $\sum_{v \in F} x(v)$ is bounded above (below) by the number of edges with at least one (with both) endpoints




FIG. 2. A balancing assignment.

in F . This yields the following bounds on the average load carried by nodes in F :

$$(1.1) \quad \frac{k-1}{k} \leq \frac{\sum_{v \in F} x(v)}{k} \leq \frac{kd+1}{k},$$

where k is the cardinality of F . As $k \rightarrow \infty$, the lower and upper bounds in (1.1) tend to 1 and d , respectively. The fact that these two limits are distinct reflects the fact that large sets of nodes in infinite tree graphs have relatively large boundaries.

To finish showing that $x(v) \in [1, d]$ for all v , we argue by contradiction and suppose that there exists $\varepsilon > 0$ such that the set $\{v: x(v) \geq d + \varepsilon\}$ is nonempty. By the observation in the previous paragraph, this set cannot have arbitrarily large connected subsets. Therefore, it must have a finite component (where a component is a maximal connected subset), which we denote by F . The definition of balanced load implies that any edge that connects F and F^c assigns its entire unit load to the node in F^c . Thus, the average load in F is $(k-1)/k$, where k is the cardinality of F , which contradicts the assumption that $x(v) \geq d + \varepsilon$ for $v \in F$. Since ε is arbitrary, $x(v) \leq d$ for all v . A similar argument shows that $x(v) \geq 1$ for all v , so that the coordinates of any balanced load vector lie within the interval $[1, d]$.

Thus, the interval $[1, d]$ is the set of possible values of the load at a given node for a balanced load vector. The reader is invited to explore further the set of all balanced load vectors—not all the vectors are constant. Roughly speaking, the two extreme balanced load vectors $x(v) \equiv 1$ and $x(v) \equiv d$ were obtained by imposing minimal and maximal, respectively, boundary conditions at ∞ . Much more complex balanced load vectors exist for this example and can be obtained by imposing mixed boundary conditions at ∞ .

Henceforth, we shall call the amount of load to be assigned that is associated with a given edge the *demand* for the edge. Thus, in our first example the demand for each edge is deterministic and is one unit. In this paper a method is given for computing the distribution of the total load at a specified node in the infinite tree graph, with certain boundary conditions at ∞ , when the edge demands are independent and identically distributed. The method works particularly well when the demand distribution is concentrated on the set of integer multiples of a positive number. For example, if the demand at each edge is 1 with probability p and 0 otherwise, we find that the distribution of load at a node is not unique (i.e., it depends on the boundary conditions at ∞) if and only if $pd > 1$. As is well known from the theory of branching processes, this condition is also necessary and sufficient for the existence of infinite connected components in the subgraph formed by the edges with unit demand. The method reduces to a simple one-dimensional fixed-point equation in case the distribution of per-edge demand is the exponential distribution. We find that for such demand distribution the distribution of load at a node is not unique if and only if $d \geq 6$.

The second example of an infinite network which we present to illustrate balanced loads in infinite networks is based on the d -dimensional rectangular

lattice, for $d \geq 1$. The set of nodes is $V = Z^d$ and the edges are the pairs of nodes at unit Euclidean distance from each other. Again, assuming one unit of load is to be assigned for each edge, by assigning half of the load of each edge to each of its endpoints, a balanced load vector is obtained with load $x(v) = d$ for all nodes v . Once again, our next question is, is the balanced load vector unique?

To gain some insight into this question let us proceed as in the case of the infinite tree graph. If V_n is an n -cube (i.e., a subset of n^d nodes which is a translation of $\{1, \dots, n\}^d$), then the average load per node in V_n satisfies the bounds

$$(1.2) \quad d - \frac{d}{n} \leq \frac{\sum_{v \in V_n} x(v)}{|V_n|} \leq d + \frac{d}{n}.$$

The lower (resp., upper) bound corresponds to assigning all load to the endpoint in V_n^c (resp., V_n) for each edge bridging V_n and V_n^c . As $n \rightarrow \infty$, the bounds in (1.2) tend to the same limit, in contrast to the limiting behavior of the bounds in (1.1). This suggests that the balanced load vector is unique. In fact, it is shown in this paper that the number of nodes in a graph with distance n from some reference node must grow geometrically with n in order for the load at the reference node (for balanced load vectors) to be nonunique. The growth rate is only polynomial in n for rectangular lattice networks, so that in such networks the balanced load vector is unique. Roughly speaking, for rectangular lattice networks, the effect of boundary conditions asymptotically vanishes as larger and larger sets of nodes are considered.

The main questions addressed in this paper can be stated in broad terms as follows. How can the set of balanced load vectors be characterized? It is not difficult to show that balanced load vectors exist, but are they unique? What is the distribution of the load at a given node for a balanced load vector when the demand vector is random? Finally, the notion of balanced load concerns local conditions. What ‘‘global’’ or long-range effects can be observed in balanced load vectors?

Long-range effects are certainly evident for the infinite tree networks, for the load at a node can depend on the boundary conditions at ∞ . Long-range effects for the d -dimensional rectangular lattice networks can still exist, but they must be more subtle, given the uniqueness of balanced load vectors. For example, given an n -cube V_n and unit demand vector, it is not difficult to show there exists a load vector such that $x(v) = d - d/n$ for all $v \in V_n$, and another such that $x(v) = d + d/n$ for all $v \in V_n$. The fact that these two load vectors are distinct for any finite n indicates that the boundary condition has an effect as long as n is finite. We call this condition *load percolation* (defined more precisely in Section 3). Load percolation occurs when there is more than one balanced load vector, but, as this example shows, load percolation can occur even if the balanced load vector is unique. It is shown in this paper that if load percolation occurs, then there is a balanced load vector and an infinite

connected set of nodes which have identical load. Intuitively, infinite components are necessarily associated with load percolation since only when neighboring nodes have identical loads can arbitrarily small changes in load at one node negate the balance condition unless the load at a neighboring node is also changed. Thus, the long-range dependence inherent in load percolation can only be “transmitted” through infinite components of equal load.

This paper leaves open the problem of whether load percolation can occur for some random, independent identically distributed loads on the edges of a rectangular lattice network with $d \geq 2$. Some computer experiments are described next which partially investigate this question. Let $0 \leq p \leq 1$ and suppose the per-edge demands are independent Bernoulli random variables with parameter p . As indicated above, the balanced load vector x is unique for each realization of the demands and if $p = 1$ load percolation occurs. Whether load percolation occurs at a given node with positive probability for some p strictly less than 1 is an open problem. Of course, p has to be large enough ($p > 0.5$ if $d = 2$) so that there is positive probability of percolation in the classical sense [4]. Simulation results concerning the problem are presented here. One approach to simulation, motivated by the fact that infinite components of identical load are associated with load percolation, is to simulate a finite rectangular grid network with different boundary conditions and vary them in order to maximize the size of the components of the network with equal load. A similar approach, which we decided to use, is to simulate the load on a torus, rather than on a finite two-dimensional grid. Intuitively speaking, using a torus is similar to using a finite grid network and imposing boundary conditions which favor the formation of large sets of nodes with equal load.

Simulation data for a 60×60 torus is shown in Figures 3 to 5 for $p = 0.6, 0.75$ and 0.90 , respectively. The figures were produced as follows. A demand vector was generated using a pseudo-random number generator, and the balanced load vector was computed. The multiplicity of each load value was determined, and the three values of highest multiplicity were identified. Finally, the nodes with load values of the three highest multiplicities are indicated in the three figures. For $p = 0.6$, note that only a small fraction of the nodes have an identical load. This was observed in other simulations for $p < 0.6$, if load values of 0 and 0.5 are neglected. These values are not associated with load percolation but are quite frequent for small values of p . At the other extreme, for $p \geq 0.90$, we find that the balanced load vector is constant over a vast majority of the nodes, as shown in Figure 5. As p was increased from 0.6 to 0.9, we observed the emergence and growth of clusters of nodes having the same load values. Figure 4 is representative of this intermediate range of values.

The simulation data lead to the conjecture that percolation of load occurs for the infinite two-dimensional lattice with Bernoulli demands if p is sufficiently close to 1.0. Moreover, the data suggest that for $d = 2$ there is a corresponding cutoff value somewhere in the interval 0.75–0.90, but it seems unwise to draw any definite conclusions. Let us endeavor to speculate even




FIG. 3. Plot of balanced load vector for 60×60 torus network with Bernoulli ($p = 0.60$) demands. The highest multiplicity values are 1.0000000, 1.3333333 and 1.2500000, with multiplicities 740, 303 and 140, respectively.

further. Load percolation is likely to occur for smaller values of p in higher dimensions. The total load within a cube of side n and d dimensions has standard deviation on the order of $n^{d/2}$, whereas the number of edges crossing the boundary of such a cube is on the order of n^{d-1} . In two dimensions these are equal, so that, intuitively speaking, the balancing capability of edges might just be able to smooth out load fluctuations,




FIG. 4. Plot of balanced load vector for 60×60 torus network with Bernoulli ($p = 0.75$) demands. The highest multiplicity values are 1.5722121, 1.5865169 and 1.5000000, with multiplicities 547, 445 and 318, respectively.




FIG. 5. Plot of balanced load vector for 60×60 torus network with Bernoulli ($p = 0.90$) demands. The highest multiplicity values are 1.8059028, 1.7986348 and 1.8053097, with multiplicities 2880, 293 and 113, respectively.

whereas in three or more dimensions there seem to be more than enough edges available to smooth the load over long ranges. Perhaps load percolation does not occur for any $p < 1$ unless $d \geq 3$.

The model in this paper was studied for finite networks in [5]. A motivation of this work is to determine the effectiveness of local balancing mechanisms. The emergence of large, even infinite, sets of nodes with identical load indicates that global balancing is induced by local adjustment. The present paper deals with a static problem. Related dynamic problems studied for finite networks are the problem of computing a balanced load vector [5] and the problem of dynamically allocating traffic in a loss network in a balanced fashion [1, 3].

