Adv. Appl. Prob. **34,** 58–84 (2002)
Printed in Northern Ireland

MARKOV RANDOM FIELD MODELS OF MULTICASTING IN TREE NETWORKS

KAVITA RAMANAN * ** AND
ANIRVAN SENGUPTA,* *Bell Laboratories, Lucent Technologies*
ILZE ZIEDINS,*** *University of Auckland*
PARTHA MITRA,* *Bell Laboratories, Lucent Technologies*

Abstract

In this paper, we analyse a model of a regular tree loss network that supports two types of calls: unicast calls that require unit capacity on a single link, and multicast calls that require unit capacity on every link emanating from a node. We study the behaviour of the distribution of calls in the core of a large network that has uniform unicast and multicast arrival rates. At sufficiently high multicast call arrival rates the network exhibits a 'phase transition', leading to unfairness due to spatial variation in the multicast blocking probabilities. We study the dependence of the phase transition on unicast arrival rates, the coordination number of the network, and the parity of the capacity of edges in the network. Numerical results suggest that the nature of phase transitions is qualitatively different when there are odd and even capacities on the links. These phenomena are seen to persist even with the introduction of nonuniform arrival rates and multihop multicast calls into the network. Finally, we also show the inadequacy of approximations such as the Erlang fixed-point approximations when multicasting is present.

Keywords: Multicasting; unicasting; broadcasting; Cayley trees; Bethe lattice; loss networks; blocking probabilities; phase transitions; Erlang fixed-point approximations; reduced load approximations; statistical mechanics; spin models; hard core models; symmetry breaking

AMS 2000 Subject Classification: Primary 60G60; 90B15
Secondary 68M20; 60K35

1. Introduction

In this paper, we present and analyse a model of multicast and unicast calls arriving at an idealized loss network which has the form of a symmetric tree.

Multicasting arises in both queueing and loss networks. Instead of having a simple end-to-end connection, a transmission is made to a group of individuals from a single site. An example of this in the loss network setting is a conference call. In a queueing setting this arises, for instance, when multiple copies of a message on the Internet are broadcast from one person to a number of sites via a mailing list distribution. In recent years there has been a growing interest in multicasting applications that require reliable data delivery such as software distribution and news broadcasts. This has given rise to several interesting problems. One major problem is the issue of how to construct connections between individuals in a way that makes communication

Received 4 December 2000; revision received 2 January 2002.
* Postal address: Bell Laboratories, Lucent Technologies, 600 Mountain Avenue, Murray Hill, NJ 07974, USA.
** Email address: kavita@lucent.com
*** Postal address: Department of Statistics, University of Auckland, Private Bag 92019, Auckland, New Zealand.

efficient and reliable and minimizes demands on network resources. Much effort has been expended on this [1], [6], [16]. Before this can be answered, however, we need to be able to analyse network performance in the presence of multicast calls, a task that is considerably more complex than the performance analysis of a purely unicast network. In this paper we study performance in the context of a symmetric loss network in the form of a tree. This is a very simple network which nevertheless exhibits interesting behaviour, and raises questions about the performance analysis of more complicated networks using traditional techniques and approximations. We expect that the insight gleaned here will be useful in more general contexts, especially in view of the fact that many of the qualitative properties observed in the case of a symmetric network persist even when asymmetry is introduced.

A general loss network without controls can be described as follows. Let J be the finite collection of resources (or links) in the network, and $\boldsymbol{C} = \{C_j,\ j \in J\}$ where C_j is the capacity of resource j. Let R denote the finite set of possible call (or customer) types in the network. Calls of each type $r \in R$ arrive as a Poisson process of rate ν_r and have identically distributed holding or service times that we assume, without loss of generality, to have mean 1 (see [5]). Each call of type r requires capacity A_{jr} from resource $j \in J$ for the duration of its holding time. If one or more of the resources (links) j for which $A_{jr} > 0$ does not have sufficient free capacity to carry the call of type r, then the call is *blocked* and considered lost. Otherwise the call is accepted. All arrival processes and holding times are assumed to be independent of one another. One of the standard measures of performance for a loss system is the probability that a call is blocked or lost (the loss, or blocking probability). Loss networks have been widely studied, as models of circuit-switched and ATM networks, and in other contexts. Excellent introductions to this general model are found in [13] and [17].

We confine our attention in this paper to a special kind of symmetric loss network with a regular tree structure. Let T_∞ be an infinite tree with $q + 1$ edges emanating from each node. Let S_L be the finite spherical subtree of radius L which consists of a single central node and all nodes no more than a distance L from it, where the distance between two nodes is the minimum number of edges in any path connecting the two nodes. There are $q + 1$ edges emanating from each node in S_L except the terminal nodes, which are attached to just one edge. We assume that each edge has the same capacity, C. We let $\mathcal{E}(T)$ denote the set of edges in a given tree T and use $\mathcal{E}(t)$ to denote the set of edges that are incident with the node t. For each node $t \in T$, we define $\mathcal{N}(t) := \{t' \in T : tt' \in \mathcal{E}(t)\}$ to be the set of immediate neighbours of t. We allow two types of calls in the network. Single link calls, requiring capacity on just a single edge, arrive at each edge as a Poisson process of rate λ. Single link calls connect two neighbouring nodes. Multicast calls, on the other hand, are centred at a node, and connect that node to *all* of its immediate neighbours. Thus a call centred at an internal node $t \in S_L$ requires capacity on each of the $q + 1$ edges in $\mathcal{N}(t)$, while a call centred at a terminal node requires capacity on just one edge. We assume that multicast calls arrive at each node t as a Poisson process of rate $\nu_t = \nu$. We will assume that all calls require just a single unit of capacity on each edge of their connection. Throughout this paper we will be concentrating on assessing the blocking probabilities for large spherical tree networks. Much of the treatment below can be applied in the more complex multirate setting, but in order to clarify the presentation we do not do so here.

Although the tree network that we have described here is clearly unrealistic, it is nevertheless a useful tool for studying some of the behaviours that multicasting might introduce to a loss network. We show below that multicasting can introduce unfairness in the symmetric spherical tree that we study, as the blocking probabilities seen by the same call types at different nodes

may vary considerably. These differences are a consequence of a phase transition effect that arises when the arrival rate for multicast calls is sufficiently high and that of unicast calls is sufficiently low. Thus we observe that the presence of unicast calls assists in maintaining fairness and homogeneity of the blocking probabilities in the system. In addition, we study the trade-off between unicast and multicast call arrival rates for a given fixed unicast blocking probability. Finally, we also shed some light on approximations of networks with multicasting. A commonly made assumption when calculating approximate blocking probabilities is that resources (edges or links) behave as though they are blocking independently of one another. This assumption underpins approximations such as the Erlang fixed-point approximation, also known as the reduced load approximation, and refinements of it—approximations that have been widely and very successfully used, even where it is clear that the independence assumption does not hold. There is now considerable theoretical justification for the use of these approximations (see e.g. [13] for details). However, it seems that when applied to multicast networks, although the approximation may still be very good for low values of the multicast arrival rate, it may considerably underestimate the point at which phase transition occurs, and cannot be applied directly beyond that point. Karvo *et al.* [10] have also applied the reduced load approximation to a multicast network, and after comparing it with simulations comment that it performs worse at higher values of the arrival rate.

The special structure of the tree network greatly simplifies the study of these questions, since the removal of any single edge from the tree splits it into two disjoint subgraphs. This in turn means that we can very easily develop recursions that give us the exact blocking probabilities. The recursions that we obtain in this paper for the blocking probabilities are related to those given in [22] and are given in the form developed in the statistical mechanics literature (see e.g. [2]).

Kelly [12], [13] previously considered the purely multicast tree network model for $C = 1$ and obtained a recursion which gives the probability of a call being accepted. An earlier paper by Spitzer [18] also gives results for tree networks that apply to the $C = 1$ case. Zachary [20], [21] analyses the more general tree network, which includes as a special case the model considered here with $C \in \mathbb{N} \cup \{\infty\}$, but $\lambda = 0$ (i.e. the pure multicast network with no unicast calls). In particular, given a homogeneous Markov specification on the Cayley tree, in [20, Theorem 4.1] Zachary establishes a one–one correspondence between Markov chains associated with that specification and solutions to an associated recursion problem, which is related to recursions considered in this paper. Georgii [8] serves as a good introduction and overview of some of these earlier results. More recently, van den Berg and Steif [19] obtain results for general bipartite graphs, but again they apply only to the equivalent of the $C = 1$ and $\lambda = 0$ case in our model. Louth [14] has studied a related model on a square lattice, and van den Berg and Steif also applied their results to the square lattice (again under a condition equivalent to the requirement that $C = 1$ and $\lambda = 0$). Brightwell and Winkler [3], [4] have applied the methods of graph theory and combinatorics to the tree network with the equivalent of $C \geq 1$ and $\lambda = 0$. Observe that in all cases, the conditions placed on these models previously considered in the literature are equivalent to assuming that there is no unicast traffic in the network.

In Section 2, we introduce the basic model and provide an explicit expression for the stationary distribution on a finite tree. In Section 2.1, we derive a recursion to calculate the normalization constant, and in Section 2.2, we provide expressions for the unicast and multicast blocking probabilities. In Section 3, we study these quantities for the collection of nodes located deep within a large network. In Sections 3.1 and 3.2, we calculate the stationary distribution and blocking probabilities for these core nodes. In Section 3.3, we specialise to the case $C = 1$ and study the trade-offs between unicast and multicast arrival rates. In Section 3.4,

we consider the case $C = 2$, and comment on the general C case in Section 3.5. We analyse the Erlang fixed-point approximation for this network in Section 4 and show that it provides an inadequate description of the network when multicasting is present. In Section 5, we show that phase transitions persist even for generalizations of the model that allow multihop multicast calls and heterogeneous arrival rates. In contrast, as discussed in the conclusions in Section 6, it is conjectured that heterogeneity destroys phase transitions for integer lattice networks [19].

2. Model description

As explained in the introduction, we consider a loss network that has the structure of a spherical Cayley tree (see, for example, Figure 2 below). The radius of such a tree is defined to be the distance (or number of edges) between the central node and any terminal node. The nodes in a tree can be divided into even and odd lattices in the following fashion. Designate an arbitrary node as being 'even'. Then all nodes at an even distance from it belong to the even lattice, and those at an odd distance from it belong to the odd lattice. For a finite tree, the reference node is often taken to be one of the terminal nodes.

