

Low Latency Multimedia Broadcast in Multi-Rate Wireless Meshes

Chun Tung Chou¹, Archan Misra²

¹School of Computer Science and Engineering, University of New South Wales, Australia

²IBM T J Watson Research Center, Hawthorne, New York, USA

Email: ctchou@cse.unsw.edu.au, archan@us.ibm.com

Abstract—In a multi-rate wireless network, a node can dynamically adjust its link transmission rate by switching between different modulation schemes. For the current IEEE802.11a/b/g standards, this rate adjustment is limited to unicast traffic only while multicast and broadcast traffic is always transmitted at the lowest possible rate. In this paper, we consider a novel type of multi-rate mesh networks where a node can dynamically adjust its link layer multicast rates to its neighbours. In particular, we consider the problem of realising low latency network-wide broadcast in this type of multi-rate wireless meshes. We will first show that the multi-rate broadcast problem is significantly different from the single-rate case. We will then present an algorithm for achieving low latency broadcast in a multi-rate mesh which exploits both wireless broadcast advantage and the multi-rate nature of the network.

I. INTRODUCTION

Wireless mesh networks are attracting significant research and commercial interest, especially as suburban and urban community-based networks. In such environments, the wireless mesh nodes act as both static *relays*, forwarding traffic to or from other mesh nodes, and *access points* providing localized first-hop connectivity to mobile or pervasive wireless devices, such as laptops and PDAs. A popular use of such a wireless mesh is to extend the footprint of wide-area connectivity to a larger community, by using the multi-hop wireless mesh to funnel traffic from an extended area to a much smaller set of gateway nodes, that connect to the Internet backbone over a high-speed wired (e.g., DSL/cable) connection.

One of the fundamental problems of existing multi-hop wireless network solutions is the sharp drop in multi-hop throughput (to as low as a few Kbps), even though individual wireless links evolve to higher speeds (such as 54Mbps or 108Mbps). To remedy this, two themes of wireless mesh research are especially popular:

- a) Use of Multi-Channel, Multi-Radio Mesh Nodes: The use of multiple radios on a single node, each tuned to possibly distinct channels, can significantly improve the spatial reuse of an individual channel, and result in higher overall capacity, by increasing the degree of concurrent transmissions in the network.
- b) Multi-rate MAC Protocols: Researchers are finally looking beyond 802.11-based single channel MACs for wireless meshes, and studying the throughput and fairness issues that arise from multi-rate MAC protocols, where adaptive modulation is used to dynamically modify the

data rate on a particular link in response to the perceived signal-to-noise (SNR) ratio.

Most of the recent research effort has, however, focused on the *unicast* traffic scenario, where each traffic flow is defined between a particular node pair. For example, [8] demonstrates the use of new unicast routing metrics for multi-channel multi-rate mesh environments, that account for both the intra-flow contention and the differential transmission rate on different links.

In this paper, we introduce the problem of efficient routing and packet distribution for *broadcast* (any by extension, multicast) traffic flows¹ in such multi-rate, multi-channel, multi-radio wireless meshes. Our primary goal is to show that the presence of multiple radios, or multi-rate adaptive modulation schemes, opens up new possibilities for broadcast traffic distribution that do not seem to have been explored before. Indeed, metrics and routing strategies defined for unicast traffic scenarios do not capture the interesting effects that broadcast traffic introduce in a wireless mesh.

We believe that the development of techniques for high-throughput, low-latency forwarding of multicast traffic is important for many of the collaborative/communal applications likely to be enabled by community wireless mesh networks. For example, wireless meshes may be used to broadcast community-specific content (such as a video feed of a neighborhood soccer game or video feeds from multiple video sensors) or even wide-area content (such as TV feeds received at a particular gateway node) to a group of receiver nodes. While routing algorithms for multicasting traffic have been well studied in multi-hop wireless networks, their focus has been largely limited to the efficient dissemination of *control* traffic, rather than the support of high bit-rate multimedia “content”. For example, routing techniques (e.g., [5]) to avoid the broadcast storm problem [14] were motivated by a desire to limit the impact of route-discovery packets broadcast by many popular reactive ad-hoc routing protocols. There appears to be little work on the impact of such multicast techniques on the achievable throughput or latency bounds. For our target applications, such as broadcasting camera feeds or transporting

¹For reasons of space, we focus purely on the broadcast problem, where a source node distributes a packet to all other mesh nodes. In general, the techniques and algorithms used for broadcasting can be applied, with minor modification, to multicast scenarios, where the traffic is intended for only a subset of the mesh.

peer-to-peer multiplayer game traffic over wireless meshes, bounding the packet distribution latency (without causing unnecessary use of channel capacity), however, is of critical importance.