The organization and other results of the paper are as follows. Sections 2 to 4 concern networks with deterministic demand vectors and Sections 5 and 6 concern networks with random demand vectors. Finite networks are considered in Section 2. First, the notation, including a more general model, is given. In the examples above a demand is associated with an edge, which allocates load to a set of two nodes, whereas in the more general model introduced in Section 2 a demand is associated with a more general entity, henceforth called a *consumer*, which is to allocate load among a finite set of nodes. The model of Section 2 is more general also in that it provides for a constraint on the maximum load that a consumer can assign to a node. It is shown in Section 2 that balanced load vectors are unique for finite networks, and are associated with solutions of a convex optimization problem. In addition, the balanced load vectors are monotone in the demands.

General considerations for infinite networks, including the introduction of boundary conditions and the definition of load percolation, are given in

Section 3. It is shown, for example, that there is a minimal and maximal load vector. Section 3 also contains the result, described above, that load percolation implies the existence of infinite components of nodes with identical load. Section 4 contains the proof that nonuniqueness of balanced load requires geometric growth in the number of nodes within distance n of a given node.

General considerations for networks with random demand are given in Section 5. Two results are given. First, the existence of minimal and maximal load vectors for deterministic demand translates into the existence of minimal and maximal distribution functions of balanced load at a given node, in case the demands are random. Furthermore, any distribution function in between the minimal and maximal distribution functions can arise as the distribution of load at the node. Second, the monotonicity of loads with demands described in Sections 3 and 4 immediately implies an FKG-type inequality in case of independently distributed demands. Among other things, the inequality implies that the load values at distinct nodes are positively correlated. This result might be anticipated on the basis of everyday experience. For example, usually either all the tellers at a bank are busy or all are lightly loaded. A similar statement may be observed for computer systems or communication networks with dynamic load balancing.

Section 6 describes how to compute the distribution of load at a given node in an infinite tree network when the per-edge demands are independent and identically distributed, as described above. Open problems are given in Section 7.

2. Notation and finite networks. A consumer-demand network is a triple (U, V, C) where U and V are finite or countably infinite sets and $C = (C_{u,v} : u \in U, v \in V)$, where $0 \leq C_{u,v} \leq +\infty$ for all u, v . Defining $N(u)$ by $N(u) = \{v \in V : C_{u,v} > 0\}$, we assume that $N(u)$ is finite for each $u \in U$, and similarly that $\{u \in U : C_{u,v} > 0\}$ is finite for each $v \in V$. Elements of U are called consumers, elements of V are called nodes and $C_{u,v}$ is an upper bound on the load of consumer u that can be assigned to node v . A demand vector m for (U, V, C) has the form $m = (m_u : u \in U)$ where $0 \leq m_u \leq \sum_v C_{u,v}$ for $u \in U$, and a baseload vector b for (U, V, C) has the form $b = (b_v : v \in V)$, where b_v is real valued for each v .

In the examples in Section 1, the edges of the graphs played the role of consumers, and $N(u)$ for a given edge was just the set consisting of the two endpoints of the edge. The baseload vector was taken to be 0.

An admissible assignment vector is a vector $f, f = (f_{u,v} : u \in U, v \in V)$, such that $0 \leq f_{u,v} \leq C_{u,v}$. The interpretation is that $f_{u,v}$ denotes the load at node v due to consumer u . An assignment vector f meets the demand if

$$(2.1) \quad \sum_{v \in V} f_{u,v} = m_u \quad \text{for } u \in U,$$

and the total load at node $v, x(v)$, is given by

$$(2.2) \quad x(v) = b_v + \sum_{u \in U} f_{u,v} \quad \text{for } v \in V.$$

Use of the baseload vector is not a big generalization—it is nearly equivalent to having more consumers, each with only one vertex in its set $N(u)$. It is not exactly equivalent since the baseload values are allowed to be negative. The notion of baseload vector is convenient for some of the proofs.

A vector $x = (x(v): v \in V)$ so arising from an admissible assignment vector f meeting the demand is called a *load vector*. A load vector x is said to be *balanced*, if, for some corresponding f , the following conditions hold. For all $u \in U$ and all $v, v' \in V$, $f_{u,v}(C_{u,v'} - f_{u,v'}) = 0$ whenever $x(v) > x(v')$. If these conditions are satisfied we say x is a balanced load vector and that f balances. Intuitively, the condition that x is a balanced means that no consumer u can shift some of its load from one node v to another node v' to reduce the absolute difference $|x(v) - x(v')|$ of loads at the two nodes. Specifically, assuming $x(v) > x(v')$, either $f_{u,v} = 0$ so that u cannot reduce the load it assigns to v , or $f_{u,v'} = C_{u,v'}$ so that u cannot increase the load it assigns to v' .

An example network is pictured in Figure 6. There are three consumers and five nodes. Suppose that the baseloads b_v are all 0 and that the capacities $C_{u,v} = 1$ whenever $v \in N(u)$. The upper picture in the figure shows an admissible assignment f and the resulting load vector. The assignment f does not balance because, for example, the first consumer could shift some load from node v_2 to node v_1 to reduce $|x(v_2) - x(v_1)|$. The lower picture in the figure shows an assignment vector \tilde{f} which balances the load. Indeed, the balance conditions are satisfied for the first consumer, since although $x(v_1)$ is less than $x(v_2)$ and $x(v_3)$, the first consumer cannot shift load to v_1 since it already assigns load to v_1 at the full unit capacity.




FIG. 6. A consumer-demand network with given demand and two assignments.

Similarly, the second consumer cannot shift more load to node v_4 , and the third consumer cannot shift more load to v_5 . Finally, the third consumer cannot shift load for v_3 to v_4 since it already assigns zero load to v_3 . Note that there are other assignment vectors that balance. For example, the first two consumers could each shift equal amounts of load between nodes v_2 and v_3 in opposite directions, yielding a new balancing assignment vector.

If $C_{u,v} \geq m_u$ for all u, v with $v \in N(u)$, then all that is relevant about C is summarized in the set of sets $N = (N(u); u \in U)$ and we call (U, V, N) a consumer-demand network with no capacity constraints. The examples of Section 1 were of this type. The balancing condition can be simplified as follows. For all $u \in U$ and all $v, v' \in N(u)$, $f_{u,v} = 0$ whenever $x(v) > x(v')$.

This completes the description of the model. The definitions above hold for finite or infinite networks, but for the remainder of this section attention is focused on *finite* consumer-demand networks. First, an optimization problem is introduced which facilitates the proof that balanced load vectors are unique for finite consumer-demand networks.

Let ϕ be a convex function on the real line and define $J(f)$ for an assignment vector f by

$$J(f) = \sum_{v \in V} \phi(x(v)),$$

where x is the load vector given by (2.2). Define a convex optimization problem P by

PROBLEM P : $\min\{J(f): f \text{ is an admissible assignment meeting demand } m\}$.

PROPOSITION 2.1. (*Assume U and V are finite.*) *If f is a solution to problem P for a strictly convex function ϕ , then f balances. Conversely, if an assignment vector f balances, it solves problem P . There exist a unique balanced load vector x and any f corresponding to x balances. In the componentwise ordering of vectors, the balanced load vector x is increasing in m and b .*

PROOF. (This result, except for the monotonicity part, was presented in [5], but the proof given here is improved. The separable form assumed for J can be weakened to the assumption that J is symmetric in its arguments [1].) Suppose that ϕ is strictly convex and that f is an admissible assignment vector meeting the demand m . If f does not balance, then there are $u \in U$ and $v, v' \in V$ such that $f_{u,v} > 0$, $C_{u,v'} > f_{u,v'}$ and $x(v) > x(v')$. Hence, decreasing $f_{u,v}$ by ε and increasing $f_{u,v'}$ by ε for a sufficiently small value of ε leads to an admissible assignment with a strictly smaller value of $J(f)$, implying that f does not solve P . Thus, if f does not balance, it does not solve problem P . Equivalently, if f solves problem P , then f balances.

Conversely, suppose that f balances. It is shown next that f solves P if ϕ is (not necessarily strictly) convex. Since ϕ can be approximated arbitrarily closely, uniformly on bounded sets, by continuously differentiable convex

functions and since the coordinates of x are bounded as f ranges over assignments meeting the demand m , it is without loss of generality that we assume ϕ to be continuously differentiable. For each $u \in U$ define $\hat{x}(u)$ and π_u by $\hat{x}(u) = \max\{x(v) : v \in V, f_{u,v} > 0\}$ and $\pi_u = \phi'(\hat{x}(u))$. [The constants $(\pi_u : u \in U)$ are Lagrange multipliers.]

Let $f + h$ be any admissible assignment meeting the demand m . To prove that $J(f + h) \geq J(f)$, since $J(f + \varepsilon h)$ is convex in ε , it is enough to show that the limit of $(J(f + \varepsilon h) - J(f))/\varepsilon$ as ε decreases to 0 is nonnegative. The limit is equal to $\sum_{u,v} h_{u,v} \phi'(x(v))$, which, since $\sum_v h_{u,v} = 0$ for $u \in U$, is equal to

$$(2.3) \quad \sum_{u,v} h_{u,v} [\phi'(x(v)) - \pi_u].$$

Consider first a pair u, v such that $h_{u,v} > 0$. Then $f_{u,v} < C_{u,v}$ since $f + h$ is admissible, so that $\hat{x}(u) \leq x(v)$ since x is balanced, which, in turn, implies that $\pi_u \leq \phi'(x(v))$. On the other hand, if $h_{u,v} < 0$, then $f_{u,v} > 0$ since $f + h$ is admissible, so that $\hat{x}(u) \geq x(v)$ by the definition of $\hat{x}(u)$, which, in turn, implies that $\pi_u \geq \phi'(x(v))$. Therefore, all terms of the sum in (2.3) are nonnegative. Thus, if f balances it solves P .

