Multicast Scaling Properties in Massively Dense Ad Hoc Networks

Philippe Jacquet, Georgios Rodolakis
INRIA and Ecole Polytechnique
France
philippe.jacquet@inria.fr, georges.rodolakis@inria.fr

Abstract

We study the benefits of multicast routing in the performance of mobile ad hoc networks. In particular we show that if a node wishes to communicate with n distinct destinations, multicast can reduce the overall network load by a factor $O\left(\sqrt{n}\right)$, when used instead of unicast. One of the implications of this scaling property consists in a significant increase of the total capacity of the network for data delivery. We discuss how these results can be taken into consideration in the operation of a multicast protocol using overlay multicast trees.

1. Introduction

Multicast routing constitutes an important challenge for ad hoc networks, mainly due to the dynamic nature of the routes in the network. In the research literature, many possible approaches have been proposed [7]. However, although the capacity of mobile ad hoc networks has been a very active research area since the seminal paper of Gupta and Kumar [9], the specific impact of multicast routing has not attracted too much attention, with the exception of [13].

One of the advantages of multicast routing is that it reduces the total bandwidth required to communicate with all group destinations, since some links can be common to several destinations. In wired networks, the gain of multicast communication has been studied in [1, 5, 11], by estimating the ratio of the number of links in a multicast tree to n destinations over the average unicast hop distance between two random nodes. The resulting normalized multicast cost has been found experimentally to scale in $n^{0.8}$. The gain of multicast is reflected by how far the normalized multicast cost deviates from linear growth. However, the topology of mobile ad hoc networks is significantly different and one would expect a much different scaling law. Indeed the average unicast hop distance in wired networks is usually of the order $\log N$, where N is the total number of nodes in the network, while in ad hoc networks the average distance grows proportionally to $\sqrt{N/\log N}$, since the optimal neighbor degree increases in $O(\log N)$ when the capacity increases.

In this paper, we establish performance bounds on the expected size of multicast trees as a function of the number of multicast destinations n, both via analytical methods and via simulation. In random mobile ad hoc networks, the gain of multicast communication compared to unicast is significantly larger than in wired networks. For instance, we show that a scaling law in $O(\sqrt{n})$ holds for the normalized multicast cost and we provide a protocol outline, to be used in conjunction with the unicast routing protocol OLSR [6]. This result can provide further motivation in supporting multicasting in mobile networks, besides the advantages of group-oriented communication. The implications of this scaling law consist in a significant increase of the total capacity of the network for data delivery, while the total amount of generated data will actually decrease (compared to the case where each node transmits data to one single destination), and both are proportional to $\sqrt{\frac{N}{\log N}}$.

The remainder of this paper is organized as follows. In Section 2 we present the network model and provide analytical results on the scaling law of the normalized multicast cost. The impact of multicasting in the capacity of the network is discussed in Section 3. In Section 4 we outline the basic operation of a multicast protocol which is based on previous observations. In Section 5 we present measurements on multicast scaling derived from simulations in generated graph models of wireless ad hoc networks. We conclude with some issues for further research in Section 6.

2. Multicast cost scaling law

In this section we will quantify the cost of multicast communication versus the average unicast cost. We assume that nodes have a complete knowledge of the network topology. In order to optimize the control traffic we will see in a further section how we can somewhat relax this hypothesis in the use of the OLSR link state routing protocol.

2.1. Model description

We assume that N nodes forming a massively dense ad hoc network are distributed according to a Poisson process in an area of arbitrary size $\mathcal{A}$. In this case, *i.e.*, when N is large, routes can be considered as continuous lines between nodes, and the number of retransmissions needed for a packet to reach its destination is $\Theta\left(\frac{d}{r}\right)$, where r is the typical radio range and d is the Euclidean distance from the source to the destination [10, 3]. Hence, we can represent a massively dense ad hoc network with an Euclidean graph, in which the edge costs are proportional to hop distances between nodes.

The result of Gupta and Kumar [9] states that the maximum bandwidth is attained when the radio range is $r = \alpha\sqrt{\frac{\log N}{N}}$, where α is a constant which depends on signal propagation and medium access control.

A source and a multicast group of size n are chosen uniformly at random among the N nodes. As a result, the $n+1$ multicast nodes are distributed in the area according to a Poisson process of intensity $\frac{n+1}{\mathcal{A}}$.