Efficient algorithms for broadcasting data in multi-hop wireless networks exploit the natural *wireless broadcast advantage* [15], whereby a single transmitting node can reach multiple one-hop neighboring nodes with a single transmission. Most work on broadcast in MANETs has focused on *energy-efficiency*, and aims to reduce the number of distinct transmissions needed to reach the entire set of receivers. Examples of such energy-efficient broadcasting algorithms include the BIP algorithm [15] for incremental construction of a broadcast tree and the EWMA algorithm [4]. Energy efficiency may be a critical metric in mobile ad-hoc networks (MANETs), but is of less concern in many mesh environments, where the nodes are relatively static (e.g., mounted on rooftops) and directly connected to regular power outlets. Our research effort instead focuses on developing low-latency, high-throughput multi-hop wireless packet broadcast algorithms, that may be leveraged by high-bit rate or interactive multicast multimedia streams.

A. The Main Questions With Broadcasting in Multi-Rate, Multi-Channel Meshes

We can formalize the issues with high-performance multi-hop broadcast by first defining a new metric of interest: *the broadcast latency*, computed as the maximum delay between the transmission of a packet by a source node and its eventual reception (over multi-hop paths) by all the intended receivers. In many cases, our goal is to minimize this worst-case path latency, since this not only indirectly appeals to a notion of more efficient packet delivery, but also translates into lower latency variation among the receivers. Constraining such latency variation may be especially important in interactive environments (e.g., to preserve temporal fairness among players in interactive multi-player games). Given such a metric, our overall research effort addresses the following questions:

- a) **Effect of Multi-Rate Links on Efficient Broadcasting:** Is it true that multicast-tree based distribution techniques outperform unicast-based strategies for broadcast traffic in such multi-rate meshes? Or, can one do better by using alternative packet distribution topologies and algorithms? Is there a benefit of allowing link-layer multicast to operate at different rates, and if so, how does one modify practical tree-based routing protocols to better exploit such rate diversity?
- b) **Sensitivity of Broadcast Topology to Traffic Generation Rate** Since we are no longer confined to low bit-rate, sporadic control traffic, does the variation in the source traffic generation rate affect the nature or topology of efficient packet broadcasting techniques? How does the choice of the broadcast distribution topology depend on the existing “traffic load” on individual nodes or links?
- c) **Effect of Multiple Radios and Channels on Efficient Broadcasting** How do we modify the multicast routing protocols to exploit the existence of multiple channels

or radios on each node? What are the appropriate routing metrics for multicast/broadcast traffic, and do channel assignment strategies need to be modified to better support multicast flows?

- d) **Architectures for Efficient Multicast Route Establishment** Unlike MANETS, mesh networks have relatively stable topologies. In such a scenario, can effective centralized or quasi-distributed route establishment protocols and architectures be designed for multicast flows? How are computed source-specific routes, and/or scheduling strategies, communicated to the individual nodes?

B. Contributions of This Paper

Given space limitations and the ongoing nature of our research, we shall principally tackle only the first question (effect of multi-rate links on broadcasting topologies) in this paper. Accordingly, the analytical and numerical results presented in this paper are restricted to a mesh network, where each node has a single radio, with all radios tuned to a common channel. However, due to adaptive modulation, the link data rate between a particular node pair varies based on the link distance, or more accurately, the SNR characteristics over this link. Our principal contributions include:

- 1) Demonstrating that the broadcast latency is not necessarily minimized by tree-based packet distribution topologies, where each intermediate node uses a single link-layer broadcast to reach its entire set of child nodes. Rather, optimal or efficient packet broadcasting is often achieved by having an intermediate node perform *multiple broadcasts*, each of which is directed towards a different subset of child nodes.
- 2) Demonstrating that the optimal traffic distribution topology, and the best broadcasting strategy, can be highly sensitive to the bit (or packet generation) rate of the broadcast flow. This suggests that effective broadcasting protocols in wireless meshes will need to have the expected bit rate of the application flow as an explicit parameter.
- 3) Designing modifications to existing wireless broadcast algorithms, such that they exploit the wireless broadcast advantage, as well as the multi-rate nature of individual links.

To our knowledge, our work is the first to illustrate the sub-optimality of a strategy that solely exploits the wireless broadcast advantage by sending a packet to all child nodes in a single transmission, and to present heuristic wireless broadcasting techniques that incorporate the multi-rate nature of the wireless links.

II. IMPACT OF MULTI-RATE LINKS ON EFFICIENT BROADCASTING

Effective packet broadcasting in a multi-rate multi-hop wireless environment depends strongly on the interaction between the routing and MAC layers. Intuitively, a pure flooding strategy, where each intermediate node re-broadcasts a received

packet, might be most robust, but can lead to significantly high broadcast latency, as the high number of redundant transmissions at the MAC layer lead to contention-induced backoffs (a.k.a, the broadcast storm problem). Accordingly, efficient broadcast strategies typically aim to build a distribution tree (or sometimes a “mesh” [11] for robustness), where redundant transmissions are eliminated or minimized. Given such a distribution tree, the simple strategy of treating each link in the forwarding tree as distinct, and thus having each intermediate node forward a packet to each of its downstream neighbors via individual unicasts, is also wasteful. By failing to exploit the wireless broadcast advantage, the all-unicast approach not only maximizes the forwarding latency at each intermediate node, but can induce additional backoff-based delay at the MAC layer due to the increased number of distinct transmissions at each node. Based on these observations, the natural solution implicit in most multicast routing protocols is that each intermediate node will *transmit its packet only once*, reaching all of its immediate downstream neighbors in a single link-level broadcast.