The function J in problem P depends on f only through the load vector x . In addition, the set of load vectors, obtained as the image of all admissible assignments f meeting the demand, is a compact, convex set. Thus, if ϕ is strictly convex and continuous, problem P is equivalent to maximizing a strictly convex, continuous function over a compact, convex set. Thus, a balanced load vector exists and it is unique. Any assignment f corresponding to the balanced load vector solves P , and hence balances.

It remains to prove the claimed monotonicity. Suppose that x and \tilde{x} are the balanced load vectors for demand vectors and baseloads (m, b) and (\tilde{m}, \tilde{b}) , respectively, such that $b \leq \tilde{b}$ and $m \leq \tilde{m}$. Let f and \tilde{f} denote corresponding balancing assignments.

It must be shown that $x \leq \tilde{x}$. For the sake of argument by contradiction, suppose $\varepsilon = \max_v [x(v) - \tilde{x}(v)] > 0$. Let $F = \{v : x(v) - \tilde{x}(v) = \varepsilon\}$. Lemma 2.1 below implies that, for $u \in U$,

$$(2.4) \quad \sum_{v \in F} f_{u,v} \leq \sum_{v \in F} \tilde{f}_{u,v}.$$

Summing each side of (2.4) over $u \in U$ and using the ordering of baseloads yields $\sum_{v \in F} x(v) \leq \sum_{v \in F} \tilde{x}(v)$, which contradicts our assumption that $\varepsilon > 0$. \square

LEMMA 2.1. *Fix $u \in U$. Suppose (x, f) satisfies the following conditions: f is an admissible assignment for (U, V, C) and $f_{u,v}(C_{u,v'} - f_{u,v'}) = 0$ whenever $x(v) > x(v')$. (The vector x need not be the load vector for f .) Suppose (\tilde{x}, \tilde{f}) satisfies the same conditions and that $m_u = \sum_v f_{u,v} \leq \tilde{m}_u = \sum_v \tilde{f}_{u,v}$. Given*

$\varepsilon \geq 0$, suppose F is a set of nodes such that

$$\begin{aligned} x(v) - \tilde{x}(v) &\geq \varepsilon & \text{if } v \in F \cap N(u), \\ x(v) - \tilde{x}(v) &< \varepsilon & \text{if } v \in F^c \cap N(u). \end{aligned}$$

Then $\sum_{v \in F} f_{u,v} \leq \sum_{v \in F} \tilde{f}_{u,v}$.

PROOF. Let $r = \max\{x(v): f_{u,v} > 0\}$. The definitions of F and r imply the following two relations:

$$(2.5) \quad \{v \in F^c: x(v) < r\} \cap N(u) \supset \{v \in F^c: \tilde{x}(v) \leq r - \varepsilon\} \cap N(u),$$

$$(2.6) \quad \{v \in F: x(v) \leq r\} \cap N(u) \subset \{v \in F: \tilde{x}(v) \leq r - \varepsilon\} \cap N(u).$$

To complete the proof, the following three inequalities are justified below:

$$(2.7) \quad \sum_{v \in F} f_{u,v} \leq \left(m_u - \sum_{v \in F^c: x(v) < r} C_{u,v} \right) \wedge \left(\sum_{v \in F: x(v) \leq r} C_{u,v} \right)$$

$$(2.8) \quad \leq \left(\tilde{m}_u - \sum_{v \in F^c: \tilde{x}(v) \leq r - \varepsilon} C_{u,v} \right) \wedge \left(\sum_{v \in F: \tilde{x}(v) \leq r - \varepsilon} C_{u,v} \right)$$

$$(2.9) \quad \leq \sum_{v \in F} \tilde{f}_{u,v}.$$

Inequality (2.7) follows from the definition of r and the balancing conditions for (x, f) . Inequality (2.8) follows from the assumption $m_u \leq \tilde{m}_u$, (2.5) and (2.6). Finally, inequality (2.9) follows from the balancing condition for (\tilde{x}, \tilde{f}) . \square

Given a threshold L and a load vector x , the load in excess of L is defined by $\sum_v (x(v) - L)_+$, the maximum load is $\max_v x(v)$ and the minimum load is $\min_v x(v)$.

COROLLARY 2.1. *The balanced load vector minimizes the load in excess of L for any L . It also minimizes the maximum load and maximizes the minimum load.*

PROOF. The first assertion of the corollary follows from Proposition 2.1 with the choice $\phi(z) = (z - L)^+$. Since the resulting value of $J(f)$ is 0 if and only if the maximum load is less than or equal to L , it follows that the balanced load vector minimizes the maximum load. Similarly, since the balanced load vector minimizes J for $\phi(z) = (L - z)^+$ and any choice of L , it must also maximize the minimum load. \square

3. Boundary conditions, existence of balanced load vectors and load percolation. A simple method to produce balanced load vectors for consumer-demand networks with infinitely many nodes is to construct balanced load vectors on finite subsets of nodes and then take convergent subsequences. Different limits can sometimes result if different boundary

conditions are placed on the finite subsets. Hence, a notion of boundary conditions is presented before the existence of balanced load vectors is considered. This approach is quite common in the theory of statistical mechanics. The section concludes with a notion of long-range influence, called load percolation, and a demonstration that load percolation can occur only if there are infinite sets of nodes with equal load for some balanced load vector.

Let (U, V, C) be a consumer-demand network with demand vector m and baseload vector b . Let f be an admissible assignment vector satisfying the demand m and let x denote the corresponding load vector. Given a subset V_0 of V , we say that f balances in V_0 (with unspecified boundary condition) if the following condition holds: for all $u \in U$ and all $v, v' \in V_0$, $f_{u,v}(C_{u,v'} - f_{u,v'}) = 0$ whenever $x(v) > x(v')$. We say that f balances in V_0 with β -boundary condition for some $\beta \in [-\infty, +\infty]$ if: for all $u \in U$ and all $v, v' \in V$, $f_{u,v}(C_{u,v'} - f_{u,v'}) = 0$ whenever one of the following three conditions holds:

$$\begin{aligned} x(v) > x(v') \quad \text{and} \quad v, v' \in V_0, \\ x(v) > \beta, \quad v \in V_0 \quad \text{and} \quad v' \in V_0^c, \\ \beta > x(v'), \quad v \in V_0^c \quad \text{and} \quad v' \in V_0. \end{aligned}$$

Note that the second and third conditions are similar to the first, except if a node is in V_0^c , then the corresponding load value is replaced by β . Roughly speaking, the type of boundary condition we impose is similar to an infinite heat bath in statistical physics, in which the nodes outside the finite set v_0 are perceived to be held at the constant value β . If f balances the load in V_0 with β -boundary condition, then x is said to be balanced in V_0 with β -boundary condition.

LEMMA 3.1. *Assume V_0 is finite. There exists a load vector x^β that balances in V_0 with β -boundary condition. The restriction of x^β to V_0 , $(x^\beta(v): v \in V_0)$, is unique (for given β) and is nondecreasing in m, b and β . If x is balanced in V_0 with unspecified boundary condition, then $x^{-\infty}(v) \leq x(v) \leq x^{+\infty}(v)$ and all $v \in V_0$. For all $v \in V_0$ and $\beta \in [-\infty, +\infty]$,*

$$(3.1) \quad x^\beta(v) = \begin{cases} x^{+\infty}(v), & \text{if } x^{+\infty}(v) \leq \beta, \\ \beta, & \text{if } x^{-\infty}(v) \leq \beta \leq x^{+\infty}(v), \\ x^{-\infty}(v), & \text{if } x^{-\infty}(v) \geq \beta. \end{cases}$$

In other words, for each $v \in V_0$, $x^\beta(v)$ is the point in the interval $[x^{-\infty}(v), x^{+\infty}(v)]$ nearest to β .

PROOF. Assume to begin that β is finite and define the function J_{β, V_0} by $J_{\beta, V_0}(f) = \sum_{v \in V_0} (x(v) - \beta)^2$. This function is a strictly convex and continuous function of $(x(v): v \in V_0)$, and the set of possible values of $(x(v): v \in V_0)$ as f varies is a compact, convex set. There thus exists a minimum, and the resulting load vector restricted to V_0 is unique. Reasoning as in the proof of Proposition 2.1, it can be shown that f balances the load in V_0 with β -boundary condition if and only if f minimizes J_{β, V_0} over the set of admissible assignments meeting demand m_u . The existence of x^β and the

uniqueness of $(x^\beta(v): v \in V_0)$ are established if β is finite. The balance conditions for $\beta = +\infty$ are equivalent to the balance conditions for any finite β larger than the maximum possible value of load for nodes in V_0 , so therefore $x^{+\infty}$ exists and its restriction $(x^{+\infty}(v): v \in V_0)$ is unique and is equal to $(x^\beta(v): v \in V_0)$ for all sufficiently large β . Similarly, $x^{-\infty}$ exists and its restriction $(x^{-\infty}(v): v \in V_0)$ is unique.

Let x be an arbitrary load vector that balances in V_0 . To establish that $x^{-\infty}(v) \leq x(v)$ for $v \in V_0$, we view the restrictions of x and $x^{-\infty}$ to V_0 as balanced load vectors for a new network with set of nodes V_0 . The demand vector \tilde{m} for x restricted to V_0 is given by $\tilde{m}_u = m_u - \sum_{v \in V_0^c} f_{u,v}$, and the demand vector $\tilde{m}^{-\infty}$ for $x^{-\infty}$ restricted to V_0 is defined similarly. The base-loads for the two load vectors are the same: the restriction of b to V_0 . The β -boundary conditions with $\beta = -\infty$ imply that $\tilde{m}_u \leq \tilde{m}_u^{-\infty}$ for all u , so by Proposition 2.1 it follows that $x^{-\infty}(v) \leq x(v)$ for all $v \in V_0$. The same argument shows that $x(v) \leq x^{+\infty}(v)$ for all $v \in V_0$.