An optimal multicast tree is a Steiner tree, *i.e.*, a tree of minimal cost connecting all of the multicast nodes via an arbitrary subset of the remaining nodes that are not in the multicast group. Therefore, the problem of finding the optimal tree is NP-complete, even in Euclidean graphs, although in this case there is a polynomial time approximation scheme [2].

Since the problem of finding the optimal tree is intractable, we will use an approximation. We consider the two more common cases of minimum spanning trees and shortest path trees.

2.2. Minimum spanning trees

First, we consider multicast trees corresponding to minimum spanning trees on the $n+1$ multicast nodes. In a minimum spanning tree branching is constrained only on multicast nodes, and the computation can be performed in polynomial time. On the other hand, in Steiner trees, branching can occur on any node (or any point in the plane in the Euclidean case).

In general graphs, the cost of a minimum spanning tree is within twice the cost of an optimal Steiner tree [14]. However, it can be shown that the Euclidean minimum spanning tree is not longer than $\frac{2}{\sqrt{3}}$ times the optimal Euclidean Steiner tree [8]. Hence, in the case of massively dense ad hoc networks, minimum spanning trees yield results which are very close to the optimal.

To proceed we will compute the expected path length (in meters) of a minimum spanning tree on $n+1$ points in an area $\mathcal{A}$. We denote this length $L(n+1)$.

From the analysis in [4, 12], in the $2-$dimensional case, it comes that an upper bound for the average path length (in meters) of a minimum spanning tree is

$$L(n+1) \leq \gamma n\sqrt{\frac{\mathcal{A}}{n+1}} \sim \gamma\sqrt{\mathcal{A}n}. \quad (1)$$

where γ is a constant that depends on the shape of the network domain. For a disk or a square we have $\gamma = \frac{1}{\sqrt{2}}$.

We now define the normalized multicast cost $R(n)$ for a multicast group of size n as

$$R(n) = \frac{\text{multicast cost}}{\text{average unicast cost}},$$

where the costs are expressed in number of hops. In other words the multicast cost is the number of links in the multicast tree, and the average unicast cost is the average route length from a random source in the multicast group to a random destination in the multicast group.

We base our analysis on the observation that routes can be considered as continuous lines between nodes, and the number of hops needed for a packet to reach its destination is $\Theta\left(\frac{d}{r}\right)$, where r is the optimal radio range as stated by Gupta and Kumar.

The expected path length of the multicast tree in number of hops is $\Theta\left(\frac{L(n+1)}{r}\right)$, while the average unicast cost is $\Theta\left(\frac{L(2)}{r}\right)$. This implies that, for the normalized multicast cost, it holds

$$R(n) \simeq \frac{L(n+1)}{L(2)}. \quad (2)$$

Quantity $L(2)$ is highly dependent on the shape of the network domain, but is of order $\sqrt{\mathcal{A}}$ when the network domain shape stays within some reasonable model. In all rigor we have $L(2) = \beta\sqrt{\mathcal{A}}$. For the disk we have $\beta = \frac{128}{45\pi^{3/2}} \approx 0.51$, for the square it becomes $\beta = \frac{1}{15}(2+\sqrt{2}+5\log(1+\sqrt{2})) \approx 0.52$.

Combining (1), (2), we get

$$R(n) \leq \frac{\gamma n}{\beta\sqrt{n+1}} = O(\sqrt{n}). \quad (3)$$

Hence, we obtain the multicast scaling law $R(n) = O(\sqrt{n})$. It comes that the gain of multicast over unicast, which is reflected by how far $R(n)$ deviates from linear growth, is also $O(\sqrt{n})$. This result is in contrast with similar comparisons in wired networks [1, 5, 11] where the gain of multicast communication is significantly smaller. In that case, the multicast cost scales, according to experimental studies, following a power law with exponent between $.8$ and $.9$.

More generally, we can show, using the same approach, that for a network spanning on a domain in dimension D

$$R(n) = O\left(n^{1-\frac{1}{D}}\right).$$

In [12] it is shown that the length of minimum spanning trees on points randomly placed in a hypercube is $O\left(n^{1-\frac{1}{D}}\right)$, even when the point distribution is not uniform, with some mild constraints. This implies that the multicast scaling law still holds when the multicast nodes are not distributed uniformly among the nodes of the network.