We first attempt to show how this central premise (i.e., that each intermediate node reaches all its neighbors in a single broadcast transmission) can lead to sub-optimal behavior in multi-rate wireless mesh environments. We implicitly assume that the MAC layer of future wireless meshes will provide some form of multicast support, where the transmitter may be able to specify the transmission rate of the MAC-layer broadcast, and either explicitly or implicitly (based on the range within such a broadcast is correctly received) the recipients of the broadcast. At present, technologies such as 802.11a or 802.11b do not offer this ability and mandate that all broadcasts proceed at the lowest possible rate (e.g., 1 Mbps for 802.11b and 6 Mbps for 802.11a), so that they may effectively reach the largest possible set of one-hop neighbors. Clearly, future MAC protocols may permit more flexibility. For example, relatively simple techniques have been proposed (e.g., [6]) to support such selective-broadcast at the wireless link layer, while the IEEE 802.16a group is considering the support of multicast traffic at the MAC layer.

A. Illustrating the Role of Differential Link Rates on the Broadcast Latency

To understand the closely coupled nature of the broadcast tree formation and the MAC layer scheduling, consider the topology shown in Figure 1 with five nodes, labelled as Nodes 1 to 5, arranged in a straight line. For simplicity, we will refer Node 1 as N_1 etc. in the text. In Figure 1, the d value between 2 nodes indicates the physical distance in meters between them. We assume each node is equipped with an IEEE 802.11b radio tuned to the same channel. By using the Qualnet simulator [13] as a reference, we find the transmission rate versus transmission range relationship in Table I assuming a free space propagation model. Note also that the interference range in Qualnet is 520m. Thus, given the network configuration in Figure 1, there are 4 links in the network. Link (1,2) has a maximum capacity of 11Mbps while




Fig. 1. Motivating example for the multi-rate network-wide broadcast problem.

Transmission rate (Mbps)	Maximum transmission range (m)
1.0	483
2.0	370
5.5	351
11.0	283

TABLE I

THIS TABLE SHOWS THE MAXIMUM TRANSMISSION RANGE IN METERS FOR DIFFERENT IEEE802.11B TRANSMISSION RATES. THE RANGES ARE OBTAINED FROM QUALNET [13] ASSUMING A FREE SPACE MODEL.

the other three links have a maximum capacity of 1 Mbps. Since our concern is packet delivery latency, we indicate the relative time required to send a packet for each link using the t value indicated in the Figure.

We assume that N_1 (i.e. Node 1) is the source node and it wants to send a packet to all the nodes in the network. Since the network is not fully connected, some nodes will need to act as a relay. We consider two different forwarding alternatives. In the first approach, which we call Alt_1 , each node is only allowed to broadcast the packet once. Due to this restriction, N_1 (the source node) must broadcast at the lower rate of 1Mbps to both N_2 and N_5 , taking a time of 11 units to transmit the packet. Note that N_1 could not possibly use other transmission rates because N_5 will not receive the packet otherwise. This results in the transmission schedule depicted in Figure 2, and leads to a broadcast latency of 33 time units.

In the second approach, which we call Alt_2 , we allow each node to broadcast the same packet more than once. Figure 3 depicts the transmission schedule. It shows the source N_1 transmitting the same packet two times. It first transmits to N_2 at 11Mbps (at time $t = 0$), taking 1 time unit. It then transmits the same packet again at time $t = 12$ to N_5 at a lower rate of 1Mbps. Note that the transmissions ($N_1 \rightarrow N_5$) and ($N_2 \rightarrow N_3$) cannot take place at the same time because of interference. In contrast to the first approach, the whole network-wide broadcast latency is now 23 time units.




Fig. 2. Alt_1 : Transmission schedule if each node can only broadcast a packet at most once.




Fig. 3. Alt_2 : Transmission schedule if each node can broadcast a packet more than once.

This examples illustrates the following important feature of broadcasting in multi-rate wireless meshes:

Property 1: If a node is to multicast to a number of its neighbouring nodes simultaneously, the maximum broadcast rate that can be used is constrained by the lowest rate to reach all these nodes independently. Accordingly, if the objective is to improve the broadcast latency, a new degree-of-freedom that can be used is *to allow a node to transmit the same packet more than once, to different subsets of its immediate downstream neighbors*.

By exploiting this degree-of-freedom, an intermediate node can transmit the packet at a higher rate to children that lie along the “more critical” sub-trees (i.e., those that might take longer to forward the packet) to their leaf nodes, and subsequently use a lower-rate transmission to a subset of the “less critical” sub-trees.