A proof that x^β is nondecreasing in m and b follows the proof of monotonicity given for Proposition 2.1. Specifically, suppose that x and \tilde{x} are balanced in V_0 with β -boundary conditions for (m, b) and (\tilde{m}, \tilde{b}) , respectively, where $b \leq \tilde{b}$ and $m \leq \tilde{m}$. Modify x and \tilde{x} by setting $x(v) = \tilde{x}(v) = \beta$ for all $v \in V_0^c$. [Of course, the modified vectors may no longer satisfy (2.2) unless $v \in V_0$, but that is not important.] It must be shown that $x \leq \tilde{x}$. The proof of this fact is word for word identical to the last two paragraphs of the proof of Proposition 2.1. The proof that x^β is nondecreasing in β follows from (3.1), which is established below.

It remains to establish (3.1). Fix a finite value of β . Let $V_1 = \{v \in V_0: x^\beta(v) > \beta\}$. The balancing conditions satisfied by x^β imply that x^β balances in V_1 with $-\infty$ -boundary condition. By the part of the proposition already proved, it follows that $x^\beta(v) \leq x(v)$ for all $v \in V_1$ and any x that balances in V_1 with unspecified boundary conditions. Taking $x = x^{-\infty}$ thus establishes that $x^\beta(v) \leq x^{-\infty}(v)$ for all $v \in V_1$. On the other hand, $x^\beta(v) \geq x^{-\infty}(v)$ for all $v \in V_1$ (in fact for all $v \in V_0$) so that $x^\beta(v) = x^{-\infty}(v)$ for all $v \in V_1$. The same argument shows that $x^\beta(v) = x^{+\infty}(v)$ for all $v \in V_0$ such that $x^\beta(v) < \beta$. Thus, (3.1) is proven for all $v \in V_0$ such that $x^\beta(v) \neq \beta$. Relation (3.1) is also true for $v \in V_0$ such that $x^\beta(v) = \beta$, by the extremality of $x^{-\infty}$ and $x^{+\infty}$ already proven. Thus, (3.1) is established for all $v \in V_0$. \square

Throughout the remainder of this section, let $(V_n: n \geq 1)$ denote a sequence of finite subsets of V such that $V_n \nearrow V$, which, by definition, means that $V_n \subset V_{n+1}$ for each n and $\bigcup_{n \geq 1} V_n = V$. By Lemma 3.1, for each n there exist load vectors $x_n^{-\infty}$, x_n^β and $x_n^{+\infty}$ which are balanced in V_n with $-\infty$ -boundary condition, β -boundary condition and $+\infty$ -boundary condition, respectively. Let $f_n^{-\infty}$, f_n^β and $f_n^{+\infty}$ denote the corresponding assignment vectors.

PROPOSITION 3.1. *Given a consumer-demand network (U, V, C) with demand vector m and base-load vector b , there exist balanced load vectors $x^{-\infty}$ and $x^{+\infty}$ such that, for any balanced load vector x , $x^{-\infty} \leq x \leq x^{+\infty}$ coordinate-*

wise. For any finite β relation (3.1) defines a balanced load vector x^β . The load vector x^β is nondecreasing in m, b and β . For any $\beta \in [-\infty, +\infty]$ and $v \in V$, $\lim_{n \rightarrow \infty} x_n^\beta(v) = x^\beta(v)$, and the convergence is monotone for large enough n .

PROOF. Since $0 \leq f_{u,v} \leq m_u$ for each u, v and any admissible assignment vector f meeting demand m , there is a subsequence of $n \rightarrow \infty$ such that, along the subsequence, each coordinate of the assignment vectors $f_n^{-\infty}, f_n^\beta$ and $f_n^{+\infty}$ converges. Let $f^{-\infty}, f^\beta$ and $f^{+\infty}$ denote the respective limits and let $x^{-\infty}, x^\beta$ and $x^{+\infty}$ denote the corresponding load vectors. The balancing conditions satisfied by $f_n^{-\infty}, f_n^\beta$ and $f_n^{+\infty}$ for each n imply that the limits $x^{-\infty}, x^\beta$ and $x^{+\infty}$ are balanced load vectors with corresponding assignment vectors $f^{-\infty}, f^\beta$ and $f^{+\infty}$. The load vectors obtained for V_n satisfy relation (3.1) for $v \in V_n$, so that (3.1) is satisfied by the limits $x^{-\infty}, x^\beta$ and $x^{+\infty}$ for all $v \in V$. Similarly, the load vectors obtained for V_n , when restricted to V_n , are nondecreasing in m and b , so that the limits $x^{-\infty}, x^\beta$ and $x^{+\infty}$ are also monotone in m and b . Equation (3.1) implies that x^β is also nondecreasing as a function of β .

Let x be an arbitrary balanced load vector. The extremality properties of $x_n^{-\infty}$ and $x_n^{+\infty}$ imply that $x_n^{-\infty} \leq x \leq x_n^{+\infty}$ on V_n for each n . In the limit $n \rightarrow \infty$ this yields $x^{-\infty} \leq x \leq x^{+\infty}$ as desired.

The extremality properties of $x_n^{-\infty}$ and $x_n^{+\infty}$ imply that $x_{n+1}^{-\infty} \geq x_n^{-\infty}$ on V_n and $x_{n+1}^{+\infty} \leq x_n^{+\infty}$ on V_n . Thus, for $v \in V$ and for n large enough that $v \in V_n$, the sequence $x_n^{-\infty}(v)$ is nondecreasing in n and the sequence $x_n^{+\infty}$ is nonincreasing in n . These facts, together with (3.1) for set V_n , imply that, for any $\beta \in [-\infty, +\infty]$ and $v \in V$, $\lim_{n \rightarrow \infty} x_n^\beta(v) = x^\beta(v)$, and the convergence is monotone for n large enough that $v \in V_n$. Of course, the corresponding assignment vectors need not converge without passing to a subsequence. \square

The concept of load percolation for the given network (U, V, C) , demand vector m and baseload vector b is defined as follows. *Load percolation* is said to occur at a fixed node v_0 if $x_n^{-\infty}(v_0) < x_n^{+\infty}(v_0)$ for all sufficiently large n (equivalently, for all n such that $v_0 \in V_n$). A sufficient but not necessary condition for load percolation at v_0 is $x^{-\infty}(v_0) < x^{+\infty}(v_0)$. The property of load percolation does not depend on which sequence of finite subsets $(V_n: n \geq 0)$ is used to define the sequences $(x_n^{-\infty})$ and $(x_n^{+\infty})$. In fact, a condition equivalent to load percolation at v_0 is that $x_A^{-\infty}(v_0) < x_A^{+\infty}(v_0)$ whenever A is a finite set of nodes containing v_0 and $x_A^{-\infty}$ and $x_A^{+\infty}$ are load vectors that balance in A with $-\infty$ - and $+\infty$ -boundary conditions, respectively.

In the remainder of this section, it is shown that, as described in the Introduction, load percolation is associated with infinite sets of nodes with identical loads. Some additional notation is introduced in order to state the result. Nodes v and v' are called neighbors if $v, v' \in N(u)$ for some consumer u . A path from v to v' is a sequence s of nodes $s = (v_1, \dots, v_{|s|})$ so that $v = v_1, v' = v_{|s|}$ and v_i and v_{i+1} are neighbors for $1 \leq i \leq |s| - 1$, and $|s|$ denotes the length of the path. The *distance* between two nodes v and v' is

the minimum length of a path from v to v' , unless no path exists in which case the distance is infinite. The boundary of a set V_0 of nodes, written ∂V_0 , is defined by

$$\partial V_0 = \{v' \in V - V_0 : v' \text{ is a neighbor of } v \text{ for some } v \in V_0\}.$$

A set V_0 is connected if, for any $v, v' \in V_0$, there is a path from v to v' . Finally, a component of a set V_0 is a maximal connected subset of V_0 .

PROPOSITION 3.2. *Suppose load percolation occurs at node v_0 . Let $\beta \in [x^{-\infty}(v_0), x^{+\infty}(v_0)]$. Then the component of $\{v : x^\beta(v) = \beta\}$ containing v_0 is infinite.*

PROOF. We prove the contrapositive. Let V^* denote the component in question and suppose that V^* is finite. The goal is to establish that load percolation does not occur at v_0 . Consider the partition of ∂V^* into two sets $\partial_+ V^*$ and $\partial_- V^*$ given by

$$(3.2) \quad \partial_+ V^* = \{v \in \partial V^* : x^\beta(v) > \beta\},$$

$$(3.3) \quad \partial_- V^* = \{v \in \partial V^* : x^\beta(v) < \beta\}.$$

Note that, by relation (3.1), $x^{+\infty}(v) = x^\beta(v)$ for $v \in \partial_- V^*$ and $x^{-\infty}(v) = x^\beta(v)$ for $v \in \partial_+ V^*$. Since $\bigcup_{n \geq 1} V_n = V$ and $x_n^{+\infty}(v) \rightarrow x^{+\infty}(v)$ as $n \rightarrow \infty$ for all $v \in V$, if n is sufficiently large, then $V^* \cup \partial V^* \subset V_n$ and $x_n^{+\infty}(v) < \beta$ for $v \in \partial_- V^*$.