2.3. Shortest path trees

A popular approach in building multicast trees in wired networks consists in pruning shortest path trees rooted at the source node. In this case, we cannot prove worst case bounds on the total cost of shortest paths trees, compared to the cost of optimal Steiner trees. However, in practice, shortest path trees achieve a satisfactory performance. Moreover, shortest path trees minimize the maximum path length from the source to any destination.

In the currently considered model of mobile ad hoc networks, when $n \ll N$, a shortest path tree is equivalent to n unicasts, since the expected number of branching nodes is very small. Hence for a small number of destination nodes, the gain of multicast communication is negligible.

On the other hand, when $n \to N$, the total number of hops in the tree also tends to N, since we consider a tree spanning on almost all the nodes. The average unicast distance in hops is $O\left(\frac{1}{r}\right)$, where the radio range $r = \alpha\sqrt{\frac{\log N}{N}}$. Hence, the normalized multicast cost tends to

$$R(N) = O(Nr) = O(\sqrt{N \log N}).$$

This is the expected behavior for any method used to construct a tree spanning on all the nodes of an ad hoc network with unit cost links.

Consequently, for large multicast group sizes we still expect to observe a multicast scaling law of the form $R(n) = O(n^{\frac{1}{2}+\varepsilon})$, for any $\varepsilon > 0$. In Section 5, we study the normalized multicast cost of shortest path trees experimentally.

3. Capacity of multicast communication

In this section we investigate the impact of the multicast cost scaling law in the capacity of the network, when all nodes communicate with multicast. We are interested in the order of magnitude of the maximum attainable bandwidth. We show that similar bounds to the ones described in [13] can be obtained without the need of particularly complex additional routing mechanisms.

In presence of traffic density of λ bits per time unit per square area unit, the typical radius of correct reception r decays in $O\left(\frac{1}{\sqrt{\lambda}}\right)$ [9, 10]. If C is the capacity generated by each node, the density of traffic generated per square unit area is $O(CN)$. The maximum bandwidth attainable for unicast traffic is $C = O\left(\frac{1}{\sqrt{N \log N}}\right)$.

We have shown that each multicast packet in a group of size n will be retransmitted $O(\sqrt{n}\frac{1}{r})$ times. This yields a traffic density (including retransmissions) $\lambda = O(CN\sqrt{n}\frac{1}{r})$. Therefore

$$r = O\left(\frac{1}{\sqrt{\lambda}}\right) = O\left(\sqrt{\frac{r}{CN\sqrt{n}}}\right)$$
$$\Rightarrow C = O\left(\frac{1}{rN\sqrt{n}}\right).$$

As a result, the maximum rate at which a node can transmit multicast data is $O\left(\frac{1}{\sqrt{nN \log N}}\right)$ and it is obtained for $r = O\left(\sqrt{\frac{\log N}{N}}\right)$. In this case, the total rate at which data is received by the n destinations in the multicast group is $O\left(\sqrt{\frac{n}{N \log N}}\right)$.

When all nodes in the network communicate in unicast (each node with one single destination), the total capacity of the network increases with network size in $O\left(\sqrt{\frac{N}{\log N}}\right)$. However, when there are $O(N)$ nodes in the network acting as multicast sources in groups of size n (e.g. in teleconferences between n users), the total rate at which data is transmitted in the network is $O\left(\sqrt{\frac{N}{n \log N}}\right)$. Similarly, the total rate at which data is received is $O\left(\sqrt{\frac{nN}{\log N}}\right)$. Hence, compared to unicast traffic, multicast traffic results in an increase by a factor $O(\sqrt{n})$ of the capacity of the network (and per node) for receiving data, although the total distinct data transmitted will in fact decrease by the same factor.

4. Protocol overview

In this section, we outline the basic operation of a multicast protocol, taking into consideration the previously derived results. However, a detailed protocol implementation is outside the scope of this paper.

4.1. Overlay tree construction

As we saw previously, it is more efficient to consider minimum spanning trees for the construction of multicast trees. We discuss here an algorithm which achieves the multicast cost $R(n) = O(\sqrt{n})$ calculated in Section 2. The algorithm does not require any more information than what is provided by a link state unicast routing protocol, like OLSR [6].

Algorithm 1 *(input: network graph, output: overlay multicast tree)*

1. Find shortest paths between all pairs of multicast nodes.

2. Build complete graph on multicast nodes with costs $c_{ij} = \{$length of shortest path between i and $j\}$.
3. Build minimum spanning tree on the complete graph, rooted at the source node.