B. Illustrating the Role of Flow Rate on the Broadcast Latency

The example above can be easily extended to demonstrate the fact that the traffic rate of the broadcast flow itself has a significant impact on the feasibility of different broadcast trees, and consequently, the achievable broadcast latency bounds. The example in Section II-A considered the best way to distribute a single broadcast packet. Now, let us consider the case of a periodic (CBR-like) broadcast flow, where a packet of the same size is generated periodically every P secs. Even in the absence of any other cross traffic in the mesh, the scheduling mechanism must now ensure the avoidance of collisions among consecutive packets of the same flow. Mathematically, this can be achieved by ensuring the transit time of a packet through any *collision domain* is lower than the inter-packet interval, i.e., all channel activity, related to a specific packet, in the neighborhood of a transmitting node is completed before the arrival of the next packet. Channel activity refers both to the transmitter’s own transmissions (since it may transmit each packet multiple times, each for a different neighbor subset), as well as the subsequent transmissions by all downstream nodes within its *interference range*.

To illustrate this, first consider the flow F_1 , which generates a packet once every 24 time units. Also, assume that the first packet arrives at time $t = 0$. In such a case, approach Alt_2 is clearly superior, since N_1 can first transmit to N_2 (transmission ending at $t = 1$, wait for the transmission $N_2 \rightarrow N_3$ (ending at $t = 12$), and then complete the transmission $N_1 \rightarrow N_5$ (ending at $t = 23$), well before the arrival of the next packet from F_1 . Now, however consider the case of a flow F_2 , transmitting at an overall rate $\sim 10\%$ higher than F_1 , with an inter-packet gap of 22 time units. In this case, it is easy to see that Alt_2 is a non-feasible packet broadcast topology, since the second packet will arrive at N_1 *before* it has completed the transmission of the first packet. Eventually, as each packet waits progressively longer in the buffer at N_1 , the flow will suffer from some loss (e.g., due to buffer overflow at N_1 , resulting solely from the intra-flow congestion caused by Alt_2). In this case, the only alternative is to follow Alt_1 , whereby N_1 first transmits to both N_5 and N_2 (completing

the transmission at $t = 11$), followed by the transmission $N_2 \rightarrow N_3$ (ending at $t = 22$), before the arrival of the next packet at N_1 . Accordingly, the maximal delivery latency of the broadcast traffic for flow F_2 is almost 50% higher than that of flow F_1 , even though the traffic load of F_2 is only $\sim 10\%$ higher! Furthermore, if we consider a flow F_3 with an inter-packet gap of 20 time units, it is clear that there exist no feasible *lossless* broadcasting topology.

Property 2: The choice of the best broadcast distribution tree, and the extent to which individual nodes reap the potential benefit of multiple independent transmissions to different sets of downstream neighbors, can be highly sensitive to the traffic load of the broadcast flow. As a consequence, the achievable broadcast latency is itself strongly dependent on the traffic characteristics of the broadcast traffic.

To the best of the authors’ knowledge, the degree-of-freedom of allowing a node to transmit the same packet more than once, and the sensitivity of the achievable latency bounds on the traffic generation rate, have not been pointed out before. *Equally importantly, our results suggest that for broadcast traffic, there is a tradeoff between the maximum achievable throughput and the packet delivery latency.* These insights suggest that when, latency is a concern (e.g., in multi-player distributed games), the construction of the broadcasting topology should factor in the expected load from the application. Note that this new degree-of-freedom can be combined with others that have already been proposed, namely multi-radio, multi-channel [8] and network coding [16]. Of course, if the objective is instead to minimize the total energy consumption, then transmitting the same packet more than once will always result in worse performance.

III. RELATED WORK

Much work has been done in achieving *efficient* network layer multicast and broadcast in multi-hop wireless networks and wireless ad hoc networks. The majority of the work measures efficiency in terms of energy consumption [4], [5], [15], the number of transmissions (which is equivalent to energy consumption if broadcast power cannot be adjusted) [12] or the amount of overhead in route discovery and management [10], [11]. However, all of these approaches are based on single-rate wireless networks.

The work that is most similar to ours is [9], which considers the problem of achieving minimum broadcast latency in a single-rate wireless ad hoc network. They show that their optimisation problem is NP-hard and provided a polynomial time algorithm to solve the problem. If each node is allowed to multicast at most once, then our problem is a generalisation of that in [9] to the multi-rate case. However, as we have argued in Section II, the multi-rate problem has a number of unique properties not present in the single-rate case.