We argue that, for such n , $x_n^{+\infty}(v) = x^\beta(v)$ for $v \in V^*$. On the one hand, $x_n^{+\infty}(v) \geq x^\beta(v)$ for $v \in V^*$, in fact for $v \in V_n$, since $x_n^{+\infty}$ dominates on V_n any load vector that balances on V_n . On the other hand, it will be shown at the end of the proof that, for any u ,

$$(3.4) \quad \sum_{v \in V^*} f_{u,v} \leq \sum_{v \in V^*} f_{u,v}^\beta,$$

where f and f^β denote assignment vectors corresponding to $x_n^{+\infty}$ and x^β , respectively. Summing each side of (3.4) over u shows that $\sum_{v \in V^*} x_n^{+\infty}(v) \leq \sum_{v \in V^*} x^\beta(v)$ for $v \in V^*$. This proves that $x_n^{+\infty}(v) = x^\beta(v)$ for $v \in V^*$ and sufficiently large n .

Similarly, $x_n^{-\infty}(v) = x^\beta(v)$ for $v \in V^*$ and sufficiently large n . Therefore, $x_n^{+\infty}(v) = x_n^{-\infty}(v)$ for sufficiently large n and all $v \in V^*$, in particular, for $v = v_0$. Thus, if V^* is finite, load percolation does not occur at v_0 .

It remains to establish (3.4) for given u . Without loss of generality, suppose $N(u) \cap V^* \neq \emptyset$. Since, f and f^β each satisfy the same demand,

$$(3.5) \quad \sum_{v \in V^*} (f_{u,v}^\beta - f_{u,v}) = T_+ + T_- - T_+^\beta - T_-^\beta,$$

where

$$\begin{aligned} T_+ &= \sum_{v \in \partial_+ V^*} f_{u,v}, & T_- &= \sum_{v \in \partial_- V^*} f_{u,v}, \\ T_+^\beta &= \sum_{v \in \partial_+ V^*} f_{u,v}^\beta, & T_-^\beta &= \sum_{v \in \partial_- V^*} f_{u,v}^\beta. \end{aligned}$$

Note that if $v \in V^* \cup \partial_+ V^*$ and $v' \in \partial_- V^*$, then $x_n^{+\infty}(v) \geq x^\beta(v) \geq \beta > x_n^{+\infty}(v')$ and $x^\beta(v) > x^\beta(v')$. Therefore, $T_- = T_-^\beta = m_u \wedge \sum_{v \in \partial_- V^*} C_{u,v}$. Note also that $x^\beta(v) > x^\beta(v')$ is $v \in \partial_+ V^*$ and $v' \in V^* \cup \partial_- V^*$. Therefore,

$$T_+ \geq T_+^\beta = \left(m_u - \sum_{v \in V^* \cup \partial_- V^*} C_{u,v} \right)_+.$$

In summary, $T_- = T_-^\beta$ and $T_+ \geq T_+^\beta$ which in view of (3.5) implies (3.4). \square

The next corollary follows immediately from Proposition 3.2 and the fact that $x^\beta \equiv x$ for all β if there is a unique balanced load vector x .

COROLLARY 3.1. *If there is a unique balanced load vector x and if load percolation occurs at a node v_0 , then the component of $\{v: x(v) = x(v_0)\}$ containing v_0 is infinite.*

4. Necessary condition for nonuniqueness. Let (U, V, C) be a consumer-demand network such that, for some constants A and B , $|N(u)| \leq A$ for $u \in U$ and $|\{u \in U: C_{u,v} > 0\}| \leq B$ for $v \in V$. Suppose a demand vector m is given with $m_u \geq 0$ for all u as usual. Throughout this section the base-load is taken to be 0. Fix a node v_0 and let S_n denote the set of nodes at distance less than or equal to n from v_0 (the notion of distance is defined just before Proposition 3.2). Let x and \tilde{x} denote balanced load vectors for (U, V, C) and m .

PROPOSITION 4.1. *If $x(v_0) > \tilde{x}(v_0)$, then there is a constant $a > 1$ such that $|S_n| \geq a^n$ for all n .*

A plausibility argument for the proposition is given in the Introduction. Our proof begins with the following lemma.

LEMMA 4.1. *Proposition 4.1 is true under the additional assumption $m_u \leq L$ for all $u \in U$ for some $L < \infty$.*

PROOF. Let $\varepsilon = x(v_0) - \tilde{x}(v_0)$, which, by assumption, is positive, and let G denote the component of the set $\{v: x(v) - \tilde{x}(v) \geq \varepsilon\}$ containing v_0 . For $n \geq 1$ set

$$s_n = \sum_{v \in S_n \cap G} x(v) - \tilde{x}(v).$$

By the relation between assignment vectors and load vectors,

$$(4.1) \quad s_n = \sum_u \left(\sum_{v \in S_n \cap G} f_{u,v} - \tilde{f}_{u,v} \right).$$

Consider the sum in parentheses on the right-hand side of (4.1). If $N(u) \cap S_n \cap G = \emptyset$, the sum is 0, so consider u such that $N(u) \cap S_n \cap G \neq \emptyset$. Then $N(u) \subset (S_n \cap G) \cup G^c$, so the sum is less than or equal to 0 by Lemma 2.1

with $F = S_n \cap G$. For any u the sum is less than or equal to m_u , which is at most L . Thus,

$$\begin{aligned} s_n &\leq L|\{u: N(u) \cap \partial S_n \cap G \neq \emptyset\}| \\ &\leq BL|\partial S_n \cap G|. \end{aligned}$$

On the other hand, $s_n \geq \varepsilon|S_n \cap G|$. Thus, $|\partial S_n \cap G| \geq (\varepsilon/BL)|S_n \cap G|$ for all $n \geq 1$. Since S_{n+1} is the union of the disjoint sets S_n and ∂S_n , $|S_{n+1} \cap G| \geq (1 + \varepsilon/BL)|S_n \cap G|$. By induction on n , $|S_n| \geq |S_n \cap G| \geq (1 + \varepsilon/BL)^n$ for all n . \square

Given $D > 0$, define $\bar{m}_u = m_u \wedge (AD)$ for $u \in U$.

LEMMA 4.2. *There exists a balanced load vector \bar{x} for demand \bar{m} such that $\bar{x}(v) \wedge D = x(v) \wedge D$ for all $v \in V$.*

PROOF. Assume without loss of generality that V is connected so that $S_n \nearrow V$. First, a (not necessarily balancing) assignment vector \hat{f} meeting demand \bar{m} is specified. Consider $u \in U$. If $m_u \leq AD$ let $\hat{f}_{u,v} = f_{u,v}$ for $v \in N(u)$. Otherwise, let $\hat{f}_{u,v} = f_{u,v} \wedge \gamma_u$ for $v \in N(u)$, where γ_u is the unique number such that $\sum_{v \in N(u)} \hat{f}_{u,v} = AD$. Observe that $\gamma_u \geq D$. The vector \hat{f} is an admissible assignment meeting demand \bar{m} . Let \hat{x} denote the load vector for \hat{f} . Although \hat{x} is not necessarily balanced, it satisfies $\hat{x}_v \wedge D = x_v \wedge D$ for all $v \in V$. Let $F = \{v: x(v) \geq D\}$. Note that if $v, v' \in N(u)$ for some u and if $v \in F$ and $v' \in F^c$, then $f_{u,v'}(C_{u,v'} - f_{u,v'}) = \hat{f}_{u,v'}(C_{u,v'} - \hat{f}_{u,v'}) = 0$ since x is balanced, $f_{u,v} \geq \hat{f}_{u,v}$ and $f_{u,v'} = \hat{f}_{u,v'}$.

For $n \geq 1$ consider the finite network with set of nodes $S_n \cap F$, set of consumers U , capacity vector equal to the restriction of C to $U \times (S_n \cap F)$ and demand $m_u - \sum_{v \in (S_n \cap F)^c} \hat{f}(u, v)$ for $u \in U$. Note that \hat{x} restricted to $S_n \cap F$ is a load vector for the finite network. Let $x^{(n)}$ be a *balanced* load vector for the finite network. Since \hat{x} satisfies $\hat{x}(v) \geq D$ for all $v \in S_n \cap F$, the same is true for $x^{(n)}$ by the minimum maximizing property (cf. Corollary 2.1) of balanced load vectors for finite networks.

Let $f^{(n)}$ denote the assignment vector for $x^{(n)}$. Since each coordinate of $f^{(n)}$ is bounded uniformly in n , there is a subsequence of $n \rightarrow \infty$ such that, along the subsequence, the vectors $f^{(n)}$, and hence also the vectors $x^{(n)}$, converge coordinatewise. Denote the limits by \bar{f} and \bar{x} . Note that $\bar{x}(v) \geq D$ for $v \in F$ since $x^{(n)}(v) \geq D$ for all $v \in F \cap S_n$. Extend \bar{f} to an assignment vector for (U, V, C) by setting $\bar{f}_{u,v} = \hat{f}_{u,v}$ ($= f_{u,v}$) if $v \in F^c$ and extend \bar{x} to be the load vector for \bar{f} . Then \bar{x} is a balanced load vector and $\bar{x}_u \wedge D = x_u \wedge D$ for all $u \in U$. The lemma is proved. \square

PROOF OF PROPOSITION 4.1. Choose $D = x(v_0)$. Then $x(v_0), \tilde{x}(v_0) \leq D$ so by Lemma 4.2, applied to both x and \tilde{x} , there exist two balanced load vectors for (U, V, C) with demand \bar{m} , one with load at v_0 equal to $x(v_0)$ and one with load at v_0 equal to $\tilde{x}(v_0)$. Hence, by Lemma 4.1 with $L = AD$, there exists $\alpha > 1$ such that $|S_n| \geq \alpha^n$ for all n . \square

5. Randomly loaded networks. Let (U, V, C) denote a consumer-demand network. Suppose that $b = (b_v: v \in V)$ is a random baseload vector and that $m = (m_u: u \in U)$ is a random demand vector such that $0 \leq m_u \leq \sum_{v \in V} C_{u,v}$ for all u with probability 1. Then for each $\beta \in [-\infty, +\infty]$ there is a balanced load vector x^β , which is random. Fix a node v_0 and let F^β denote the probability distribution function of $s^\beta(v_0)$. Write $F < G$ for two distribution functions F and G if $F(c) \geq G(c)$ for all real c .