The construction of the minimum spanning tree (step 3) can be implemented using Prim's algorithm. The resulting tree is an overlay multicast tree, since it consists only of multicast nodes and its links are in fact tunnels in the actual network. Multicasting is achieved when each node forwards multicast packets to its successors in the overlay tree.

Step 1 corresponds to n Dijkstra algorithm iterations. Therefore, the total complexity is $O(n(M + N \log N))$, where n is the multicast group size, N and M are the number of nodes and edges in the network, respectively. The algorithm's expected complexity can be improved because it is not necessary in practice to compute all shortest paths from each node to all other nodes to build the minimum spanning tree.

Due to the fact that the distance between two nodes in number of hops is proportional to the Euclidean distance, the resulting multicast tree can be considered as an approximation of the Euclidean minimum spanning tree on the multicast nodes.

The advantage of this algorithm is that only the multicast nodes need to participate in the construction of the multicast tree, while the other nodes serve merely as relays and are not necessarily aware of the multicast communication. This fact facilitates the development of a peer to peer protocol which can be run only by the participating multicast nodes, hence it could be downloaded dynamically by a node whenever it decides to join a multicast communication.

However, a practical implementation of a protocol using this approach must still face the difficulties introduced by the nodes' mobility. The protocol should provide a mechanism for nodes to communicate reliably among them their participation in the multicast communication.

4.2. Implementation on a link state protocol

The Optimized Link State Routing protocol is using topology control optimization. Broadcast traffic is relayed via Multi-Point Relay (MPR) nodes. MPR nodes are elected by their neighbors because they cover their two-hop neighborhood. That way broadcast traffic consumes less resources in order to be forwarded to all destinations. In order to save more on control traffic, nodes have the possibility to advertize a small subset of their neighbor links. The advertized link set can be limited to MPR links, *i.e.*, the neighbors that have elected this node as an MPR. In this case the nodes have only a partial knowledge of the network topology. The fact that any given node can compute a shortest path to any arbitrary destination comes from the fact that the node knows its own neighbor list. In order to allow all multicast group members to compute a shortest path between any other pair of nodes in the multicast group means that multicast group members must set their advertized link set to their whole neighbor list. This behavior is an option in OLSR.

Therefore, the nodes can compute the minimum spanning overlay tree by using their local information. The multicast tree computation leads to the determination of the overlay neighbors of each node, *i.e.*, the nodes in the multicast group to which the node is connected via tunnel in the multicast tree. Upon reception of a multicast packet, the node transmits one copy to each of its overlay neighbors via IP tunnel over OLSR (except towards the node from which it has received the multicast packet).

To proceed to the correct computation of the overlay tree, multicast nodes need to know the membership of their multicast group. The membership knowledge can be maintained via a periodic Standby message transmitted via the multicast tree or by a Join message periodically transmitted via the optimized broadcast transport protocol of OLSR. When a node leaves the multicast group it can advertize it by emitting a Leave packet on the multicast tree, so that the multicast tree can be updated. If a mutlicast node leaves the network, this will be detected by OLSR, leading also to an update of the multicast tree.

Once the overlay tree is computed, a node can communicate with its logical neighbors using shortest paths, which can be continuously updated by the link state protocol. Hence we obtain a protocol which is robust with respect to mobility. Satisfactory performance can be ensured by performing the multicast tree construction algorithm at an appropriate frequency. When the overlay tree is modified due to mobility or to membership changes, it may be safe for a member to keep during a short time both the old and the new overlay neighbors.

5. Simulations

In this section, we present simulations that verify the theoretical results on generated graph models of mobile ad hoc networks. We measure the normalized multicast cost $R(n)$ for various sources and multicast groups, and take the average for each group size n. The results are plotted in $\log \log$ scale, and compared to a line which corresponds to the predicted asymptotic growth.

The graphs are generated by placing nodes randomly in a square for $2-\mathrm{D}$ networks and in a cube for $3-\mathrm{D}$ networks, and then connecting the nodes which are in a distance smaller than the critical radius for connectivity r, such that the average number of neighbors for each node is $\log N$.