The problem of routing unicast flows in multi-rate multi-hop wireless networks has previously been studied in [1], [3], [8]. The authors of [8] proposed a routing metric which can be used for a multi-channel multi-hop wireless network. Their metric takes different transmission rates into account by

having the metric inversely proportional to the transmission rate. The authors in [3] used simulation to study the end-to-end UDP and TCP throughput of a multi-rate multi-hop path. Their simulation study revealed a number of interesting findings, for example, they found that a 2-hop path of 11 Mbps links can have very different throughput from that of 4-hop paths of 5.5Mbps links. The work in [1] shows that if the interference range is infinity, then the unicast routing path that minimizes the total path delay will also maximises the throughput between the source and destination. To deal with multi-rate links, [1] defines the medium-time metric (MTM) for each transmission rate. MTM essentially measures the time it takes to transmit a packet over a multi-rate links. It takes into account transmission delay (i.e. frame size divided by transmission rate) and the overheads, which in the case of IEEE802.11 includes RTS/CTS/ACK frames and channel contention. Note that the inclusion of channel contention is needed to account for intra-flow interference.

IV. OPTIMAL NETWORK-WIDE BROADCAST IN A MULTI-RATE WIRELESS MESH NETWORK

In this section we formulate the problem of finding the optimal network-wide broadcast topology in a multi-rate, multi-hop wireless mesh, i.e., the topology that minimizes the broadcast latency. The formulation essentially boils down to an integer programming problem that simultaneously determines a) the broadcasting topology, i.e., the packet distribution tree, and b) the broadcast scheduling, i.e., when, and to which subset of downstream neighbors, a node transmits a packet.

A. The modelling assumptions

The modelling assumptions are:

- 1) Each node in the network is equipped with one radio, with all radios tuned to a common channel.
- 2) By adjusting the modulation scheme, a node can multicast at different data rates. The same transmission power is used for all data rates. As a result, the transmission range a decreasing function of the data rate. Let s_{\max} denote the maximum transmission range. Also, we use a disc model for the transmission range.
- 3) A node's neighbours are all the nodes that can be reached using the lowest possible transmission rate.
- 4) A node can multicast at different rates to different subsets of its neighbours. Let $\{i_1, \dots, i_k\}$ be a subset of the neighbours of node j and the maximum rates which node j can use to reach these nodes independently are r_1, \dots, r_k respectively. If node j wants to multicast to i_1, \dots, i_k in one go, this can only be done at a rate of $\min(r_1, \dots, r_k)$ or lower.
- 5) We assume a binary interference model, as follows: If while a node k is receiving a frame, a node j within a radius κs_{\max} from node k transmits a frame, then the frame that k is receiving is assumed to be corrupted and lost. This corresponds to the interference model of [17]; a typical value of κ is 1.7. We will refer to " κs_{\max} " as the interference range. This is consistent

with the fact that although different links have different transmission ranges, the *transmission power* level does not vary with the link distance (the rate variation is due to the employment of different modulation schemes for different receiver power levels).

- 6) We assume an ideal MAC layer, as follows: Two nodes i and j can multicast at the same time if and only if node i 's multicast does not interfere with the intended recipients of node j 's multicast and vice versa.
- 7) We assume a centralised entity which schedules these multicasts so that, under the ideal MAC layer assumption, no two multicasts will interfere with each other.
- 8) Each node can multicast the same packet up to m_{\max} times, clearly to different subsets of its neighbors. $m_{\max} = 1$ corresponds to the conventional use of broadcast trees, where each node reaches all its child nodes in a single transmission.

Note that the basic building block of achieving the network-wide broadcast is a sequence of link layer multicasts (to a subset of the neighboring nodes) instead of link layer broadcasts.

B. Optimisation problem

Due to lack of space, the actual formulation of the optimization as an integer programming problem is reported in [7]. The key decisions in this optimisation problem are: (1) Whether a node should multicast and if so, to which of its neighbours; (2) The timing of these multicasts. To determine the timings of these multicast, we must make sure that a node can only multicast a packet after it has received it. Also, when some multicasts cannot take place at the same time because they interfere with each other, they must be scheduled so as to minimize the broadcast latency. Not surprisingly, this multi-rate broadcast problem is NP-hard.

Theorem 1: The minimum latency network-wide broadcast problem in a multi-rate wireless mesh network is NP-hard.

Proof: This follows from the fact that the minimum latency network-wide broadcast problem in a single-rate wireless mesh network, which is a special case of the multi-rate case, is NP-hard. The NP-hardness result for the single-rate case is given in [9]. \square

Given the hardness of the problem, the optimization tool can be executed only for relatively small and simple topologies. However, comparing the broadcast latency achieved by the optimization technique with that achieved by a conventional tree-formation algorithm (that does not exploit the rate diversity of different links) will provide a sense of the degree of improvement that may be achieved by a multi-rate aware algorithm. In Section V, we will also propose a set of heuristic measures to solve this problem. Comparisons of the integer programming formulation with existing algorithms and new proposed heuristics are provided in Section VI-A.

V. HEURISTIC FOR LOW LATENCY BROADCAST TREE

Due to lack of space, we will only present an algorithm for the case where each node may broadcast a packet at most

once. For the algorithm which allows multiple transmission per node, the reader can refer to [7].