PROPOSITION 5.1. *A distribution function F is the distribution of load at v_0 for some balanced load vector if and only if $F^{-\infty} < F < F^{+\infty}$ [assuming the probability space is large enough to support a uniform random variable independent of (m, b) for the “if” statement]. In particular, $F^\beta(c) = F^{+\infty}(c)I_{\{c < \beta\}} + F^{-\infty}(c)I_{\{c \geq \beta\}}$.*

PROOF. The fact that F is the distribution of $x(v_0)$ for some balanced load vector only if $F^{-\infty} < F < F^{+\infty}$ follows from the minimality of $x^{-\infty}$ and the maximality of $x^{+\infty}$. Conversely, let F be any distribution function satisfying $F^{-\infty} < F < F^{+\infty}$. Then by the well-known connection between stochastic ordering and stochastic coupling, if the underlying probability space supports a uniform random variable independent of (m, b) , there exists a random variable B such that $B \in [x^{-\infty}(v_0), x^{+\infty}(v_0)]$ with probability 1 and B has distribution function F . Letting x^B denote x^β evaluated at $B = \beta$, we have that x^B is a balanced load vector and $x^B(v_0) = B$ with probability 1 so that $x^B(v_0)$ has distribution function F . The final statement in the proposition follows from relation (3.1). \square

The monotonicity property stated in Proposition 3.1 immediately implies an FKG-type inequality for networks with independent demand and baseload variables.

PROPOSITION 5.2. *Suppose $b_v, v \in V$, and $m_u, u \in U$, are mutually independent random variables. Fix $\beta \in [-\infty, +\infty]$ (or let β be random and independent of m and b). Let g and h be nondecreasing, Borel measurable functions mapping R^V to R such that $E[g(x^\beta)^2]$ and $E[h(x^\beta)^2]$ are finite. Then $E[g(x^\beta)h(x^\beta)] \geq E[g(x^\beta)]E[h(x^\beta)]$.*

Taking $g(x) = x(v)$ and $h(x) = x(v')$ for two nodes v and v' , Proposition 5.2 yields that $\text{cov}(x^\beta(v), x^\beta(v')) \geq 0$. Of course, if there is a unique balanced load vector x with probability 1, then the FKG inequality holds for x^β replaced by x .

PROOF OF PROPOSITION 5.2. Since by Proposition 3.1 the balanced load vector is a nondecreasing function of b, m and β , the random variables $g(x^\beta)$ and $h(x^\beta)$ are also nondecreasing functions of b, m and β . The proposition hence follows from the FKG inequality for nondecreasing functions of independent random variables (see [2] and [4], Theorem 2.4). \square

6. Tree networks with random demand. Tree networks as considered in Section 1 are investigated in this section with independent random demands on the edges. First, the notion of τ -surplus of a node in a finite network is introduced. Intuitively, the τ -surplus is the quantity of load that must be “removed” from the node in order to make the load at the node (for the new balanced load vector) equal to τ . It is shown that the τ -surplus of nodes for finite subsets with boundary conditions can be computed recursively, yielding a method to compute numerically the distribution of load at a given node for finite or infinite tree networks. The basic technique was introduced in [5] for analysis of a random tree network with Poisson degree and unit demands. Two examples are investigated more closely, as outlined in Section 1—Bernoulli loads and exponentially distributed loads.

The τ -surplus of a node. Let (U, V, C) be a finite consumer-demand network with demand vector m and baseload b . Given a real number τ and a node $v_0 \in V$, define the τ -surplus of v_0 to be the unique value y so that if the baseload b_{v_0} is changed to $b_{v_0} - y$, then the load at v_0 for a balanced assignment is τ . The load at v_0 under a balanced assignment for the original baseload vector b is greater than (resp., less than, equal to) τ if and only if the τ -surplus of v_0 is greater than (resp., less than, equal to) 0. Hence, the τ -surplus of v_0 for all τ determines the load at v_0 for a balanced assignment.

Next suppose V is infinite. Fix a node v_0 and take a sequence of finite subsets $(v_n; n \geq 0)$ with $V_n \nearrow V$. Consider the τ -surplus at v_0 for the balancing problem in V_n as $n \rightarrow \infty$. Specifically, Let $Y_{\tau, n}^{+\infty}$ (resp., $Y_{\tau, n}^{-\infty}$) denote the τ -surplus of node v_0 for the balancing problem in V_n with $+\infty$ -boundary condition (resp., $-\infty$ -boundary condition). The variables $Y_{\tau, n}^{+\infty}$ (resp., $Y_{\tau, n}^{-\infty}$) are nonincreasing (resp., nondecreasing) in n and are bounded. Let $Y_{\tau}^{+\infty}$ (resp., $Y_{\tau}^{-\infty}$) denote the corresponding limits. We thus can write $Y_{\tau, n}^{+\infty} \searrow Y_{\tau}^{+\infty}$ and $Y_{\tau, n}^{-\infty} \nearrow Y_{\tau}^{-\infty}$, where the notation $a_n \searrow a$ (resp., $a_n \nearrow a$) denotes that the sequence (a_n) is nonincreasing (resp., nondecreasing) with limit a . Observe that $Y_{\tau, n}^{+\infty} \geq Y_{\tau, n}^{-\infty}$ for all n , so that $Y_{\tau}^{+\infty} \geq Y_{\tau}^{-\infty}$.

PROPOSITION 6.1. *For any finite τ ,*

$$(6.1) \quad P[x^{+\infty}(v_0) < \tau] = P[Y_{\tau}^{+\infty} < 0] = \lim_{n \rightarrow \infty} P[Y_{\tau, n}^{+\infty} < 0],$$

$$(6.2) \quad P[x^{-\infty}(v_0) \leq \tau] = P[Y_{\tau}^{-\infty} \leq 0] = \lim_{n \rightarrow \infty} P[Y_{\tau, n}^{-\infty} \leq 0].$$

PROOF. As in the proof of Proposition 3.1, let $x_n^{+\infty}$ denote a load vector that is balanced in V_n with $+\infty$ -boundary condition. As discussed after the proof of Proposition 3.1, $x_n^{+\infty}(v_0) \searrow x^{+\infty}(v_0)$. We thus have

$$\{x_n^{+\infty}(v_0) < \tau\} \nearrow \{x^{+\infty}(v_0) < \tau\}$$

and

$$\{x_n^{+\infty}(v_0) < \tau\} = \{Y_{\tau, n}^{+\infty} < 0\} \nearrow \{Y_{\tau}^{+\infty} < 0\}.$$

Consequently, $x^{+\infty}(v_0) < \tau$ if and only if $Y_{\tau}^{+\infty} < 0$, and $\{Y_{\tau, n}^{+\infty} < 0\} \nearrow \{Y_{\tau}^{+\infty} < 0\}$, which implies (6.1), and (6.2) is established similarly. \square

COROLLARY 6.1. *The distribution of $x(v_0)$ is not unique (i.e., not the same for all equilibrium load vectors x) if and only if there exists an interval of positive length such that $P[Y_\tau^{+\infty} < 0] < P[Y_\tau^{-\infty} \leq 0]$ for all τ in the interval. In particular, the distribution of $x(v_0)$ is not unique if there exists an interval of positive length such that $P[Y_\tau^{+\infty} < 0] < P[Y_\tau^{-\infty} < 0]$ for all τ in the interval.*

PROOF. The distribution of $x(v_0)$ is unique if and only if $x^{+\infty}(v_0)$ and $x^{-\infty}(v_0)$ have the same distribution function. If they do have the same distribution function, then whenever τ is not a point of discontinuity of the function, $P[Y_\tau^{+\infty} < 0] = P[x^{+\infty}(v_0) < \tau] = P[x^{-\infty}(v_0) \leq \tau] = P[Y_\tau^{-\infty} \leq 0]$. Since a distribution function has at most countably many points of discontinuity, there is no interval of positive length such that $P[Y_\tau^{+\infty} < 0] < P[Y_\tau^{-\infty} \leq 0]$ for all τ in the interval. Conversely, if $x^{+\infty}(v_0)$ and $x^{-\infty}(v_0)$ have different distribution functions, then there is an interval of positive length such that the left-continuous modification of the distribution function of $x^{+\infty}(v_0)$ lies strictly below the distribution function of $x^{-\infty}(v_0)$ throughout the interval. Consequently, $P[Y_\tau^{+\infty} < 0] < P[Y_\tau^{-\infty} \leq 0]$ for all τ in that interval. The first statement in the corollary is proved. The second statement in the corollary is an immediate consequence of the first and the fact that $P[Y_\tau^{-\infty} < 0] \leq P[Y_\tau^{-\infty} \leq 0]$ for all τ . \square

The following lemma is useful for calculating the τ -surplus of nodes in a tree network. Roughly speaking, it shows that if the τ -surplus of one specified node in each of several distinct networks is known and if a new network is formed by linking each of the specified nodes together to some new node, then the τ -surplus of the new node can be easily computed. The result applies to tree networks in a straightforward manner, since larger and larger trees can be formed by using this construction repeatedly.