The resources in the network T are its edges $\mathcal{E}(T)$, each of which has capacity C. Note that the capacity C may also be even or odd. The total number of types of calls in any Cayley tree T is given by $M_T := |T| + |\mathcal{E}(T)|$, comprising multicast calls that arrive at each node t with Poisson rate ν (and require unit capacity from each edge incident to that node), and unicast calls that arrive at an edge with Poisson rate λ (and require unit capacity from that edge for the period of its holding time). Let $\mathcal{I} := \{0, 1, \ldots, C\}$. The set of feasible configurations on T (subject to the blocking constraint) is given by

$$\Omega_T := \{n \in \mathcal{I}^{M_T} : n_{t'} + n_t + n_{tt'} \leq C \text{ for every } t', t \in T, t' \in \mathcal{N}(t)\}. \tag{2.1}$$

The dynamics of the model can then be described by a Markov process $\boldsymbol{n}(s)$, where for each $s \in [0, \infty)$, $\boldsymbol{n}(s) := \{\boldsymbol{n}_t(s), \boldsymbol{n}_e(s),\ t \in T, e \in \mathcal{E}(T)\} \in \Omega_T$ and $\boldsymbol{n}_t(s)$ and $\boldsymbol{n}_e(s)$ are, respectively the random variables whose distributions describe the number of multicast and unicast calls in progress at node t and edge e at time s. Under the assumptions stated above, it is well known [13] that the process $\boldsymbol{n}(s)$ has a stationary distribution π, which has the truncated product form

$$\pi(\boldsymbol{n}) = \frac{1}{\mathcal{Z}_\pi} \prod_{t \in T} \frac{\nu^{n_t}}{n_t!} \prod_{e \in \mathcal{E}(T)} \frac{\lambda^{n_e}}{n_e!}, \qquad \boldsymbol{n} \in \Omega_T, \tag{2.2}$$

where $\mathcal{Z}_\pi$ is the normalizing constant

$$\mathcal{Z}_\pi := \sum_{n \in \Omega_T} \prod_{t \in T} \frac{\nu^{n_t}}{n_t!} \prod_{e \in \mathcal{E}(T)} \frac{\lambda_{n_e}}{n_e!}. \tag{2.3}$$

2.1. Recursion relation

As discussed above, we have an explicit expression (2.2) for the stationary distribution of interest. However, this expression is not always very useful because the calculation of the normalization constant in (2.3) often becomes infeasible for networks of realistic size. Note that the stationary distribution π defines a Markov random field on T (see [22] for a discussion on the connection between loss networks and Markov random fields). Therefore, not surprisingly, analogous problems of calculating normalization constants of the form (2.3) often arise in statistical mechanics in the calculation of Gibbs measures for Markov random fields. From

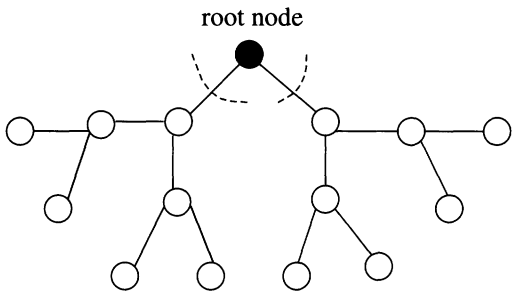

FIGURE 1: A rooted (2,3) Cayley tree.

the statistical mechanics perspective, the network model described above (with $\lambda = 0$) can be interpreted as a 'spin model' on the Cayley tree with hard constraints [2], [15], and is analogous to the hard-core models studied in [19]. Given a tree T, and $\boldsymbol{n} \in \Omega_T$, the weight of the configuration $\boldsymbol{n}$ for this model takes the form

$$W(\boldsymbol{n}) := \prod_{t \in T} g(n_t) \prod_{tt' \in \mathcal{E}(T)} h(n_t, n_{t'}, n_{tt'}),$$

where

$$g(i) := \frac{\nu^i}{i!}, \qquad h(i, j, k) := \frac{\lambda^k}{k!} \chi(C - (i + j + k)),$$

and

$$\chi(x) := \begin{cases} 0 & \text{if } x < 0, \\ 1 & \text{otherwise.} \end{cases}$$

The weight is proportional to the probability of a configuration. The normalization constant defined in (2.3) can be expressed in terms of the weight as $Z_\pi = \sum_{\boldsymbol{n} \in \Omega_T} W(\boldsymbol{n})$, and is referred to in the physics literature as the partition function.

In order to simplify the calculation of the stationary distribution, we introduce another subgraph T_m, which we refer to as a rooted (q, m) Cayley tree. The graph T_m is a tree of size m which comprises a distinguished node or root O, internal nodes, and terminal nodes that are a distance m from the root O. The distinguished node has q edges emanating from it, the terminal nodes have just one, and all other nodes are incident to $q + 1$ edges (see Figure 1). (We will also use T_m to denote the set of nodes of this finite subgraph.) In order to derive a recursion relation which characterizes the stationary distribution, we exploit the property that the removal of any internal edge splits the tree T_m into two disjoint trees. Let $\Omega_m := \Omega_{T_m}$ denote the set of feasible configurations on the tree T_m as defined in (2.1). For $i \in \mathcal{I}$, we define $Z_m(i)$ to be the weighted sum of all feasible configurations in T_m that have i multicast calls at the root node O. More precisely,

$$Z_m(i) := \sum_{n \in \Omega_m : n_0 = i} \prod_{t \in T_m} \frac{\nu^{n_t}}{n_t!} \prod_{e \in \mathcal{E}(T_m)} \frac{\lambda^{n_e}}{n_e!}.$$

Then the configuration vector $(Z_{m+1}(i),\ i \in \mathcal{I})$ for the tree T_{m+1} can be expressed in terms of that for T_m as follows:

$$Z_{m+1}(i) = \frac{\nu^i}{i!}\left(\sum_{j=0}^{C-i}\frac{\lambda^j}{j!}\sum_{k=0}^{C-i-j} Z_m(k)\right)^q. \tag{2.4}$$

Since a spherical tree of radius L can be decomposed into a central node with $q+1$ edges, and $q+1$ rooted trees each of size $L-1$, the normalization constant (or partition function) Z_π for the tree S_L is given by

$$Z_\pi = \sum_{i=0}^{C} Z_\pi(i), \tag{2.5}$$

where

$$Z_\pi(i) = \frac{\nu^i}{i!}\left(\sum_{j=0}^{C-i}\frac{\lambda^j}{j!}\sum_{k=0}^{C-i-j} Z_{L-1}(k)\right)^{q+1}. \tag{2.6}$$

Thus for a finite spherical tree S_L, given boundary conditions for the external nodes, the partition function Z_π can be calculated from the iteration (2.4) and expressions (2.5) and (2.6). Free boundary conditions correspond to setting

$$Z_0(i) = \frac{\nu^i}{i!},$$

for $i = 0, \ldots, C$. The problem of computing the stationary distribution π on a finite tree T is thus reduced to that of analysing the recursion relation (2.4).

2.2. Blocking probabilities

In this section we derive expressions for the stationary blocking probabilities in S_L, the spherical tree of radius L.

We begin with the stationary unicast blocking probabilities. Let s be the central node, s' a neighbouring node and α_e^L be the probability that a unicast call arriving at the internal edge $e = ss'$ is blocked. Removal of the internal edge ss' partitions the tree into two rooted trees, one of size L, with root s, and another of size $L-1$, with root s' (see Figure 2). Then the weight of a configuration with i multicast calls at node s, j calls at node s' and k unicast calls on ss' is

$$\frac{\lambda^k}{k!} Z_L(i) Z_{L-1}(j),$$

provided $i + j + k \le C$. The probability of call acceptance on the edge ss' is the weighted sum of configurations that have one unit of free capacity on the edge ss' divided by the weighted sum of all configurations. Therefore,

$$1 - \alpha_e^L = \frac{\sum_{k=0}^{C-1}(\lambda^k/k!)\sum_{i=0}^{C-1-k}\sum_{j=0}^{C-1-k-i} Z_L(i)Z_{L-1}(j)}{\sum_{k=0}^{C}(\lambda^k/k!)\sum_{i=0}^{C-k}\sum_{j=0}^{C-k-i} Z_L(i)Z_{L-1}(j)}. \tag{2.7}$$

Now consider the stationary multicast blocking probability, β_s^L say, for a call arriving at the central node s. Since the node s is central, the spherical Cayley tree can be partitioned into the central node s with $q+1$ edges emanating from it, and $q+1$ rooted Cayley trees of size $L-1$,

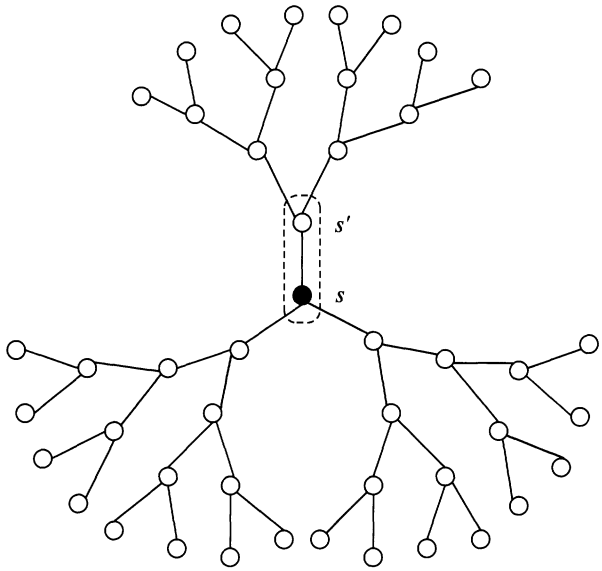

FIGURE 2: A spherical Cayley tree with radius 4 ($q = 2$).

each having as its root one of the nodes adjacent to the node s. Note that if there are i multicast calls at node s, then there can be at most $C - i$ unicast calls on each edge emanating from it; and if there are then j unicast calls on the edge ss', then there can be at most $C - i - j$ multicast calls at node s'. Hence, the weight of a configuration with i multicast calls at node s is

$$\frac{\nu^i}{i!}\left(\sum_{j=0}^{C-i}\frac{\lambda^j}{j!}\sum_{k=0}^{C-i-j} Z_{L-1}(k)\right)^{q+1}.$$

Consequently,

$$1-\beta_s^L = \frac{\sum_{i=0}^{C-1}(\nu^i/i!)(\sum_{j=0}^{C-i-1}(\lambda^j/j!)\sum_{k=0}^{C-i-j-1} Z_{L-1}(k))^{q+1}}{\sum_{i=0}^{C}(\nu^i/i!)(\sum_{j=0}^{C-i}(\lambda^j/j!)\sum_{k=0}^{C-i-j} Z_{L-1}(k))^{q+1}}. \tag{2.8}$$

Related expressions for other networks can be found in [22].

3. Large networks

In this section, we examine call distributions and blocking probabilities for the collection of nodes that are located deep within the core of a large network. This collection of deep nodes in a large spherical Cayley tree is often referred to as the Bethe lattice. In Sections 3.1 and 3.2, we study the limit of the recursion and blocking probabilities described in Sections 2.1 and 2.2 respectively. In Sections 3.3 and 3.4, we consider the special cases of $C = 1$ and $C = 2$, where some analytical results can be obtained. We present numerical results for higher C in Section 3.5.

3.1. Limit of the finite tree stationary distributions

In Section 2.1, we showed that given a boundary condition Z_0, we can use the recursion (2.4) to calculate the stationary distribution π on any spherical network. Insight into the stationary distribution for large networks can be obtained by analysing properties of the recursion (2.4) in the asymptotic limit as $m \to \infty$.

For $i = 1, \ldots, C$, and $m \geq 0$, define

$$\xi_m(i) := \frac{Z_m(i)}{Z_m(0)} \quad \text{and} \quad \xi_m := (\xi_m(1), \ldots, \xi_m(C)).$$

(Note that ξ_m is well defined since $Z_m(0)$ is always strictly positive for any m.) We can then express the recursion (2.4) in the more convenient form

$$\xi_{m+1} = \Phi^{\nu,\lambda}(\xi_m), \tag{3.1}$$

where for $\xi \in \mathbb{R}_+^C$, the two parameter family of mappings $\Phi^{\nu,\lambda} : \mathbb{R}_+^C \to \mathbb{R}_+^C$, $\nu, \lambda > 0$, is defined by $\Phi^{\nu,\lambda}(\xi) = (\Phi_1^{\nu,\lambda}(\xi), \Phi_2^{\nu,\lambda}(\xi), \ldots, \Phi_C^{\nu,\lambda}(\xi))$, such that for $k = 1, \ldots, C$,

$$\Phi_k^{\nu,\lambda}(\xi) := \frac{\nu^k}{k!}\left(\frac{\sum_{i=0}^{C-k}(\lambda^i/i!)[1 + \sum_{j=1}^{C-k-i}\xi(j)]}{\sum_{i=0}^{C}(\lambda^i/i!)[1 + \sum_{j=1}^{C-i}\xi(j)]}\right)^q. \tag{3.2}$$

We will refer to the mappings $\Phi^{\nu,\lambda}$ as random field maps, to differentiate them from the Erlang fixed point maps f introduced in Section 4. For the case $\lambda = 0$, related maps associated with a general homogeneous Markov specification can be found in [20, p. 899], [21, (2.11)] and [22].