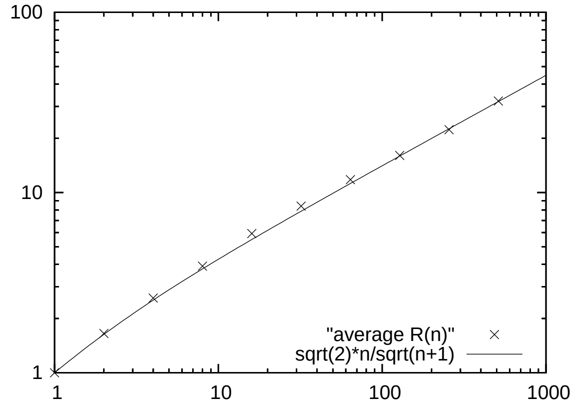

(a) 2-D network with 1500 nodes

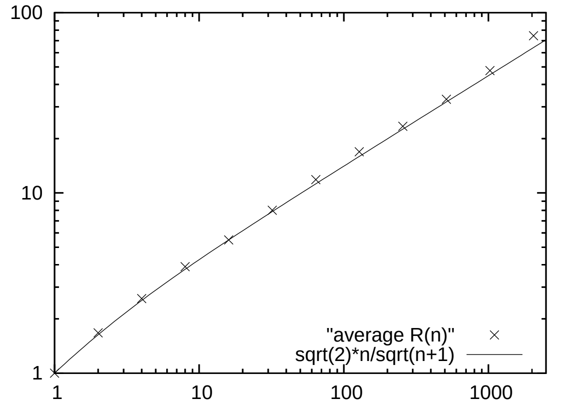

(b) 2-D network with 4500 nodes

Figure 1. Normalized multicast cost $R(n)$ versus multicast group size n, for minimum spanning trees.

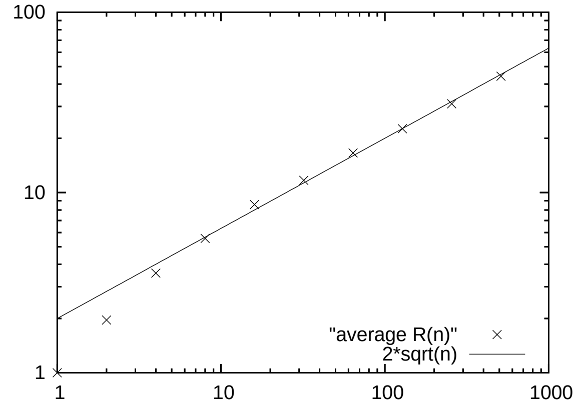

(a) 2-D network with 1500 nodes

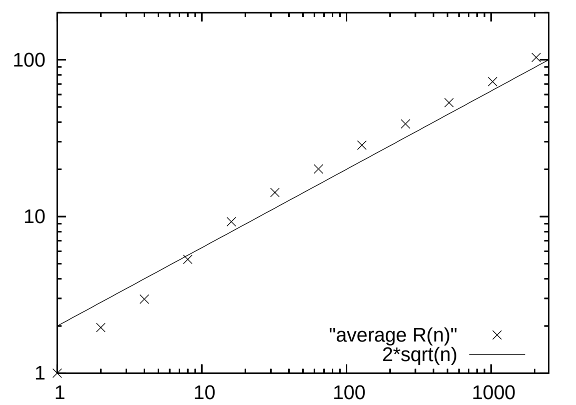

(b) 2-D network with 4500 nodes

Figure 2. Normalized multicast cost $R(n)$ versus multicast group size n, for shortest path trees.

In figure 1, we present results corresponding to minimum spanning trees, constructed according to the algorithm of the previous section. The measured cost is compared to the function $\frac{n\sqrt{2}}{\sqrt{n+1}}$, obtained from equation (3) by setting $\gamma = \frac{1}{\sqrt{2}}$ and $\beta = 1/2$, which corresponds to the approximate case where we do not consider border effects.

Figure 2 depicts measurements of the normalized cost of shortest path trees. The multicast cost $R(n)$ is compared to function $2\sqrt{n}$ (where the constant 2 was chosen empirically). Observe that the cost is always higher than in minimum spanning trees, although the plot grows linearly with a slope close to 0.5 for large n.

In figure 3, we present simulation results in the case of a $3-$dimensional ad hoc network for both minimum spanning trees and shortest path trees. In this case, the normalized multicast cost scales in $O(n^{\frac{2}{3}})$.