LEMMA 6.1. *Let (U, V, C) be a finite network with no capacity constraints. Suppose there is a partition $\{u_1, \dots, u_d\}, U_1, U_2, \dots, U_d$ of U and a partition $\{v_0, V_1, V_2, \dots, V_d\}$ of V such that $N(u_i) = \{v_0, v_i\}$, where $v_i \in V_i$ for all i , and $N(u) \subset V_i$ for all $u \in U_i$ and all i . Let τ be a real number. Let y denote the τ -surplus of v_0 relative to (U, V) and, for $1 \leq i \leq d$, let y_i denote the τ -surplus of the node v_i relative to the subnetwork (U_i, V_i) . For brevity, let $m_i = m_{u_i}$. Then*

$$(6.3) \quad y = b_{v_0} + \sum_{i=1}^d [y_i + m_i]_0^{m_i} - \tau,$$

where we use the notation $[x]_a^b$ for the number in $[a, b]$ closest to x .

PROOF. Construct an assignment vector f for (U, V, N) meeting the demand and balanced with baseload \bar{b} defined by $\bar{b}_{v_0} = b_{v_0} - y$ and $\bar{b}_v = b_v$ for $v \neq v_0$, as follows. Set $f_{u_i, v_i} = [-y_i]_0^{m_i}$ and $f_{u_i, v_0} = m_i - [-y_i]_0^{m_i} = [y + m_i]_0^{m_i}$ for $1 \leq i \leq d$. Choose $f_{u, v}$ for $u \in U_i, v \in V_i$ for each i so that the demand is

met for each $u \in U_i$ and so that the load is balanced in V_i for each i (with unspecified boundary conditions). We claim that f is a balanced assignment for (U, V, N) . By construction, f is an admissible assignment vector that meets the demand m . Since the load is balanced within V_i for each i , it remains to check the balance conditions for u_1, \dots, u_d . The load at node v_0 for assignment f and baseload \bar{b} is $b_{v_0} - y + \sum_{i=1}^d f_{u_i, v_0}$, which by (6.3) is equal to τ . For each i one of the following three applies: (1) if $y_i < 0$, then $f_{u_i, v_i} = 0$ and $x(v_i) > \tau$; (2) if $y_i < -m_i$, then $f_{u_i, v_i} = m_i$ and $x(v_i) < \tau$; or (3) if $-m_i \leq y_i \leq 0$, then $x(v_i) = \tau$. In any case, the balance condition holds for u_i . Thus, f balances with baseload \bar{b} and demand m as claimed, and the load at v_0 is τ . The proof of the lemma is complete. \square

For the remainder of this section, suppose (U, V, C) is a tree network with degree $d + 1$ where $d \geq 1$. Suppose the baseload is 0 and that demand on each edge has distribution function F_m and the demands on different edges are mutually independent.

A generalized distribution function F corresponds to a probability measure on the interval $[-\infty, +\infty]$. Such a function is assumed to be nondecreasing and right continuous, and the limits $\lim_{x \rightarrow -\infty} F(x)$ and $\lim_{x \rightarrow +\infty} [1 - F(x)]$ must be nonnegative—the values of the limits are the probability assigned to $-\infty$ and $+\infty$, respectively. Write $F < G$ if $F(x) \geq G(x)$ for all x . We write $I_{\{x \leq -\infty\}}$ (resp. $I_{\{x \leq +\infty\}}$) to denote the minimal (resp. maximal) generalized distribution function.

Given $\tau \geq 0$ and an integer $d \geq 1$, let $\Gamma_{d, \tau}$ denote a mapping from the space of generalized distribution functions to the space of ordinary distribution functions defined as follows. Given a generalized distribution function F , let $\Gamma_{d, \tau} F$ denote the distribution function of Y , where

$$(6.4) \quad Y = \sum_{i=1}^d [Y_i + m_i]_0^{m_i} - \tau,$$

$m_1, \dots, m_d, Y_1, \dots, Y_d$ are mutually independent, and, for each i , m_i has distribution function F_m and Y_i has distribution function F . Since Y in (6.4) is a nondecreasing, continuous function of Y_i for each i , it follows that the mapping $\Gamma_{d, \tau}$ is monotone, in the sense that if $F < G$, then $\Gamma_{d, \tau} F < \Gamma_{d, \tau} G$. Also $\Gamma_{d, \tau}$ is a continuous mapping if the space of generalized distribution functions is endowed with the usual weak topology. Thus, there is a unique minimal solution and a unique maximal solution to the fixed-point equation $F = \Gamma_{d, \tau} F$.

Fix a node $v_0 \in V$ and let V_n be the set of nodes at distance less than or equal to n from V_0 . By Lemma 6.1 and induction on n , it follows that $Y_{\tau, n}^{+\infty}$ has distribution function $\Gamma_{d+1, \tau} \Gamma_{d, \tau}^{n*} I_{\{x \leq +\infty\}}$, where $\Gamma_{d, \tau}^{n*}$ denotes the n -fold composition of $\Gamma_{d, \tau}$ with itself. The function $I_{\{x \leq +\infty\}}$ appears here since the $+\infty$ -boundary condition on V_n corresponds to τ -surplus $+\infty$ for all nodes in V_n^c . Let n converge to ∞ to obtain the following proposition.

PROPOSITION 6.2. For given τ , $Y_\tau^{+\infty}$ has distribution $\Gamma_{d+1,\tau}F^{+\infty}$, where $F^{+\infty}$ is the maximal solution to $\Gamma_{d,\tau}F = F$. Similarly, $Y_\tau^{-\infty}$ has distribution $\Gamma_{d+1,\tau}F^{-\infty}$, where $F^{-\infty}$ is the minimal solution to $\Gamma_{d,\tau}F = F$.

Trees with exponentially distributed loads. Suppose that the distribution of the demand m_u for each edge u in the tree has the exponential distribution with mean 1: $F_m(c) = 1 - e^{-c}$ for $c \geq 0$. The following relation is helpful in determining the set of solutions F to the fixed-point equation $F = \Gamma_{d,\tau}F$. If m is exponentially distributed with mean 1 and Y is independent of m and has an arbitrary distribution, then, for any $c > 0$,

$$\begin{aligned} P[[Y + m]_0^m \geq c] &= P[[Y + m]_0^m \geq c | Y \geq 0]P[Y \geq 0] \\ &\quad + P[[Y + m]_0^m \geq c | 0 > Y \geq -m]P[0 > Y \geq -m] \\ &= P[m \geq c | Y \geq 0]P[Y \geq 0] \\ &\quad + P[Y + m \geq c | 0 > Y \geq -m]P[0 > Y \geq -m] \\ &= e^{-c}P[Y \geq 0 \text{ or } 0 > Y \geq -m] \\ &= e^{-c}P[Y + m \geq 0]. \end{aligned}$$

Thus, the distribution of $[Y + m]_0^m$ depends on the distribution of Y only through the parameter $q = P[Y + m \geq 0]$. If F_n is a sequence of distribution functions such that $F_{n+1} = \Gamma_{d,\tau}F_n$ and if q^n is the parameter q associated with F_n , then $q^{n+1} = \gamma_{d,\tau}(q^n)$, where

$$\gamma_{d,\tau}(q) = P[m + Z_1 + \dots + Z_d - r \geq 0],$$

Z_1, \dots, Z_d are random variables with $P[Z_i \geq c] = qe^{-c}$ for $c > 0$ and $P[Z_i = 0] = 1 - q$, m is an exponentially distributed random variable with mean 1 and m, Z_1, \dots, Z_d are mutually independent. If U denotes the number of i such that $Z_i > 0$, then U has the binomial distribution with parameters d and q . The probability that the sum of $u + 1$ exponentially distributed random variables is greater than or equal to τ is equal to $P[\text{Poi}(\tau) \leq u]$, where $\text{Poi}(\tau)$ represents a Poisson random variable with mean τ . Therefore,

$$\gamma_{d,\tau}(q) = P[\text{Poi}(\tau) \leq \text{Bi}(d, q)],$$

where $\text{Bi}(d, q)$ represents a Poisson random variable with parameters d and q which is independent of $\text{Poi}(\tau)$.

Let $q^{+\infty}$ (resp., $q^{-\infty}$) denote the maximal (resp., minimal) fixed point of $\gamma_{d,\tau}$. Then $q^{+\infty}$ and $q^{-\infty}$ are the parameters associated with $F^{+\infty}$ and $F^{-\infty}$, the maximal and minimal solutions to $F = \Gamma_{d,\tau}F$. Therefore, $Y_\tau^{+\infty}$ has the same distribution as $Z_1 + Z_2 + \dots + Z_{d+1} - \tau$, where Z_1, \dots, Z_{d+1} are independent random variables with $P[Z_i \geq c] = q^{+\infty}e^{-c}$ for $c > 0$ and $P[Z_i = 0] = 1 - q^{+\infty}$, so that

$$P[Y_\tau^{+\infty} < 0] = 1 - P[Z_1 + Z_2 + \dots + Z_{d+1} \geq \tau] = 1 - \bar{\gamma}_{d+1,\tau}(q^{+\infty}),$$

where

$$\bar{\gamma}_{d+1,\tau}(q) = P[\text{Poi}(\tau) < \text{Bi}(d + 1, q)].$$

Combining this and a similar equation for $Y_\tau^{-\infty}$ with Proposition 6.1 yields the following corollary to Proposition 6.2.

COROLLARY 6.2. *For a tree network with independent, exponentially distributed demands of mean 1 on the edges,*

$$(6.5) \quad P[x^{+\infty}(v_0) < \tau] = 1 - \bar{\gamma}_{d+1, \tau}(q^{+\infty}),$$

$$(6.6) \quad P[x^{-\infty}(v_0) \leq \tau] = 1 - \bar{\gamma}_{d+1, \tau}(q^{-\infty}).$$

Thus, the distributions of $x^{+\infty}(v_0)$ and $x^{-\infty}(v_0)$ are different if and only if the minimal and maximal fixed points of $\gamma_{d, \tau}$ are distinct for τ in some interval of nonzero length.