Suppose $\mathbb{R}^C$ is endowed with the usual Euclidean metric and define the compact subset $K_C^\nu \subset \mathbb{R}^C$ to be the Cartesian product of the sets $[0, \nu^i/i!]$, $i = 1, \ldots, C$. Observe that $\Phi^{\nu,\lambda}(\mathbb{R}_+^C) \subset (K_C^\nu)$ and the functions $\Phi_k^{\nu,\lambda}$ have continuous partial derivatives with respect to $\xi(j)$, $j = 1, \ldots, C$. This shows that there exists a continuous function $g(\nu, \lambda) : \mathbb{R}^2 \to \mathbb{R}^2$ that is nondecreasing in ν with $g(0, 0) = 0$, which satisfies

$$\|\Phi^{\nu,\lambda}(\xi) - \Phi^{\nu,\lambda}(\eta)\| \leq g(\nu, \lambda)\|\xi - \eta\|.$$

Let $S \subset \mathbb{R}^2$ be the open set defined by $\{x \in \mathbb{R}^2 : g(x) < 1\}$. This implies the following theorem, which is related to results of Zachary [20, Theorem 4.2], who showed that there exists at least one fixed point to the related map in [20], which is equivalent to the case when $\lambda = 0$ here.

Theorem 3.1. *Let $\Phi^{\nu,\lambda} : \mathbb{R}_+^C \to \mathbb{R}_+^C$, $\nu, \lambda > 0$, be the family of mappings defined in (3.2), and let $S \subset \mathbb{R}^2$ be the associated set defined above. Then for any $(\nu, \lambda) \in S$, the equation $\Phi^{\nu,\lambda}(\xi) = \xi$ has a unique solution $x^{\nu,\lambda} \in K_C^\nu$. Moreover, given any initial condition $\xi_0 \in \mathbb{R}^C$, the sequence defined iteratively by $\xi_n := \Phi^{\nu,\lambda}(\xi_{n-1})$ converges to $x^{\nu,\lambda}$ and*

$$\|\xi_{n+1} - x^{\nu,\lambda}\| \leq \frac{g(\nu, \lambda)}{1 - g(\nu, \lambda)}\|\xi_n - \xi_{n+1}\|.$$

Proof. The metric space $\mathbb{R}_+^C$ is complete, and for every $\nu > 0$, K_C^ν is compact. Moreover, the fact that the mappings $\Phi^{\nu,\lambda}$ are Lipschitz continuous with constant $g(\nu, \lambda)$ shows that the mapping $\Phi^{\nu,\lambda} : K_C^\nu \to K_C^\nu$ is strictly contractive on S. The result is then a simple consequence of the Banach fixed-point theorem [23, Theorem 1.A].

Theorem 3.1 indicates that for small enough ν and λ, the sequence of iterates of the map $\Phi^{\nu,\lambda}$ has a unique limit that is independent of the initial condition. Consequently, for such values of ν and λ the marginals of the stationary distribution at nodes in the core of sufficiently large networks would be almost identical. However, increasing ν or λ can give rise to *phase transitions*, defined here to be the existence of ν, λ for which the sequences of iterates of $\Phi^{\nu,\lambda}$ starting at different initial conditions have multiple limit points. Phase transitions destroy the spatial homogeneity of the stationary distribution at the core. For a given λ, we denote the lowest value of ν at which a phase transition occurs by $\nu_{\mathrm{pt}}(\lambda)$. As we see from the discussion below, for $C = 1$ the greater the value of λ, the greater the value of $\nu_{\mathrm{pt}}(\lambda)$. These phenomena are illustrated in more detail for the cases $C = 1$ and $C = 2$ in Sections 3.3 and 3.4 respectively.

Remark 3.1. It is possible to establish a rigourous correspondence between the fixed points of the random field maps $\Phi^{\nu,\lambda}$ and Markov chains on the infinite tree T that are invariant to tree isomorphisms (as defined in [20, p. 895]) or, equivalently, simple invariant Gibbs measures on Ω_T, the set of feasible configurations on T (see [3, Definition 3.7]). The case $\lambda = 0$ can be deduced from [3], [20]. The proof for the case $\lambda > 0$, however, is more involved and requires the introduction of notation and concepts not central to this work. Hence, we defer the details of the proof to a subsequent paper.

3.2. Limit of the blocking probabilities

Recall the expression (2.7) that characterizes α_e^L, the stationary unicast blocking probability of an edge comprising a central node and its neighbour in S_L, the spherical Cayley tree of radius L. Define α_e to be the limiting probability as $L \to \infty$, so that α_e provides an approximation for the blocking probability of edges that are deep in the interior of a large Cayley tree network. Similarly, recall from (2.8) that β_s^L is the probability that a multicast call is blocked from the central node s of S_L, and let β_s be the limit of β_s^L as $L \to \infty$.

Lemma 3.1. *For (ν, λ) such that $\nu < \nu_{\mathrm{pt}}(\lambda)$, let $\xi^* = \xi^*(\nu, \lambda)$ be the unique fixed point of the map $\Phi^{\nu,\lambda}$. Then*

$$1 - \alpha_e = \frac{\sum_{k=0}^{C-1}(\lambda^k/k!)[1 + 2\sum_{i=1}^{C-1-k}\xi^*(i) + \sum_{i=1}^{C-1-k}\sum_{j=1}^{C-1-k-i}\xi^*(i)\xi^*(j)]}{\sum_{k=0}^{C}(\lambda^k/k!)[1 + 2\sum_{i=1}^{C-k}\xi^*(i) + \sum_{i=1}^{C-k}\sum_{j=1}^{C-k-i}\xi^*(i)\xi^*(j)]} \tag{3.3}$$

and

$$1 - \beta_s = \frac{\sum_{i=0}^{C-1}(\nu^i/i!)[\sum_{j=0}^{C-1-i}(\lambda^j/j!)(1 + \sum_{k=1}^{C-1-i-j}\xi^*(k))]^{q+1}}{\sum_{i=0}^{C}(\nu^i/i!)[\sum_{j=0}^{C-i}(\lambda^j/j!)(1 + \sum_{k=1}^{C-i-j}\xi^*(k))]^{q+1}}. \tag{3.4}$$

Proof. Using the definition of ξ_m, we can rewrite (2.7) as

$$1 - \alpha_e^L = \frac{\sum_{k=0}^{C-1}(\lambda^k/k!)[1+\sum_{i=1}^{C-1-k}\xi_L(i)+\sum_{j=1}^{C-1-k}\xi_{L-1}(j)+\sum_{i=1}^{C-1-k}\sum_{j=1}^{C-1-k-i}\xi_L(i)\xi_{L-1}(j)]}{\sum_{k=0}^{C}(\lambda^k/k!)[1+\sum_{i=1}^{C-k}\xi_L(i)+\sum_{j=1}^{C-k}\xi_{L-1}(j)+\sum_{i=1}^{C-k}\sum_{j=1}^{C-k-i}\xi_L(i)\xi_{L-1}(j)]}. \tag{3.5}$$

The expression for $1 - \alpha_e$ follows directly from (3.5) and the fact that $\xi_L \to \xi^*$ as $L \to \infty$. Analogous calculations yield the expression for $1 - \beta_s$.

3.3. The case $C = 1$

In this section, we analyse the performance of the model for the case $C = 1$. As mentioned earlier, the limiting stationary distribution for $C = 1$ with $\lambda = 0$ was obtained in [12], [13]. Here we are also interested in studying the trade-off between unicast blocking and the capacity to support multicast calls.

3.3.1. *The recursion for* $C = 1$. For $C = 1$, (3.2) reduces to the one-dimensional recursion

$$\xi_{m+1} = \frac{\nu}{(1 + \lambda + \xi_m)^q}.$$

The critical curve L_{cr} in (ν, λ)-space of bifurcation points of the two-parameter family of mappings $\Phi^{\nu,\lambda}$ (see e.g. [7]) is defined by the relation

$$\left| \frac{\partial \Phi^{\nu,\lambda}}{\partial \xi}(\xi^*) \right|_{(\nu,\lambda)\in L_{\mathrm{cr}}} = 1,$$

where, as in the previous section, $\xi^* = \xi^*(\nu, \lambda)$ is the fixed point for the mapping $\Phi^{\nu,\lambda}$. Since

$$\left| \frac{\partial \Phi^{\nu,\lambda}}{\partial \xi}(\xi) \right| = \frac{q\nu}{(1 + \lambda + \xi)^{q+1}}$$

and

$$\xi^* = \frac{\nu}{(1 + \lambda + \xi^*)^q}, \tag{3.6}$$

simple algebraic manipulations show that, at criticality, the fixed point is

$$\xi^*_{\mathrm{cr}} = \frac{1 + \lambda}{q - 1},$$

and hence the critical curve L_{cr} is described by the set of points $(\nu_{\mathrm{cr}}(\lambda), \lambda)$ where

$$\nu_{\mathrm{cr}}(\lambda) = \frac{q^q}{(q-1)^{q+1}}(1 + \lambda)^{q+1}. \tag{3.7}$$

The value of the critical points of bifurcation for the family of mappings $\Phi^{\nu,\lambda}$ for $C = 1$ was obtained in the absence of unicast calls ($\lambda = 0$) in [12], [13]. We define the subcritical region to be the set of points (ν, λ) that satisfy $\nu < \nu_{\mathrm{cr}}(\lambda)$ and the supercritical region to be the set of points that satisfy $\nu > \nu_{\mathrm{cr}}(\lambda)$. To see that bistable behaviour cannot exist in the subcritical region, consider the iterated map $\Phi^{\nu,\lambda} \circ \Phi^{\nu,\lambda}(\xi)$. Then it is straightforward to check that (a) $\Phi^{\nu,\lambda} \circ \Phi^{\nu,\lambda}(0) > 0$; (b) $\Phi^{\nu,\lambda} \circ \Phi^{\nu,\lambda}(\cdot)$ is increasing in ξ; (c) $\Phi^{\nu,\lambda} \circ \Phi^{\nu,\lambda}(\cdot)$ has a point of inflection (i.e. its second derivative vanishes and changes sign) at $\tilde{\xi} = [\nu(q-1)/(1+\lambda)]^{1/q} - (1 + \lambda)$; and finally (d) for $\nu < \nu_{\mathrm{cr}}$, $|\partial \Phi^{\nu,\lambda} \circ \Phi^{\nu,\lambda}(\tilde{\xi})/\partial \xi| < 1$. From (d) we can conclude that, for $\nu < \nu_{\mathrm{cr}}$, $|\partial \Phi^{\nu,\lambda} \circ \Phi^{\nu,\lambda}(\xi)/\partial \xi| < 1$ for all ξ, and hence that the iterated map has a unique fixed point for $\nu < \nu_{\mathrm{cr}}$. This yields the following result.

Theorem 3.2. *For* $C = 1$, *the region where there is no phase transition is precisely the subcritical region* $\{(\nu, \lambda) \in \mathbb{R}^2_+ : \nu < \nu_{\mathrm{cr}}(\lambda)\}$, *where* $\nu_{\mathrm{cr}}(\lambda)$ *is given by (3.7).*

In Section 4, we will show that the critical point for the Erlang fixed-point approximation is $\frac{1}{2}\nu_{\mathrm{cr}}$, which provides evidence of the inadequacy of this approximation in this context, although we note that for $\nu < \frac{1}{2}\nu_{\mathrm{cr}}$ the approximation is still very good.