We now present an heuristic algorithm to create efficient delivery trees for broadcast packets in a variable-rate mesh network. Broadly speaking, any heuristic algorithm has to make two choices: 1) Whether a node should multicast and if so, to which of its neighbours; (2) The timing of these multicasts. Note that these two decisions are closely coupled since a multicasting node can only multicast after it has received the packet and radio interference means that the multicasts must be scheduled so that interfering multicasts do not take place at the same time. Given the hardness of the problem, we decompose the algorithm into two logically independent steps:

- **Topology Construction:** In this step, the aim is to compute a broadcast tree (or a spanning tree) T . This step merely decides the hierarchy of the broadcast tree, i.e., identifies at each intermediate nodes, the child nodes that it is responsible for.
- **Transmission Scheduling:** The second step schedules (for now, we conceptually assume a centralized scheduler) the independent transmissions by each node, taking into account that (i) a node can only multicast after it has received the packet, and (ii) interfering multicasts cannot occur concurrently.

Clearly, this decomposition of the overall optimization problem is not optimal. However, as already noted, a joint optimization is computationally infeasible, except for trivially small mesh topologies. The heuristic for the Topology Construction phase is called the Broadcast Incremental Bandwidth (BIB) technique, and is presented in Section V-A. This is followed by Section V-B where we present the scheduling heuristic, that takes into account the conflict graph of the underlying tree topology. Note that the scheduling heuristic is independent of the Topology Construction algorithm and can take any spanning tree as its input.

Let us first introduce some common mathematical notation. The entire wireless mesh is represented as a graph (V, E) , with the mesh nodes forming the vertices and the edges representing the direct link between any two nodes. Accordingly, $(i, j) \in E$ denotes the direct unicast link between nodes i and j . Based on the distance between such a node pair, each link (i, j) can be associated with a transmission rate R_{ij} which is the maximum transmission rate that can be used between the two nodes. The transmission rate $R_{ij} = 0$ if i and j are not one-hop neighbors, i.e., if j cannot correctly receive a packet from i even if i transmits at the lowest rate and maximum power.

A. The Broadcast Incremental Bandwidth (BIB) Topology Construction Algorithm

We first compute the tree from a source node s to all the other nodes $V - \{s\}$ in the wireless mesh. Any candidate algorithm should obviously exploit the wireless multicast advantage [15] to reach multiple neighbours in a single transmission. The algorithm must also take into account the multi-rate nature of the problem, for example, considering if a node

should reach another neighbour using their direct hop (at a lower transmission rate) or via two hops of higher rates. In addition, the algorithm should be aware of the interference between neighboring multicasts. For example, if a number of multicasts are within the interference range of each other, they can only take place one after another. As a special case, if all the transmissions interfere with one other (i.e. the interference radius is infinity) then only one multicast can take place at a time. In such a case, minimizing the broadcast latency is identical to minimizing the total transmission time of all the multicasts, i.e., the resulting tree should be a radio analogue of the wired minimum spanning tree.

The Broadcast Incremental Bandwidth (BIB) algorithm is very similar to the BIP algorithm [15] in that both use a modified version of Prim's algorithm, greedily adding links to an existing tree such that the incremental cost is minimized. However, while BIP focuses on the development of low-energy packet distribution trees, BIB primarily aims to choose high-rate links, since the transmission of a packet by a transmitter to its neighbors is constrained by the slowest of the point-to-point links between the transmitter node and each individual neighbor.

For any particular packet forwarding topology, let $N(x)$ denote the one-hop neighbors of node x and $Neigh(x) (\subset N(x))$ the designated downstream neighbors. In other words, node x must broadcast the packet so that it is correctly received by all nodes $y : y \in Neigh(x)$. Clearly, the transmission rate $R(x)$ of node x is given by the slowest downstream link, i.e.,

$$RN_{Neigh(x)} = \min_{k \in Neigh(x)} R_{xk}. \quad (1)$$

To apply a minimum-cost tree construction algorithm such as Prim's, the *transmission latency cost* TL_{ij} for a link (i, j) between two nodes i and j is initially set to be the inverse of the transmission rate, i.e., $TL_{ij} = \frac{1}{R_{ij}}$. The algorithm is initiated with a tree T that initially contains only the source node s as the root, with the cost of any other node x , $C(x)$, set to TL_{sx} . Each node in the tree has a *tree node cost* $TC(\cdot)$ reflecting the cost of forwarding the packet to the "slowest" child node; at the beginning, $TC(s) = 0$. In each subsequent step, the node x with the current minimum cost is added to the tree. Let P_x denote the parent of the chosen node x ; clearly, P_x is already part of the tree. The tree node cost for P_x , $TC(P_x)$, is then incremented by the cost associated with node x . Additionally, for each neighbor y of P_x that is not already in the tree, its cost is dynamically updated to be the difference between the transmission rate cost $TL_{P_x y}$ (which does not change) and the tree node cost $TC(P_x)$ (which might change with each iteration) if it is more favourable to reach y by using incremental broadcast from P_x . This dynamic modification of the link cost at each iteration distinguishes this approach from the basic Prim's algorithm and is designed to reflect the wireless broadcast advantage. The pseudocode for the BIB algorithm is presented in Figure 5.