Fixed points of the function $\gamma_{d, \tau}$ are easily computed numerically for small d and the qualitative behavior of the fixed points is easy to determine for large d . The distribution functions are pictured in Figure 7 for $1 \leq d \leq 8$, and they shift to the right as d increases. For each integer d in the range $1 \leq d \leq 5$, the upper and lower distributions are identical, and for $d \geq 6$ the upper and lower distributions are different in some interval. A proof of this fact by elementary numerical analysis is straightforward but tedious, and is




FIG. 7. *Distribution functions of load for $1 \leq d \leq 8$. For $d = 6, 7, 8$ the distribution functions for both the minimal and maximal loads are shown, with dashed lines indicating where they diverge from each other.*

omitted. The function $\gamma_{\tau,d}$ is pictured in Figures 8 and 9 for the two “borderline” cases, $d = 5$ and $d = 6$, respectively, for selected values of τ . The derivative of the distribution function for $d = 5$ is less than 35 throughout the interval $[0, 1]$. The function $\gamma_{d,\tau}(q)$ is decreasing in τ for each value of d and q fixed. There are either one, two or three fixed points of $\gamma_{d,\tau}$ for d, τ fixed.

Trees with Bernoulli loads. Suppose the demand m_u is 1 with probability p and 0 with probability $1 - p$ for all u , and that demands for distinct consumers are independent.

COROLLARY 6.3. *If $pd \leq 1$, then, with probability 1 there is a unique balanced load vector and load percolation does not occur at any node. If $pd > 1$ then, for any fixed node v_0 , the load at v_0 is not unique with positive probability.*

PROOF. If $pd \leq 1$, then the components of the graph induced by edges u with $m_u = 1$ are all finite, with probability 1, as is well known from the theory of branching processes. Hence, the balanced load vector is unique and load percolation does not occur, with probability 1. The first statement of the proposition is proved.




FIG. 8. Graph of $\gamma_{\tau,d}$ for $d = 5$ and $\tau = 3.0, 3.05, 3.10, 3.15, 3.20$ and 3.25 .




FIG. 9. Graph of $\gamma_{\tau,d}$ for $d = 6$ and $\tau = 3.4, 3.5, 3.6$ and 3.7 .

It is shown next that $P[x^{-\infty}(v_0) > 1] = 0$ for any node v_0 and any value of p with $0 \leq p \leq 1$. If, in (6.4), $\tau = 1, Y_i \leq -1$ and $m_i \in \{0, 1\}$ for $1 \leq i \leq d$, then $Y = -1$. Thus, the minimal solution to the equation $F = \Gamma_{d,\tau}F$ for $\tau = 1$ is $F^{-\infty}(x) = I_{\{x \geq -\tau\}}$. Moreover, $F^{-\infty} = \Gamma_{d+1,\tau}F^{-\infty}$. Therefore, by (6.2) and Proposition 6.2, $P[x^{-\infty}(v_0) \leq 1] = F^{-\infty}(0) = 1$, as was to be proved. (Another proof of this result is the following. It was proved for $p = 1$ in Section 1, and it can thus be proved for all p by exploiting the fact that $x^{-\infty}$ is stochastically increasing in p .)

To complete the proof of the proposition, we show that $P[x^{+\infty}(v_0) > 1] > 0$ if $pd > 1$. Suppose that $\tau = (2t + 1)/2t$ for some integer t , and $d \geq 2$. Define a new mapping $\tilde{\Gamma}_{d,\tau}$ as follows. Given a distribution function F , let $\tilde{\Gamma}_{d,\tau}F$ denote the distribution function of \tilde{Y} , where

$$\tilde{Y} = \begin{cases} 0, & \text{if } [Y_i - m_i]_0^{m_i} \geq \frac{t+1}{2t} \text{ for exactly two values of } i, \\ -\frac{k+1}{2t}, & \text{if } [Y_i + m_i]_0^{m_i} = 1 - \frac{k}{2t} \text{ for some } i \text{ and some } k \in \{0, \dots, t-2\} \\ & \text{and } [Y_j + m_j]_0^{m_i} \leq \frac{t}{2t} \text{ for } j \neq i, \\ -\tau, & \text{otherwise,} \end{cases}$$

and where $m_1, \dots, m_d, Y_1, \dots, Y_d$ are the same as in the definition of $\Gamma_{d,\tau}F$. Note that $\tilde{Y} \leq Y$ for any value of $m_1, \dots, m_d, Y_1, \dots, Y_d$ and that \tilde{Y} , like Y , is a nondecreasing function of Y_1, \dots, Y_d . Hence, if $F < G$, then $\tilde{\Gamma}_{d,\tau}F < \tilde{\Gamma}_{d,\tau}G$. It follows that $\tilde{\Gamma}_{d,\tau}$ has a maximal fixed point, $\tilde{F}_{d,\tau}^{+\infty}$, and that $\tilde{F}_{d,\tau}^{+\infty}$ is stochastically smaller than the maximal fixed point of $\Gamma_{d,\tau}$. It will be shown that $\tilde{F}_{d,\tau}^{+\infty}(0) > 0$ if $pd > 1$ so that indeed $P[x^{+\infty}(v_0) > 1] > 0$ if $pd > 1$.

The probability mass corresponding to the distribution function $\tilde{F}_{d,\tau}^{+\infty}$ is concentrated on the finite set $\{0, -1/2t, \dots, -(t-1)/2t\} \cup \{-\tau\}$. Setting π_j equal to the probability mass assigned by $\tilde{F}_{d,\tau}^{+\infty}$ to $-j/2t$, the fixed-point equation of $\tilde{\Gamma}_{d,\tau}$ becomes

$$(6.7) \quad \pi_0 = \binom{d}{2} (ps)^2 (1-ps)^{d-2},$$

$$(6.8) \quad \pi_{j+1} = \pi_j [dp(1-ps)^{d-1}], \quad 0 \leq j \leq t-2,$$

$$(6.9) \quad s = \pi_0 + \dots + \pi_{t-1}.$$

Solve for π_1, \dots, π_{t-1} in terms of π_0 . Assume that $dp(1-ps)^{d-1} \neq 1$ and sum a partial geometric series to obtain

$$s = \frac{1 - [dp(1-ps)^{d-1}]^t}{1 - [dp(1-ps)^{d-1}]} \pi_0.$$

Using (6.7) to substitute in for π_0 and canceling a factor of s , we obtain the following equation for s :

$$(6.10) \quad sp^2 \binom{d}{2} (1-ps)^{d-2} \{1 - [dp(1-ps)^{d-1}]^t\} = 1 - dp(1-ps)^{d-1}.$$

Let s_0 denote the unique value in the interval $(0, 1)$ satisfying $dp(1-ps_0)^{d-1} = 1$. Such s_0 exists assuming that $dp > 1$. Then $s = s_0$ is a solution to (6.10). However, in deriving (6.10) we assumed that $dp(1-ps)^{d-1} \neq 1$, so the value s_0 does not necessarily represent a solution to the original equations (6.7)–(6.9). Let D_{left} (resp. D_{right}) denote the derivative of the quantity on the left-hand side (resp. right-hand side) of (6.10), evaluated at s_0 . The left-hand side of (6.10) is 0 at $s = 0$ and the right-hand side is negative at $s = 0$. Thus, if $D_{\text{left}} > D_{\text{right}}$, then there must exist a solution s to (6.10) in the interval $(0, s_0)$. By straightforward calculation, find that $D_{\text{left}} > D_{\text{right}}$ if and only if

$$(6.11) \quad t > \frac{2(1-s_0p)}{(d-1)s_0p}.$$

Thus, for t satisfying (6.11), the maximal fixed point $F^{+\infty}$ of $\Gamma_{d,\tau}$, where $\tau = (2t+1)/2t$, assigns positive probability to $[0, \infty)$. Hence, $P[x^{+\infty}(v_0) \geq \tau] > 0$ if $dp > 1$ and $\tau = (2t+1)/2t$ where t satisfies (6.11). \square

7. Open problems. A number of questions arise. Is there a way to compute the distribution of the load at node 0 in a d -dimensional lattice network for some $d \geq 2$ and independent demands on edges with some

interesting distribution? Does load percolation occur with positive probability for a d -dimensional rectangular lattice and Bernoulli loads for some $p < 1$? Whether or not load percolation occurs, what is the typical size or shape of the connected components consisting of nodes with equal load for a balanced load vector?

Some questions related to a $d + 1$ -dimensional tree network with no baseload and independent, identically distributed demands on edges are the following. Is it true that for any $d_0 \geq 1$ there is a distribution of demands such that the balanced load vector is unique for $d \leq d_0$? Stephen Turner [6] of Cambridge University conjectures that for any given demand distribution there is nonuniqueness for sufficiently large d . How can one characterize the possible sets of fixed points of $\gamma_{d,\tau}$ as the demand distribution varies? Is there a simple test for when load percolation occurs in tree networks in cases where the balanced load vector is unique, in particular, for exponentially distributed demands with $2 \leq d \leq 5$?

REFERENCES

- [1] ALANYALI, M. and HAJEK, B. (1995). Analysis of simple algorithms for dynamic load balancing. Unpublished manuscript.
- [2] ESARY, J. D., PROSCHAN, F. and WALKUP, D. W. (1967). Association of random variables, with applications. *Ann. Math. Statist.* **38** 1466–1474.
- [3] GIBBENS, R. J., KELLY, F. P. and TURNER, S. R. E. (1993). Dynamic routing in multiparented networks. *IEEE Trans. Networking* **1** 261–270.
- [4] GRIMMETT, G. (1989). *Percolation*. Springer, New York.
- [5] HAJEK, B. (1990). Global load balancing by local adjustment. *IEEE Trans. Inform. Theory* **36** 1398–1414.
- [6] TURNER, S. (1993). Personal communication.

COORDINATED SCIENCE LABORATORY
1308 WEST MAIN STREET
URBANA, ILLINOIS 61801