Note that in the asymptotic regime as $q \to \infty$, for the case of a pure multicast network ($\lambda = 0$) we have

$$\lim_{q\to\infty} \nu_{\mathrm{cr}} \approx \lim_{q\to\infty} \frac{\mathrm{e}}{q} = 0.$$

This suggests that in pure multicast networks with a higher degree of connectivity, heterogeneous multicast blocking probabilities in the network will arise at lower arrival rates.

Beyond the critical point, an orbit of period 2 appears, and there is no longer a unique limit point for the sequence of iterates of the map $\Phi^{\nu,\lambda}$. Let the points of this orbit be denoted by ξ_1^* and ξ_2^*. Then we observe that

$$\Phi^{\nu,\lambda}(\xi_1^*) = \xi_2^* \quad \text{and} \quad \Phi^{\nu,\lambda}(\xi_2^*) = \xi_1^*, \tag{3.8}$$

and so, for $i = 1, 2$,

$$\Phi^{\nu,\lambda} \circ \Phi^{\nu,\lambda}(\xi_i^*) = \xi_i^*.$$

Spitzer [18, Theorem 9] and Zachary [20, Theorems 4.1 and 4.3] showed that given any homogeneous Markov specification on a regular Cayley tree, there exists an associated recursion map such that the existence of a 2-period orbit (that is not a fixed point) for the map is equivalent to the existence of a Markov chain that is not invariant under isomorphisms of the tree, but that is invariant under isomorphisms that map the even and odd lattices onto themselves. Thus, ν_{cr} denotes the point at which there appears a Markov chain on the infinite tree that is not translation invariant, but has an alternating pattern on the even and odd lattices.

3.3.2. *Blocking v. capacity trade-offs for* $C = 1$. We now calculate the trade-off in the subcritical region between the volume of multicast and unicast calls that can be supported in the network for a given blocking probability for unicast calls. For simplicity of notation, we let $\alpha = \alpha_e$ be the unicast blocking probability. Then for (ν, λ) in the subcritical region, from the expression (3.3) it follows that

$$1 - \alpha = \frac{1}{1 + \lambda + 2\xi^*},$$

where ξ^* is the unique fixed point of the map $\Phi^{\nu,\lambda}$. Rearranging the above equation yields

$$\xi^* = \frac{\alpha}{2(1-\alpha)} - \frac{\lambda}{2}.$$

Using the fact that ξ^* satisfies the fixed point equation (3.6), for (ν, λ) in the subcritical region we obtain

$$\nu = \left(\frac{\alpha}{2(1-\alpha)} - \frac{\lambda}{2}\right)\left(1 + \frac{\alpha}{2(1-\alpha)} + \frac{\lambda}{2}\right)^q. \tag{3.9}$$

The above equation is plotted for different values of α for the case $q = 2$ in Figure 3. In the subcritical region, the curves characterize the trade-off between multicast and unicast arrival rates for fixed α and are marked as solid curves, whereas in the supercritical region the curves no longer have this physical interpretation, and are thus marked with dashed lines.

It is intuitively clear that as the unicast arrival rate decreases, the multicast arrival rate must increase in order for the unicast blocking probability to be maintained at α. Indeed, as illustrated in Figure 3, the ν–λ trade-off curves are monotone. In addition they are also convex, which means that, for every fixed α, the rate of increase of the multicast arrival rate decreases with decrease in the unicast arrival rate. Let $\bar{\lambda} = \bar{\lambda}(\alpha)$ be the point at which $\partial\nu/\partial\lambda$ becomes zero,

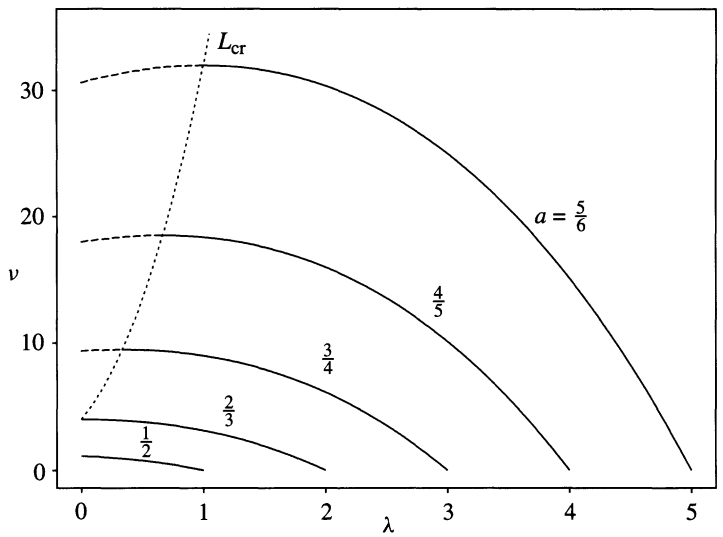

FIGURE 3: Trade-off between unicast and multicast calls ($C = 1, q = 2$).

and let $\bar{\nu} = \bar{\nu}(\alpha)$ be the corresponding value of ν at which this happens. Let $\theta = \alpha/(1-\alpha)$, and note that

$$\frac{\partial \nu}{\partial \lambda} = -\frac{1}{2}\left(1 + \frac{\theta}{2} + \frac{\lambda}{2}\right)^q + \frac{q}{4}(\theta - \lambda)\left(1 + \frac{\theta}{2} + \frac{\lambda}{2}\right)^{q-1}.$$

Then by the above equation we conclude that

$$\bar{\lambda} = \frac{(q-1)\theta}{q+1} - \frac{2}{q+1},$$

and substituting this back in (3.9), we see that

$$\bar{\nu} = \frac{(\theta+1)^{q+1}}{(q+1)^{q+1}} q^q.$$

Since

$$1 + \bar{\lambda} = \frac{(q-1)(\theta+1)}{q+1},$$

it follows that $\bar{\nu}$ satisfies

$$\bar{\nu} = \frac{q^q}{(q-1)^{q+1}}(1+\bar{\lambda})^{q+1}.$$

Comparing this with (3.7), we see that the locus of points $(\bar{\lambda}(\alpha), \bar{\nu}(\alpha))$ for $\alpha \in (0, 1)$ coincides with the critical phase transition curve L_{cr}. Thus we note the remarkable feature that (at least for the case $C = 1$) phase transitions occur precisely at points $(\bar{\lambda}, \bar{\nu})$ at which the rate of increase of the multicast arrival rate with decrease in the unicast arrival rate (for a fixed unicast blocking probability) is zero.

Now, for (ν, λ) beyond the critical curve, let ξ_1^*, ξ_2^* be the two period orbit defined in (3.8). Taking $L \to \infty$ in (3.5), it can be seen that the blocking probability α in the supercritical region is given by

$$1 - \alpha = \frac{1}{1 + \lambda + \xi_1^* + \xi_2^*}.$$

Thus, although there is a bifurcation of solutions for large enough values of (ν, λ), the unicast blocking probabilities are still homogeneous throughout the tree. Moreover, we observe that since the two points ξ_1^* and ξ_2^* move in a continuous fashion away from the fixed point as the critical curve is crossed, the unicast blocking probabilities change continuously through the phase transition. As we show later, this is not the case when the link capacity $C = 2$ (and, more generally, for even link capacities C), where we could have an abrupt change in the unicast blocking probabilities at certain values (ν, λ) in the subcritical region of the parameter space (i.e. where the fixed point is still locally stable).

We now consider the multicast blocking probability. For $C = 1$, from (3.4) we have, for (ν, λ) in the subcritical region,

$$1 - \beta = \frac{1}{\nu + [1 + \lambda + \xi^*]^{q+1}} = \frac{1}{\nu[1 + (1 + \lambda + \xi^*)/\xi^*]},$$

where the last equality follows from the fixed-point equation (3.6). Rearranging, we find that

$$\xi^* = \frac{\nu(\lambda + 1)(1 - \beta)}{1 - 2\nu(1 - \beta)},$$

and therefore

$$1 + \lambda + \xi^* = \frac{(1 + \lambda)(1 - \nu(1 - \beta))}{1 - 2\nu(1 - \beta)}.$$

Note that the above two equations imply that, in the subcritical region, $1 - 2\nu(1 - \beta) > 0$. Substituting the last two displays into the fixed point equation (3.6), we see that

$$\lambda = -1 + (1 - \beta)^{-1/q+1} \frac{1 - 2\nu(1 - \beta)}{[1 - \nu(1 - \beta)]^{q/(q+1)}}.$$

Analogous to the trade-off curves for fixed unicast blocking probabilities, here too it is clear that with a decrease in the multicast arrival rate, to maintain the same multicast blocking probability, the unicast arrival rate should increase. In this case, however, beyond the phase transition curve, the multicast blocking probabilities bifurcate, and take on different values depending on whether the node lies on the 'even' or 'odd' lattice.

3.4. The case $C = 2$

In this section, we analyse the performance of the model for the case where the capacity of each edge (link) is $C = 2$. As we will see below, this map displays behaviour that is not found for the $C = 1$ case. In particular, it is possible to observe a discontinuous change in the blocking probabilities with a change in parameter values. As discussed in Section 3.5, maps with even link capacity C seem to behave like $C = 2$, while maps with odd link capacity C behave like $C = 1$.

3.4.1. *The recursion for $C = 2$.* For $C = 2$ the recursion defined in (3.1) yields the following two-dimensional map:

$$\Phi_1^{\nu,\lambda}(\xi) = \nu\left[\frac{1+\lambda+\xi(1)}{1+\lambda+\frac{1}{2}\lambda^2+(1+\lambda)\xi(1)+\xi(2)}\right]^q,$$
$$\Phi_2^{\nu,\lambda}(\xi) = \frac{\nu^2}{2}\left[\frac{1}{1+\lambda+\frac{1}{2}\lambda^2+(1+\lambda)\xi(1)+\xi(2)}\right]^q.$$

Setting $\xi(1) = \xi$ and $\xi(2) = \eta$, and computing partial derivatives of the map, we see that the Jacobian at (ξ, η) is given by

$$J(\xi,\eta) = \frac{-q\nu}{(1+\lambda+\frac{1}{2}\lambda^2+(1+\lambda)\xi+\eta)^{q+1}}\begin{bmatrix}(1+\lambda+\xi)^{q-1}(\lambda+\frac{1}{2}\lambda^2-\eta) & (1+\lambda+\xi)^q\\ \frac{1}{2}\nu(1+\lambda) & \frac{1}{2}\nu\end{bmatrix}.$$

Now we know from Theorem 3.1 that the map has a unique fixed point

$$(\xi^*, \eta^*) = (\xi^*(\nu,\lambda), \eta^*(\nu,\lambda))$$

for all sufficiently small ν and λ. At the fixed point, the Jacobian has the value

$$J(\xi^*,\eta^*) = \frac{-q}{1+\lambda+\frac{1}{2}\lambda^2+(1+\lambda)\xi^*+\eta^*}\begin{bmatrix}\frac{\xi^*(\lambda+\frac{1}{2}\lambda^2-\eta^*)}{1+\lambda+\xi^*} & \xi^*\\ \eta^*(1+\lambda) & \eta^*\end{bmatrix}.$$

Thus the eigenvalues at the fixed point (ξ^*, η^*) are given by

$$\frac{q}{2[1+\lambda+\frac{1}{2}\lambda^2+(1+\lambda)\xi^*+\eta^*]}\times\left[-\frac{\eta^*(1+\lambda)+\xi^*\lambda(1+\frac{1}{2}\lambda)}{1+\lambda+\xi^*}\right.$$
$$\left.\pm\sqrt{\frac{(\eta^*(1+\lambda)+\xi^*\lambda(1+\frac{1}{2}\lambda))^2}{(1+\lambda+\xi^*)^2}+\frac{4[\xi^*\eta^{*2}+\xi^*\eta^*+\xi^{*2}\eta^*+\xi^*\eta^*\lambda(1+\xi^*+\frac{1}{2}\lambda)]}{1+\lambda+\xi^*}}\right].$$

In the pure multicast case, when $\lambda = 0$ the eigenvalues are given by

$$\frac{q}{2(1+\xi^*+\eta^*)}\left[-\frac{\eta^*}{1+\xi^*}\pm\sqrt{\frac{\eta^{*2}}{(1+\xi^*)^2}+\frac{4\xi^*\eta^{*2}}{1+\xi^*}+4\xi^*\eta^*}\right].$$

We know from Theorem 3.1 that for small ν the fixed point (ξ^*, η^*) is globally stable in the sense that, starting from any initial condition, the iterates converge to that fixed point. In order to determine when this fixed point loses its local stability, it is necessary to calculate the bifurcation point ν_{cr}, which is the point at which $|\mathrm{Re}(\kappa)| = 1$, where κ is the eigenvalue of the Jacobian. Since in this case the eigenvalues are real, this corresponds to finding the point at which $|\kappa| = 1$. If ξ^*, η^* could be expressed as a function of ν, then by setting the absolute value of the above expression equal to 1, we could determine ν_{cr}. However, since such an explicit expression does not seem feasible, we use numerics and approximations to discern the behaviour of the map.