6. Conclusion and future work

In this work, we established performance estimates for multicast routing versus unicast in massively dense ad hoc networks. We showed that multicasting can reduce the overall network load by a factor $O\left(\sqrt{n}\right)$, for n multicast group members. Consequently, the total capacity of the network for data delivery is significantly increased. Although we used geometric arguments to justify this behavior analytically, we also proposed a protocol outline which uses only the information provided by a link state routing protocol. Interesting directions for further work consist in a thorough investigation of the issues related with the efficient operation of the overlay tree based protocol in practice, and in a complete implementation on top of OLSR.

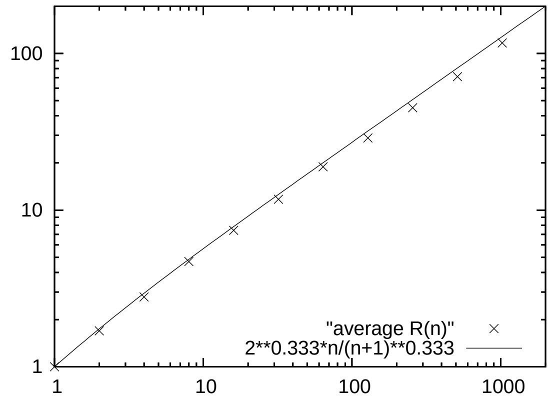

(a) Minimum spanning trees

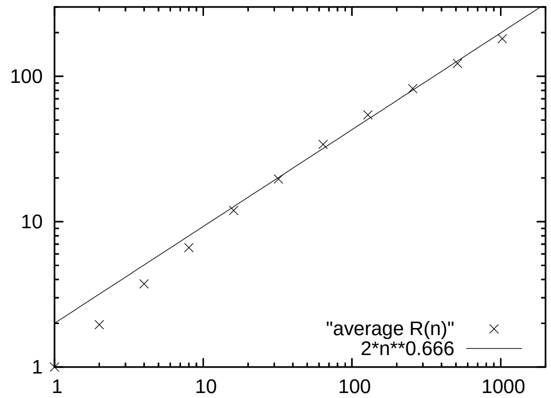

(b) Shortest path trees

Figure 3. Normalized multicast cost $R(n)$ versus multicast group size n, in a 3-D network with 3000 nodes.

References

[1] C. Adjih, L. Georgiadis, P. Jacquet, and W. Szpankowski. Is the internet fractal: The multicast power law revisited. In *SODA*, 2002.

[2] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.

[3] F. Baccelli, K. Tchoumatchenko, and S. Zuyev. Markov paths on the Poisson-Delaunay graph with applications to routing in mobile networks. *Adv. Appl. Probab.*, 32(1):1–18, 2000.

[4] D. Bertsimas and G. V. Ryzin. An asymptotic determination of the minimum spanning tree and minimum matching constants in geometrical probability. *Operations Research Letters*, 9:223–231, 1990.

[5] J. C.-I. Chuang and M. A. Sirbu. Pricing multicast communication: A cost-based approach. *Telecommunication Systems*, 17(3):281–297, 2001.

[6] T. Clausen, P. J. (editors), C. Adjih, A. Laouiti, P. Minet, P. Muhlethaler, A. Qayyum, and L.Viennot. Optimized link state routing protocol (olsr). RFC 3626, October 2003. Network Working Group.

[7] C. Cordeiro, H. Gossain, and D. Agrawal. Multicast over wireless mobile ad hoc networks: present and future directions, 2003.

[8] D. Z. Du and F. K. Hwang. A proof of gilbert-pollak's conjecture on the steiner ratio. *Algorithmica*, 45:121–135, 1992.

[9] P. Gupta and P. R. Kumar. Capacity of wireless networks. *IEEE Transactions on Information Theory*, IT-46(2):388–404, 2000.

[10] P. Jacquet. Geometry of information propagation in massively dense ad hoc networks. In *MobiHoc*, 2004.

[11] G. Phillips, S. Shenker, and H. Tangmunarunkit. Scaling of multicast trees: Comments on the chuang-sirbu scaling law. In *SIGCOMM*, pages 41–51, 1999.

[12] M. Steele. Growth rates of euclidean minimal spanning trees with power weighted edges. *Annals of Probability*, 16:1767–1787, 1988.

[13] S. Toumpis and A. J. Goldsmith. Performance bounds for large wireless networks with mobile nodes and multicast traffic. In *Inter. Workshop on Wireless Ad Hoc Networks, Oulu, Finland*, 2004.

[14] V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.