The basic BIB algorithm may also be enhanced (using ideas presented in [2]) for high-performance *reliable* broadcasting

```

Procedure:  $UpdateCost((i, j), k)$ 

/* Let  $Neigh(i)$  be the current downstream
one-hop neighbors of  $i$ . */
 $c \leftarrow \frac{1}{R_{ik}} - \frac{1}{RN_{Neigh(i)}}$ 
if ( $C(k) > c$ )
     $C(k) \leftarrow c$ ;  $P_k \leftarrow i$ .

```

Fig. 4. The cost of node k is modified to reflect only the additional (incremental) cost that would be incurred if k subsequently became a child (downstream neighbor) of node i .

```

Procedure:  $BIB(s, V)$ 

Set  $T = \{s\}$ ,  $S = V - \{s\}$ 
For ( $x \in S$ ),  $C(x) = \frac{1}{R_{sx}}$ 
while ( $S \neq \emptyset$ )
     $x \leftarrow MinCostNode(S)$ 
     $T \leftarrow T \cup \{(P_x, x)\}$ 
     $S \leftarrow S \setminus \{x\}$ 
    for ( $y \in \{N(x) \cap S\}$ )
        if ( $C(y) > \frac{1}{R_{xy}}$ )
             $C(y) \leftarrow \frac{1}{R_{xy}}$ ;  $P_y \leftarrow x$ .
    for ( $y \in \{N(P_x) \cap S\}$ )
         $UpdateCost((P_x, x), y)$ 
end-while

```

Fig. 5. The BIB algorithm. The function $MinCostNode(S)$ returns the node x in S with the minimum cost and also deletes it from the set S . The broadcast tree is the set of eventual links in T .

trees, where the tree construction process considers the quality of each link (and the resulting retransmissions needed) in addition to the link rate. Since our focus is primarily in exploring the tradeoffs in constructing high-performance packet broadcasting mechanisms, we do not explore this aspect further in this paper.

B. The Scheduling of Transmissions

While a broadcast tree determines the *sequence* of transmission (as a child can multicast a packet only after receiving it from its parent), the exact timing of the various multicasts (especially relative to different branches of the tree) still needs to be determined. We will approach this problem by formulating it as a scheduling problem with precedence constraints (which enforces that a node can only multicast after it has received the packet) and conflict graph (which models the interference between different transmissions).

Let $V_b = \{b_1, b_2, \dots, b_k\} \subset V$ be the set of all the branch points in the broadcast tree T . (Note that T in principle can be any spanning tree of the graph $G = (V, E)$.) We further assume that b_1 is the source node. The packet delivery sequence can be modelled by a directed graph (tree) $G_b = (V_b, E_b)$ such that $(b_i, b_j) \in E_b$ if and only if it is an edge in the tree T . For each node $b_i \in V_b$, we assign a cost $t(b_i)$ which is the minimum multicast transmission time it takes the node b_i to transmit a fixed size packet of p bits to all its children in the broadcast tree T . We also define an undirected conflict graph $G_c = (V_c, E_c)$ such that $V_c = V_b$ and $(b_i, b_j) \in E_c$ if and only if the multicast of b_i interferes with the reception of

the children of b_j in T or vice versa (this is consistent with the fact that the transmit power remains fixed, irrespective of the link distance).

Formally, a schedule can be defined as a mapping $\tau : V_b \rightarrow \mathbb{R}$ which gives the transmission time of node $b_i \in V_b$. Given G_b , $t(b_i)$ and G_c , a valid schedule is one which meets the following constraints:

- 1) The source multicasts at time zero: $\tau(b_1) = 0$.
- 2) A node can only multicast after it has received the packet: $\forall (b_i, b_j) \in E_b$, $\tau(b_j) \geq \tau(b_i) + t(b_i)$
- 3) For any edge $(b_i, b_j) \in G_c$, we have $(\tau(b_i), \tau(b_i) + t(b_i)) \cap (\tau(b_j), \tau(b_j) + t(b_j)) = \emptyset$. Note that (\cdot, \cdot) here also denotes an open interval in \mathbb{R} . Although the same notation is used to denote both an open interval and an edge of a graph, the usage should be clear from the context.

The objective of the scheduling algorithm is to find a valid schedule τ which minimizes the broadcast latency $\max_{b_i \in V_b} (\tau(b_i) + t(b_i))$.

We have designed a greedy algorithm to solve the above scheduling problem. The details of the algorithm are described in Figure 6. The basic idea is that, in each round of the algorithm, there are a number of qualified nodes $Q = \{q_1, q_2, \dots, q_m\}$ (Note: a node is qualified if it has already received the packet.) and that the earliest possible multicast times for these nodes, denoted as $e(q_i)$, are known. For each qualified node in Q , we compute a priority measure $f(q_i)$ — there are many possible choices for $f(q_i)$ and we will describe later one which estimates the worst case latency required to reach all the descendants of q_i in T . In each round, the node q_i that has the largest value of $f(q_i)$ is chosen (ties are broken arbitrarily) and for this particular q_i , we set $\tau(q_i) = e(q_i)$. (We will discuss how $e(q_i)$ is maintained by the algorithm later.)