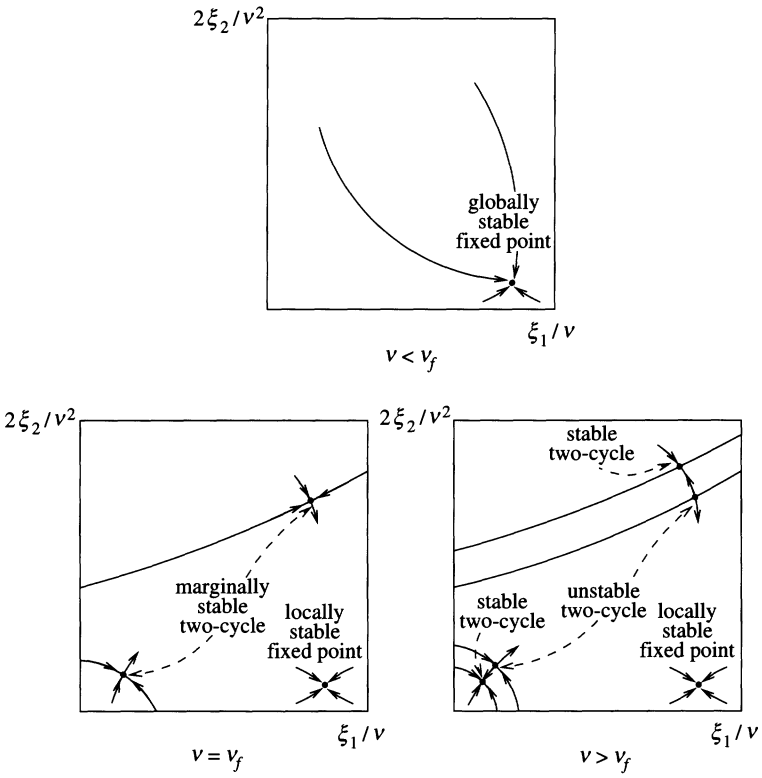

FIGURE 4: First-order phase transitions ($C = 2, \lambda = 0$).

For large ν, it is easy to see that the fixed point is approximately given by $(\nu, \frac{1}{2}\nu^{2-q})$. Thus, as $\nu \to \infty$, the absolute value of both eigenvalues tends to zero—implying that locally the fixed point becomes extremely stable, and thus that there is no finite ν_{cr}. However, although the fixed point maintains its local stability for arbitrary large values of ν, it loses global stability at a point ν_f when a marginally stable two-period orbit appears in addition to the locally stable fixed point. As ν becomes large, the marginal two-cycle splits into a stable and an unstable cycle. The domain of attraction of the unstable two-cycle defines the boundary between the basins of attraction of the fixed point and the stable two-cycle. This phenomenon is reminiscent of what is referred to as a first order transition in statistical physics, in which (at $\nu = \nu_f$) a new stable state appears whose orbits are far away from the prior stable state described by the fixed point (refer to Figure 4). Hence for $C = 2$, the phase transition point ν_{pt} coincides with ν_f. In contrast, recall that for the case $C = 1$ the phase transition occurs at ν_{cr}, the value at which the stable fixed point loses even its local stability and bifurcates into a stable two-cycle whose orbits lie close to the fixed point. Such gradual transitions are referred to as second-order transitions in the statistical mechanics literature.

3.5. Higher link capacity

We have examined the cases $C = 1$ and $C = 2$ with some care. Clearly, in a realistically sized network, C will be considerably larger. However, we conjecture that in some respects the

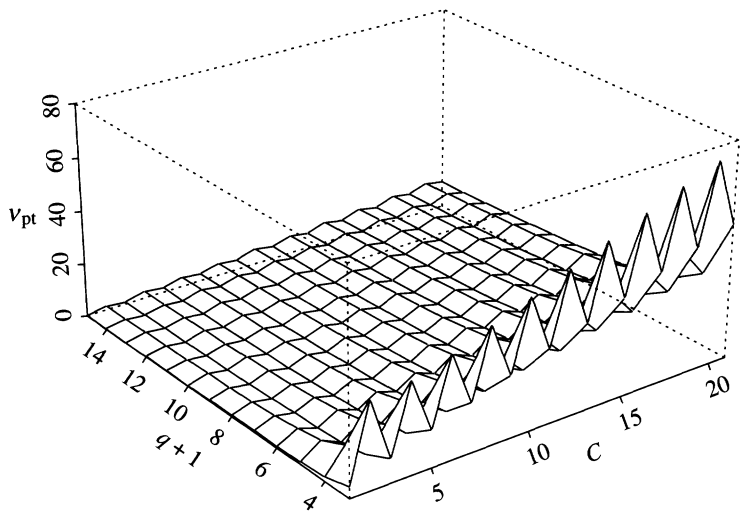

FIGURE 5: Plot of ν_{pt} for varying q and C values ($\lambda = 0$).

qualitative behaviour will be similar. In particular we conjecture that for odd C, entry to the phase transition region will be with a second order phase transition ($\nu_{pt} = \nu_{cr}$), as with $C = 1$; while for even C, the entry to the phase transition region will be accomplished via a first order phase transition ($\nu_{pt} = \nu_f$). (Recall that ν_{pt} is defined to be the lowest value of ν at which a phase transition occurs.) Numerical results seem to indicate that for higher C there will be a number of phase transitions, corresponding to modes of the distribution where, say, nodes in the even lattice carry i calls and nodes in the odd lattice carry j calls, with $i + j = C$. In Figure 5, we plot numerically obtained values for ν_{pt} for $\lambda = 0$ and varying C and q. We see that these numerical results indicate that, as q increases, the phase transition point decreases, and that ν_{pt} is of a lower order of magnitude for odd C, when compared to even C, particularly for small q. The latter behaviour may be linked with our conjecture about the differing nature of the phase transitions for odd and even C.

4. The Erlang fixed-point approximation

We have observed in previous sections that, for (ν, λ) in the region prior to phase transition, the blocking probability in the centre of S_L converges as $L \to \infty$ to the fixed point of a map, and that link blocking probabilities in the centre of the tree are homogeneous. This suggests that the natural formulation of the Erlang fixed-point (EFP) approximation, at least initially, is one where the blocking probability on every link is the same, assuming that the tree is infinite. This then gives the blocking probability, B, for a single link as a fixed point of the map

$$f^{\nu,\lambda}(B) = E(\lambda + 2\nu(1 - B)^q, C), \tag{4.1}$$

where $E(\nu, C) := (\sum_{i=0}^{C} \nu^i / i!)^{-1} \nu^C / C!$ is the Erlang-B blocking formula [13], [17]. Thus, each link sees a reduced load of $\lambda + 2\nu(1 - B)^q$ offered to it. Recall that λ is the arrival rate or load of unicast calls at every link, and there are also two streams of multicast calls offered to every link, each centred at one end of the link, each arriving at rate ν and thinned by a factor of $(1 - B)^q$.

It would, of course, also be possible to consider the Erlang fixed-point approximation to the blocking probabilities at the centre of finite trees, as they become large, without *a priori* assuming that they are the same on every link. For the tree S_L, this involves finding the fixed point of an L-dimensional map. Numerical results indicate that, for (ν, λ) in the subcritical region of the map (4.1), the approximation to the blocking probability of a central link in S_L obtained from the full L-dimensional map converges to the fixed point of (4.1). We therefore confine our attention to (4.1) for the rest of this section. We note, however, that our numerical results indicate that beyond the critical point of the EFP map, the single-link blocking probability corresponding to that of the central link for the L-dimensional EFP approximation converges to the unstable fixed point of the map (4.1).

Let $B^*(\nu, \lambda)$ be the unique fixed point of the mapping (4.1), which is guaranteed to exist for small enough λ and ν (by analogy with Theorem 3.1). As in Section 3.3, by definition the values (ν, λ) on the critical curve M_{cr} of bifurcation points of the family of mappings $f^{\nu,\lambda}$ (see e.g. [7]) must satisfy

$$\left| \frac{\partial f^{\nu,\lambda}}{\partial B}(B^*) \right|_{(\nu,\lambda)\in M_{\text{cr}}} = 1.$$

When $C = 1$,

$$f^{\nu,\lambda}(B) = \frac{\lambda + 2\nu(1-B)^q}{1 + \lambda + 2\nu(1-B)^q},$$

so that

$$\left| \frac{\partial f^{\nu,\lambda}}{\partial B} \right| = \frac{2q\nu(1-B)^{q-1}}{(1 + \lambda + 2\nu(1-B)^q)^2},$$

and at the fixed point $B^* = B^*(\nu, \lambda)$ we obtain

$$1 - B^* = \frac{1}{1 + \lambda + 2\nu(1-B^*)^q}.$$

Together these show, after simple algebraic manipulations, that at criticality

$$B^*_{\text{cr}} = \frac{1 + \lambda q}{q(1+\lambda)},$$

and the critical curve $L^{\text{EFP}}_{\text{cr}}$ for the EFP map is described by the set of points $(\nu^{\text{EFP}}_{\text{cr}}(\lambda), \lambda)$, where

$$\nu^{\text{EFP}}_{\text{cr}}(\lambda) = \frac{1}{2q}\left(1 - \frac{1}{q}\right)^{-(q+1)} (1+\lambda)^{q+1},$$

which is equal to exactly half the expression (3.7) for $\nu_{\text{cr}}(\lambda)$ of the random field map $\Phi^{\nu,\lambda}$ found in Section 3.

Figure 6 illustrates the attracting solutions to the various maps for the case $C = 1$, with $\lambda = 0$, $q = 2$ and ν ranging between 0 and 6. Figure 6(a) gives the multicast blocking probabilities β for the random field maps $\Phi^{\nu,0}$ (which when $C = 1$ can be found in Section 3.3), as well as exact blocking probabilities at the centre of S_L, when the radius $L = 10, 11$ (β_s^{10} and β_s^{11}) and when the radius $L = 40, 41$ (β_s^{40} and β_s^{41}). These all assume free boundary conditions. We note that as the radius L increases, the blocking probabilities at the centre of S_L approach those of the random field map, following either the upper or lower branch, depending on whether L, the radius of the tree, is odd or even. Figure 6(b) gives the multicast blocking

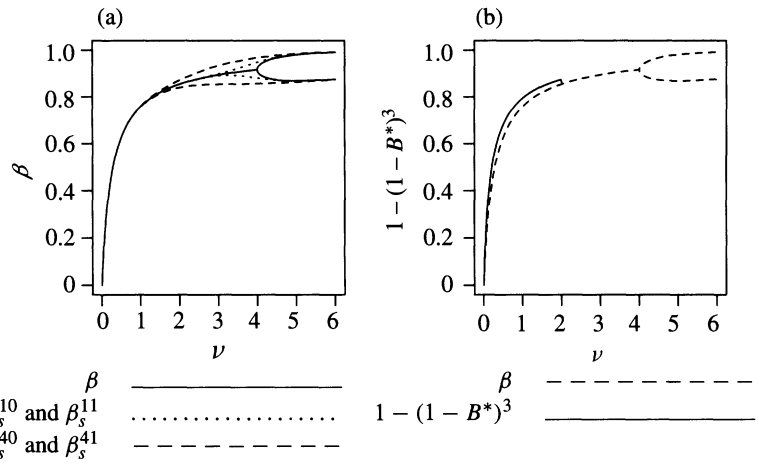

FIGURE 6: Comparison of blocking probabilities for the random field and EFP maps ($\lambda = 0$, $C = 1$, $q = 2$): (a) multicast blocking probability, β, obtained from the random field map, (b) multicast blocking probability, $1-(1-B^*)^3$, obtained from the EFP approximation.

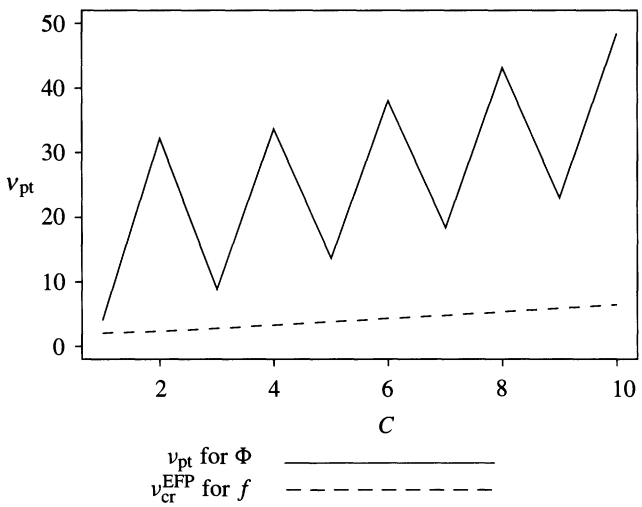

FIGURE 7: Comparison of phase transitions for the random field and EFP maps ($\lambda = 0$, $q = 2$).

probabilities obtained from the Erlang fixed-point approximation below criticality of that map (for comparison purposes the blocking probabilities obtained for the maps $\Phi^{\nu,0}$ have also been plotted in Figure 6(b)). Note that if B^* is a fixed point of the EFP map, then the multicast blocking probability is given by $1-(1-B^*)^{q+1}$. We observe that the family of random field maps $\Phi^{\nu,0}$ bifurcates at $\nu = 4$, while the family of EFP maps $f^{\nu,0}$ bifurcates at $\nu = 2$, in agreement with the results of this section and those of Section 3.

For higher C too, there are inadequacies in the EFP approximation. The most striking is that the bifurcation point for the EFP approximation occurs at much lower values than for the random field map $\Phi^{\nu,0}$, particularly for even C. Figure 7 shows ν_{pt}, the lowest value of ν for which a phase transition occurs, for both the random field maps $\Phi^{\nu,0}$ and the EFP maps $f^{\nu,0}$

when $q = 2$, and for C ranging from 1 to 10. We note that, unlike the case of the random field map, for the EFP map this *increases monotonically* with C, and no difference is apparent in the behaviour for odd and even values of the link capacity C. In addition, for values of ν past the bifurcation point for the EFP maps (i.e. $\nu > \nu_{\text{cr}}^{\text{EFP}}$), the stable solution to the map $f^{\nu,\lambda}$ in (4.1) is a two-cycle. However, we have shown in Section 3.3 that single-link blocking probabilities are homogeneous in the centre of the tree even past criticality, and do not oscillate. It appears that beyond the critical region, approximations to the single-link blocking probabilities should be obtained from the *unstable fixed point* of the EFP approximation, but it is no longer clear how this can be used to approximate the multicast blocking probabilities.

5. Some generalizations

A simple symmetric network such as the one we have studied above is, of course, unrealistic. The model could be generalized in many ways. Two of these we study in greater detail below. The first is to allow longer multihop calls into the system. This is a natural extension of the model, and essentially only complicates the formulation of normalizing constants, blocking probabilities etc. The second generalization that we study breaks the symmetry of the model by allowing heterogeneous ν. The interest here is in seeing whether phase transitions still occur. Van den Berg and Steif [19] have conjectured that if arrival rates at odd and even nodes on the *cubic* lattice $\mathbb{Z}^d$ are not equal, then the model has a unique Gibbs measure, and they give a partial proof of this conjecture. However, we find that on a tree network, phase transitions can still occur with heterogeneous ν.

5.1. Multihop multicast calls

In this section, we indicate how to generalize this model to multicasts beyond nearest neighbours on the tree. For simplicity, we consider a model where there are no unicast calls ($\lambda = 0$), but only multicast calls with the property that a multicast call originating at a node t is broadcast to all its nearest neighbours $t' \in \mathcal{N}(t)$ and next-nearest neighbours $t'' \in \mathcal{N}_2(t) := \bigcup_{t' \in \mathcal{N}(t)} \mathcal{N}(t') \setminus t$. The number of calls originating at a node t is denoted n_t, as usual, and the total number of types of calls is $M_T = |T|$, the number of nodes in the tree T. The set of feasible configurations is then

$$\Omega_T := \left\{ n \in \mathcal{I}^{M_T} : \text{ for each } (t_1 t_2) \in \mathcal{E}(T), \sum_{t \in \mathcal{N}(t_1) \cup \mathcal{N}(t_2)} n_t \leq C \right\}.$$

Once more, the stationary distribution π has the truncated product form

$$\pi(n) = \frac{1}{\mathcal{Z}_\pi} \prod_{t \in T} \frac{\nu^{n_t}}{n_t!}$$

for $n \in \Omega_T$, where $\mathcal{Z}_\pi$ is the normalizing constant

$$\mathcal{Z}_\pi := \sum_{n \in \Omega_T} \prod_{t \in T} \frac{\nu^{n_t}}{n_t!}.$$

As before, we calculate $\mathcal{Z}_\pi$ in terms of some partial sums, to be defined below, and then derive recursion relations for these partial sums. To show how to derive the recursion relations for this model, we restrict ourselves to the case $C = 1$ and $q = 2$. Let t_0 be the central node in a spherical tree, and let t_{-1}, t_1 and t_2 be its neighbours, with t_{-2} and t_{-3} being neighbours of t_{-1}

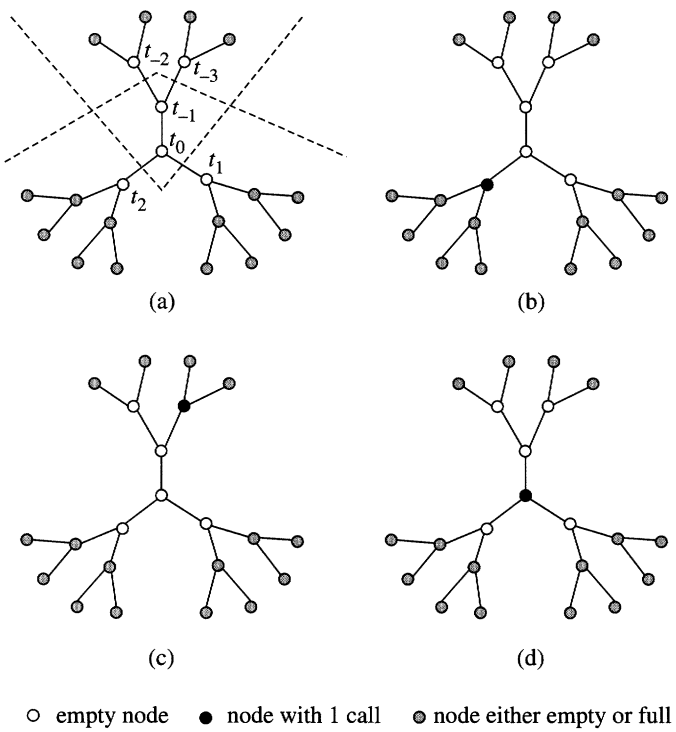

FIGURE 8: Multihop multicasting on S_3 ($C = 1$, $q = 2$): (a) augmented trees of size 4 and 3 and configuration with weight $A_{L-1}A_L$; (b) $A_{L-1}\Gamma_L$; (c) $A_L\Gamma_{L-1}$; (d) $\Delta_L B_{L-1}/\nu$.

(see for example Figure 8). There are only two possible types of configurations of calls on these sites: either there is no call on any site, or there is only one call originating from exactly one of these six sites. These possibilities give us seven different configurations.

We can find a rooted tree T_m for some m, such that T_m is rooted at t_0. Without loss of generality, we could assume that $t_1, t_2 \in T_m$. Now we define

$$Z_m(i_{-1}, i_0, i_1, i_2) := \sum_{\substack{n \in \Omega_m : n_{t_a} = i_a, \\ a=-1,\ldots,2}} \prod_{t \in T_m} \frac{\nu^{n_t}}{n_t!},$$

so that Z_m is the weighted sum of all configurations in the graph comprising the rooted tree T_m with root t_0, augmented by the node t_{-1}. As shown in Figure 8(a), any spherical tree of radius L with central node t_0 can be decomposed into two overlapping augmented trees of sizes L and $L-1$ with roots t_0 and t_{-1} respectively. Then Z_π can be expressed as a weighted sum of the seven configurations on the six nodes $t_{-3}, t_{-2}, t_{-1}, t_0, t_1, t_2$, corresponding to the empty configuration with no calls on any of the nodes, and the six configurations associated with a call on one of the nodes. Introducing the notation $A_m = Z_m(0,0,0,0)$, $B_m = Z_m(1,0,0,0)$, $\Gamma_m = Z_m(0,0,1,0) = Z_m(0,0,0,1)$, $\Delta_m = Z_m(0,1,0,0)$, it is clear that $A_L A_{L-1}$ represents the weight for the empty configuration shown in Figure 8(a), $A_{L-1}\Gamma_L$ and $A_L\Gamma_{L-1}$ represent the weights for the configurations given in Figures 8(b) and 8(c) respectively. Likewise, $\Delta_L B_{L-1}/\nu$

corresponds to the weight of the configuration in Figure 8(d), where the dividing factor ν arises due to the fact that the occupied node t_0 lies in both augmented trees. Arguing similarly for the remaining configurations, it follows that the normalization constant $\mathcal{Z}_\pi$ is given by

$$\mathcal{Z}_\pi = A_{L-1}A_L + 2(\Gamma_{L-1}A_L + A_{L-1}\Gamma_L) + \nu^{-1}(B_{L-1}\Delta_L + \Delta_{L-1}B_L)$$

for a spherical tree of radius L.

It is easy to see that the recursion relations are:

$$\begin{aligned} A_{m+1} &= (A_m + 2\Gamma_m)^2, \\ B_{m+1} &= \nu A_m^2, \\ \Gamma_{m+1} &= \Delta_m A_m, \\ \Delta_{m+1} &= \frac{B_m^2}{\nu}. \end{aligned}$$

The free boundary condition is imposed by taking $A_0 = 1, \Gamma_0 = 0, B_0 = \Delta_0 = \nu$.

As usual, the recursion relations are homogeneous in $A_m, B_m, \Gamma_m, \Delta_m$, and we could instead obtain recursions on the three ratios $\beta_m = B_m/A_m, \gamma_m = \Gamma_m/A_m, \delta_m = \Delta_m/A_m$ as follows:

$$\begin{aligned} \beta_{m+1} &= \frac{\nu}{(1+2\gamma_m)^2}, \\ \gamma_{m+1} &= \frac{\delta_m}{(1+2\gamma_m)^2}, \\ \delta_{m+1} &= \frac{\beta_m^2}{\nu(1+2\gamma_m)^2}. \end{aligned}$$

Once more, for small ν, these recursion relations lead to a unique fixed point. At large ν, there is a four-cycle leading to a phase transition. This corresponds to an ordered state which breaks homogeneity (i.e. translational symmetry).

5.2. Heterogeneous ν

In this section, we consider the model with heterogeneous ν. We assume that multicast calls arrive at either rate ν_1 or ν_2 at each node, and that these arrival rates alternate across the tree. In particular, we assume that in the spherical tree, nodes at a distance m from the nearest terminal node have arrival rate ν_1 if m is odd and ν_2 if m is even, and that terminal nodes have arrival rate ν_2. Thus, each link has multicast streams of rate ν_1 and of rate ν_2 offered to it.

The expressions obtained in Section 3 then need to be modified in the obvious fashion. The recursions (3.1) now become $\xi_{m+1} = \Phi^{\nu_1,\lambda}(\xi_m)$ if $m+1$ is odd and $\xi_{m+1} = \Phi^{\nu_2,\lambda}(\xi_m)$ if $m+1$ is even. Let ξ_1^* and ξ_2^* be fixed points for the recursions $\Phi^{\nu_1,\lambda} \circ \Phi^{\nu_2,\lambda}$ and $\Phi^{\nu_2,\lambda} \circ \Phi^{\nu_1,\lambda}$ respectively. (Note that due to Theorem 3.1, they certainly exist and are globally stable for $\nu_1 = \nu_2 = \nu$ sufficiently small.) Then in the subcritical region the limiting link blocking probability in (3.3) satisfies

$$1-\alpha = \frac{\sum_{k=0}^{C-1}(\lambda^k/k!)[1+\sum_{i=1}^{C-1-k}\xi_1^*(i)+\sum_{j=1}^{C-1-k}\xi_2^*(j)+\sum_{i=1}^{C-1-k}\sum_{j=1}^{C-1-k-i}\xi_1^*(i)\xi_2^*(j)]}{\sum_{k=0}^{C}(\lambda^k/k!)[1+\sum_{i=1}^{C-k}\xi_1^*(i)+\sum_{j=1}^{C-k}\xi_2^*(j)+\sum_{i=1}^{C-k}\sum_{j=1}^{C-k-i}\xi_1^*(i)\xi_2^*(j)]}.$$

We note that this is symmetric in ξ_1^* and ξ_2^*, thus implying that even with these heterogeneous multicast arrival rates, the link blocking probabilities are homogeneous. The limiting multicast call acceptance probabilities in the subcritical region are now given by

$$1-\beta_1 = \frac{\sum_{i=0}^{C-1}(\nu_1^i/i!)[\sum_{j=0}^{C-1-i}(\lambda^j/j!)(1+\sum_{k=1}^{C-1-i-j}\xi_2^*(k))]^{q+1}}{\sum_{i=0}^{C}(\nu_1^i/i!)[\sum_{j=0}^{C-i}(\lambda^j/j!)(1+\sum_{k=1}^{C-i-j}\xi_2^*(k))]^{q+1}}$$

for odd nodes, and

$$1-\beta_2 = \frac{\sum_{i=0}^{C-1}(\nu_2^i/i!)[\sum_{j=0}^{C-1-i}(\lambda^j/j!)(1+\sum_{k=1}^{C-1-i-j}\xi_1^*(k))]^{q+1}}{\sum_{i=0}^{C}(\nu_2^i/i!)[\sum_{j=0}^{C-i}(\lambda^j/j!)(1+\sum_{k=1}^{C-i-j}\xi_1^*(k))]^{q+1}}$$

for even nodes.

When $C = 1$, these equations are particularly simple. We use Φ_1 and Φ_2 to denote $\Phi^{\nu_1,0}$ and $\Phi^{\nu_2,0}$ respectively, and define $\Phi := \Phi_1 \circ \Phi_2$. Then note that by the definition given above, ξ_1^* is a fixed point of Φ. The recursions reduce to

$$\Phi_1(\xi) = \frac{\nu_1}{(1+\lambda+\xi)^q} \quad \text{and} \quad \Phi_2(\xi) = \frac{\nu_2}{(1+\lambda+\xi)^q},$$

which together then yield

$$\Phi(\xi) = \frac{\nu_1}{(1+\nu_2/(1+\xi)^q)^q}.$$

The limiting acceptance probabilities are then given by

$$1-\alpha = \frac{1}{1+\lambda+\xi_1^*(1)+\xi_2^*(1)},$$
$$1-\beta_1 = \frac{1}{\nu_1+(1+\lambda+\xi_2^*(1))^{q+1}},$$
$$1-\beta_2 = \frac{1}{\nu_2+(1+\lambda+\xi_1^*(1))^{q+1}}.$$

One question of interest here is whether the phase transition (i.e. the existence of multiple limit points for a sequence of iterates of Φ) that we saw with homogeneous ν can also occur when $\nu_1 \neq \nu_2$. The following theorem shows that indeed it can.

Theorem 5.1. *For $C = 1$ there exists an open set $\mathcal{O}$ such that, for $(\nu_1, \nu_2) \in \mathcal{O}$, there exists a sequence of iterates of the map $\Phi = \Phi_1^{\nu_1,0} \circ \Phi_2^{\nu_2,0}$ that has multiple limit points.*

Proof. Assume that $\nu_1, \nu_2 > 0$. Note first that $\Phi(\cdot)$ has first and second derivatives

$$\Phi'(\xi) := \frac{d\Phi}{d\xi}(\xi) = \frac{q^2\nu_1\nu_2}{(1+\nu_2/(1+\xi)^q)^{q+1}} \frac{1}{(1+\xi)^{q+1}},$$

and

$$\frac{d^2\Phi}{d\xi^2}(\xi) = \frac{-q^2(q+1)\nu_1\nu_2}{(1+\xi+\nu_2/(1+\xi)^{q-1})^{q+2}}\left(1-\frac{\nu_2(q-1)}{(1+\xi)^q}\right).$$

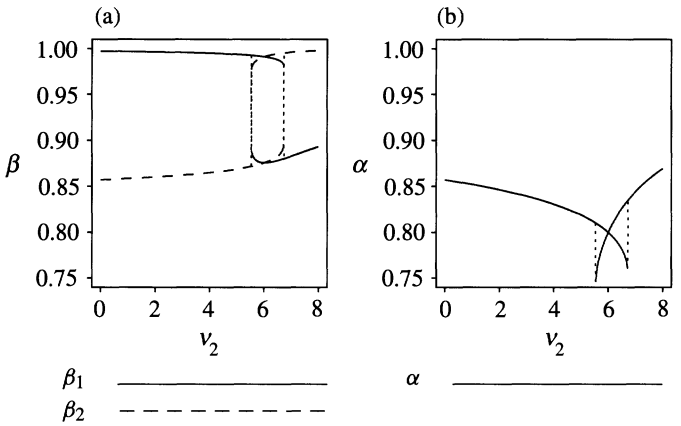

FIGURE 9: Blocking probabilities with heterogeneous ν ($\lambda = 0$, $C = 1$, $q = 2$, $\nu_1 = 6$): (a) multicast (node) blocking probability, (b) unicast (link) blocking probability.

Thus $\Phi(0) = \nu_1/(1+\nu_2)^q > 0$, $\lim_{\xi\to\infty}\Phi(\xi) = \nu_1$ and $\Phi'(\xi) > 0$ for $\xi \geq 0$. When $\nu_2 \leq 1/(q-1)$, $\Phi''(\xi) < 0$ for $\xi > 0$, and so there is a unique fixed point for the map Φ. When $\nu_2 > 1/(q-1)$, however, $\Phi''(0) > 0$, and $\Phi(\cdot)$ has a single point of inflection at the point $\tilde{\xi} = (\nu_2(q-1))^{1/q} - 1$, which is increasing in ν_2. (Note that the point of inflection ξ is such that Φ is convex for $\xi < \tilde{\xi}$ and concave for $\xi > \tilde{\xi}$.) At the point of inflection $\Phi(\tilde{\xi}) = \nu_1(1-1/q)^q$, which does not depend on ν_2, and $\Phi'(\tilde{\xi}) = \nu_1(q-1)^{q-1/q}/(\nu_2^{1/q}q^{q-1})$. If $\Phi'(\tilde{\xi}) < 1$, then a single fixed point is guaranteed, since the derivative is greatest at the point of inflection $\tilde{\xi}$.

If $\tilde{\xi}$ is also a fixed point of $\Phi(\cdot)$, then simple algebra shows that $\tilde{\xi} = [\nu_2(q-1)]^{1/q} - 1 = \nu_1(q-1)/q$, and $\Phi'(\tilde{\xi}) = q/(1+q^q/(\nu_1(q-1)^q)) = q - q/[\nu_2(q-1)]^{1/q}$. Moreover, $\Phi'(\tilde{\xi}) > 1$ if and only if $\nu_1 > q^q/((q-1)^{q+1})$ (or, equivalently, $\nu_2 > q^q/((q-1)^{q+1})$), and in this case $\tilde{\xi}$ is clearly an unstable fixed point. The properties of Φ described above then imply that there also exist two stable fixed points $\xi_1^* < \tilde{\xi} < \xi_2^*$. Thus, given any $\nu_1 > q^q/(q-1)^{q+1}$, by continuity it is clear there exists an open interval (a_1, a_2) around $\nu_2 = (\nu_1(q-1)/q+1)^q/(q-1)$ (i.e. the value at which $\tilde{\xi}$ is a fixed point) such that there are two stable points $\xi_1^* < \tilde{\xi} < \xi_2^*$ and an unstable fixed point $\xi_u^* \in (\xi_1^*, \xi_2^*)$. In addition, due to the form of the function Φ, it is clear that for $\nu_2 < a_1$ the unique fixed point $\xi^* < \tilde{\xi}$, and for $\nu_2 > a_2$, the unique fixed point $\xi^* > \tilde{\xi}$.

We illustrate the above theorem with an example. Figure 9(a) gives the limiting blocking probabilities for multicast calls at odd nodes and even nodes (β_1 and β_2 respectively) for fixed $\nu_1 = 6.0$, and ν_2 varying from 0.0 to 8.0, with $C = 1$ and $\lambda = 0$. Figure 9(b) gives the link blocking probability for the same range of parameters. Again, all of these were obtained numerically under the assumption of free boundary conditions. We note that, as ν_2 increases, the blocking probability for multicast calls at odd nodes, β_1, increases. However, the blocking probability for multicast calls at even nodes, β_2, is very close to 1.0 when ν_2 is 0 and then *decreases* initially as ν_2 increases. The explanation for this is that increasing ν_2 from 0 breaks the packing symmetry for the odd calls, thus giving a slightly lower blocking probability for the even calls. We note that this initial decrease does not occur for small ν_1 (e.g. less than 0.1). There is evidence of bistable behaviour (i.e. multiple solutions for β_1, β_2 and α) over a range of values of ν_2, including the point $\nu_2 = 6.0$.

6. Conclusions

In this paper we developed a framework in which to analyse the effect of multicast calls in large unicast tree networks. Although most of our analysis is carried out for the case of single-hop multicast calls in a regular tree network with uniform arrival rates throughout the network, we showed in Section 5 that the qualitative phenomena observed persist even in the case of multihop calls and asymmetrical arrival rates, and are thus of relevance to more general networks.

6.1. Phase transitions, fairness and capacity

Our main finding is that, at high enough multicast call arrival rates, the stationary call distribution at the core of large networks loses its spatial homogeneity, even though the call arrival pattern is homogeneous. For any finite network, the distribution of calls near the boundary of the network will clearly be influenced by the boundary conditions (i.e. the distribution of calls at the edge of the network). However, as our results show, for small values of the multicast arrival rate, ν, the stationary distribution at nodes in the core of a large network is insensitive to the boundary conditions. In addition, for the core of sufficiently large networks, this unique stationary distribution approaches homogeneity, in the sense that the marginal distributions of calls at different nodes in the core network approach a common distribution. Higher values of ν, however, lead to a phase transition or symmetry breaking in the sense that the marginal call distribution at a core node now depends on whether the distance of the node from the boundary is odd or even. (Thus, for high ν, the particular distribution realized in the core of the network depends on the network boundary conditions.) This leads to *unfairness* in the sense that even in a symmetric network with uniform loading, calls arriving at even and odd nodes in the core may experience different blocking probabilities. This phenomenon is particularly relevant in networks with a higher degree of connectivity (i.e. larger q), which manifest this phase transition at lower arrival rates, possibly well below the capacity of the network. The presence of uncorrelated unicast calls in the network seems to somewhat mitigate this effect as it raises the multicast arrival rate at which the phase transition takes place (at least for $C = 1$). Thus, if an operator would like to avoid 'unfairness', it may be desirable to work away from the region of phase transition.

6.2. Networks with tree-like structure v. lattices

Markov random fields on a Cayley tree behave quite differently from those defined on regular lattices, and thus it is instructive to contrast the two structures. A bounded square lattice $Z_n \times Z_n$ has n^2 points, of which $4n-4$ lie on the boundary. A spherical tree S_L with coordination number q has $q^L + q^{L-1}$ boundary nodes and $(q^{L+1} + q^L - 2)/(q-1)$ nodes in total. Note that the ratio of the number of boundary nodes to the total number of nodes, in the limit as $L \to \infty$, converges to zero for the lattice, but stays finite in the case of the tree. In other words, in the thermodynamic limit, regular lattices have negligibly few boundary nodes compared to the total number of nodes, whereas the Cayley tree has a nonzero fraction of nodes on the boundary. As a result, the Cayley tree system is much more susceptible to boundary conditions than a regular lattice.

To make this more concrete, we consider the particular question of whether there exist symmetry broken (i.e. non-translationally invariant) stationary distributions when $C = 2$ and $\lambda = 0$. The results of Section 3.4 indicate that the answer is affirmative for Cayley trees, whereas the results of [11, p. 5] seem to suggest that the answer is negative for lattice models. We argue heuristically to provide some intuition as to why this may be the case. For both

graphs, we define the odd lattice to be the collection of V_1 nodes that lie an odd distance away from the terminal nodes and the even lattice to be the collection of V_2 nodes that lie an even distance from the terminal nodes. Then for large ν, it is reasonable to expect that the stationary distribution would be dominated by those configurations that sustain the largest total number of calls in the network, subject to the capacity constraint. Consider two candidate configurations: (i) *homogeneous*, with one call per site, and (ii) *symmetry broken*, two calls on one sublattice, and no calls on the other. Then the homogeneous configuration has a weight $\nu^{V_1+V_2}$, while the symmetry broken one (say with calls supported only on the even lattice) has a weight $(\frac{1}{2}\nu^2)^{V_2}$. Let $V = V_1 + V_2$ be the total number of nodes in the network. On the regular lattice, $V_1 \approx V_2 \approx V$, and thus both configurations have weights with the same power of ν. However, for all values of ν, the weight of the homogeneous configuration is greater because it has a larger prefactor. The same configurations on the spherical Cayley tree have very different weights. For a tree with $q = 3$ (which is equal to the coordination number of a regular square lattice), since V_2 includes the boundaries, it is easy to see that $V_2 \approx 3V_1 \approx \frac{3}{4}V$. Hence, for small ν the homogeneous configuration has larger weight, while for ν large enough such that $(\frac{1}{2}\nu^2)^{3/4} > \nu$, the symmetry broken configuration has greater weight. Consequently, we would expect the Cayley tree model (for $C = 2$) to show a phase transition for large enough ν.

Another difference between lattice and tree models was pointed out in Section 5.2, where it was shown that phase transitions can occur on the tree even when the arrival rates on the odd and even lattices differ. In contrast, this is not expected to happen in the case of the regular square lattice [19].

In general, when we analyse the recursion relations for the Cayley tree model, we often find fixed points of fixed cycles with particular basins of attraction (as illustrated, for example, in Figure 4). These basins of attraction correspond to initial conditions for the recursion relations, which in turn correspond to specific boundary conditions. The fact that we find more phases in the Cayley tree model than the regular lattice is not surprising. It is due to the strong influence of the boundary. In physics models of spin systems, this is often an artefact of the model. However, in actual telecommunication networks the connectivity structure is often tree-like, especially in the outlying or 'access' regions of the networks, and this may significantly affect the performance in the 'core' network. Hence a strong influence of the boundary on the core may be a welcome element in a network model.

6.3. Dependence on parity of network capacity C

An interesting by-product of our analysis is the dependence of the nature and value of the phase transition on the parity of the capacity C. The phase transitions seem to be fundamentally different for odd and even C—with odd C showing a second order transition and lower phase transition values, and even C showing a first order transition that occurs at higher values. The results of [11, p. 5] suggest that in lattice models too the nature of phase transitions depends on the parity of the capacity. However, in that case, for even C there is no phase transition, as there is a unique Gibbs distribution for all values of ν, while for odd C there is a phase transition leading to multiple Gibbs distributions for large enough arrival rates ν.

6.4. Problems with traditional approximations for multicast networks

Our analysis indicates that traditional techniques such as the EFP approximation for estimating link blocking probabilities may not be very effective in networks supporting multicast calls beyond the critical point of the EFP map. This is probably due to the fact that as we approach phase transition, the correlations introduced by multicasting seem to seriously undermine the independence assumption underpinning such approximations. It would thus be of interest to

develop simple and effective approximation methods that can be effectively applied to the multicast setting near and beyond the phase transition. Zachary and Ziedins [22] outline a general framework for refining the EFP approximation, which could be applied in the multicast setting for more general networks provided no controls, such as trunk reservation, are in force.

6.5. Implications for transient behaviour in finite networks

We discussed earlier how phase transitions give rise to qualitative differences in stationary measures in the core of large networks for different regions of the parameter space. This phenomenon of different limiting quantities for different parameter values often also applies to transient behaviour. Analysis of dynamic phenomena, even for such a simple model, is rather hard. However, the existence of a phase transition suggests that finite networks would exhibit transient behaviour that spends large amounts of time in these different phases and decays slowly to the equilibrium distribution. This kind of behaviour is reasonably well understood for lattice systems, and thus it would be worthwhile to carry out a similar analysis for tree networks.

6.6. Control of the network and future directions

It would be interesting to study the effect of controls akin to trunk reservation on the phase transition effects—in the past trunk reservation has been used to control bistable behaviour (see, for instance, [9], [13]), and we believe that it or a related control could be used here to similar effect. We would also like to consider more realistic settings in which the arrival rate at each node is not uniform (or picked from a set of two), but is instead picked from a distribution.

Acknowledgements

Ilze Ziedins thanks Bell Labs, Lucent Technologies, for its generous hospitality while preparing much of this paper. Thanks are also due to Stan Zachary for his comments.

References

[1] Ballardie, A. J., Francis, P. F. and Crowcroft, J. (1993). Core based trees. In *Proc. SIGCOMM '93, Annual Tech. Conf.*, Association for Computing Machinery, New York, pp. 85–95.

[2] Baxter, R. J. (1982). *Exactly Solved Models in Statistical Mechanics*. Academic Press, London.

[3] Brightwell, G. R. and Winkler, P. (1999). Graph homomorphisms and phase transitions. *J. Combinatorial Theory B* **77,** 415–435.

[4] Brightwell, G. R. and Winkler, P. (2000). Gibbs measures and dismantlable graphs. *J. Combinatorial Theory B* **78,** 141–169.

[5] Burman, D. Y., Lehoczky, J. P. and Lim, Y. (1984). Insensitivity of blocking probabilities in a circuit-switched network. *J. Appl. Prob.* **21,** 850–859.

[6] Chiu, D. M., Hurst, S., Kadansky, M. and Wesley, J. (1998). TRAM: A tree-based reliable multicast protocol. Tech. Rep. TR-98–66, Sun Microsystems.

[7] Devaney, R. L. (1989). *An Introduction to Chaotic Dynamical Systems*, 2nd edn. Addison-Wesley, Redwood City, CA.

[8] Georgii, H.-O. (1988). *Gibbs Measures and Phase Transitions*. Walter de Gruyter, Berlin.

[9] Gibbens, R. J., Hunt, P. J. and Kelly, F. P. (1990). Bistability in communication networks. In *Disorder in Physical Systems*, eds G. R. Grimmett and D. J. A. Welsh, Oxford University Press, pp. 113–127.

[10] Karvo, J., Virtamo, J., Aalto, S. and Martikainen, O. Blocking of dynamic multicast connections. *Telecommun. Systems* **16,** 467–481.

[11] Kelbert, M. Y. and Suhov, Y. M. (1990). Statistical physics and network theory. Unpublished manuscript.

[12] Kelly, F. P. (1985). Stochastic models of computer communication systems. *J. R. Statist. Soc. B* **47,** 379–395.

[13] Kelly, F. P. (1991). Loss networks. *Ann. App. Prob.* **1,** 319–378.

[14] Louth, G. M. (1990). Stochastic networks: complexity, dependence and routing. Doctoral Thesis, University of Cambridge.

[15] MARTIN, P. P. (1991). *Potts Models and Related Problems in Statistical Mechanics* (Ser. Adv. Statist. Mechanics **5**). World Scientific, Teaneck, NJ.
[16] PAUL, S., SABNANI, K. K., LIN, J. C. AND BHATTACHARYA, S. (1997). Reliable multicast transport protocol (RMTP). *IEEE J. Sel. Commun.* **15,** 407–421.
[17] ROSS, K. W. (1995). *Multiservice Loss Models for Broadband Telecommunications Networks*. Springer, Berlin.
[18] SPITZER, F. (1975). Markov random fields on an infinite tree. *Ann. Prob.* **3,** 387–398.
[19] VAN DEN BERG, J. AND STEIF, J. E. (1995). Percolation and the hard-core lattice gas model. *Stoch. Process. Appl.* **49,** 179–197.
[20] ZACHARY, S. (1983). Countable state space Markov random fields and Markov chains on trees. *Ann. Prob.* **11,** 894–903.
[21] ZACHARY, S. (1985). Bounded, attractive and repulsive Markov specifications on trees and on the one-dimensional lattice. *Stoch. Proc. Appl.* **20,** 247–256.
[22] ZACHARY, S. AND ZIEDINS, I. (1999). Loss networks and Markov random fields. *J. Appl. Prob.* **36,** 403–414.
[23] ZEIDLER, E. (1985). *Nonlinear Functional Analysis and Its Applications I.* Springer, Berlin.