In order to describe our choice of priority measure $f(\cdot)$, we first need to define some additional notation. Let $D(b_i)$ denote the set of all descendants of b_i in the directed graph G_b . For any $x \in D(b_i)$, let $P(b_i, x)$ denote the set of nodes on the path from b_i to x in G_b (inclusive of both b_i and x). We define $w(b_i)$ as the time needed to reach all the descendants of b_i in T in the absence of radio interference, formally we write

$$w(b_i) = \max_{x \in D(b_i)} \sum_{y \in P(b_i, x)} t(y) \quad (2)$$

Note that $w(b_i)$ is a lower bound on the time required to reach all descendants of b_i . In our algorithm, for a qualified node q_i , we define $f(q_i)$ as follows:

$$f(q_i) = e(q_i) + \sum_{q_j : (q_i, q_j) \in G_c} t(q_j) + w(q_i) \quad (3)$$

Thus $f(q_i)$ is an estimate of the downstream latency to reach all the descendants of q_i in the *worst possible scenario*, where the node q_i can only transmit after *all* the qualified nodes that interferes with q_i have transmitted.

```

Procedure: Schedule( $G_b, \{t(b_j)\}, G_c$ )
Set  $Q = \{b_1\}$ ,  $e(b_1) = 0$ 
For ( $b \in V_b$ ), set  $PTIME(b) = [0, \infty)$ 
while ( $Q \neq \emptyset$ )
  For ( $b \in Q$ ), compute  $f(b)$  according to Eq. (3)
   $x \leftarrow \arg \max_{b \in Q} f(b)$ 
  Set  $\tau(x) = e(x)$ 
   $C(x) \leftarrow \{ \text{children of } x \text{ in } G_b \}$ 
   $Q \leftarrow (Q \setminus x) \cup C(x)$ 
  For ( $b \in V_b \setminus \{x\}$ )
    If ( $(b, x) \in E_c$ )
       $PTIME(b) \leftarrow PTIME(b) \setminus (\tau(x), \tau(x) + t(x))$ 
  For ( $b \in C(x)$ )
     $PTIME(b) \leftarrow PTIME(b) \setminus [0, \tau(x) + t(x))$ 
  For ( $b \in Q$ )
    For any interval  $\mathcal{I}$  in  $PTIME(b)$ 
      If  $\text{length}(\mathcal{I}) < t(b)$ 
        Remove  $\mathcal{I}$  from  $PTIME(b)$ 
    Update  $e(b)$  based on  $PTIME(b)$ 
end-while

```

Fig. 6. The scheduling algorithm

To keep track of the permissible time for a node $b_i \in V_b$ to multicast, we define a data structure $PTIME(b_i)$, which maintains the set of all time intervals on the real line over which node b_i is allowed to multicast. If a node b_j is scheduled to multicast in the interval $[\tau(b_j), \tau(b_j) + t(b_j)]$ and b_j 's multicast interferes with b_i 's (i.e. $(b_i, b_j) \in G_c$), then the open interval $(\tau(b_j), \tau(b_j) + t(b_j))$ will be removed from $PTIME(b_i)$ by computing the set difference $PTIME(b_i) \setminus (\tau(b_j), \tau(b_j) + t(b_j))$. The earliest possible multicast time for a qualified node q_i can be obtained easily from $PTIME(q_i)$.

C. Maximum end-to-end throughput

The above discussion of the tree construction and scheduling algorithms focused on the case of a *single* packet, attempting to minimize the broadcast latency for a single packet. This approach is clearly directly applicable when the data rate of the broadcast stream is low enough (e.g., for control traffic), where one can safely assume the absence of interference/scheduling conflicts among successive packets of the same flow. For higher rate data flows, it is important to compute the maximum achievable throughput of a broadcast tree, defined as the maximum data rate that can be sustained without their being any scheduling-related conflicts between packets of the same flow.

Given the broadcast tree T computed in the Topology Construction phase, it is possible to compute the maximum achievable throughput for this multicast tree, by essentially computing the minimum permissible gap between successive packets. Using the definition as in Section V-B, given the set of branch points $V_b = \{b_1, b_2, \dots, b_k\}$, the multicast transmission time $t(b_j)$ of branch point b_j for a fixed size packet of p bits and the conflict graph of the multicasts $G_c = (V_c, E_c)$, the maximum throughput ϕ is given by

$$\phi \leq \frac{p}{t(b_i) + \sum_{b_j: b_j \neq b_i \ \& \ (b_i, b_j) \in E_c} t(b_j)} \quad \forall b_i \in V_b \quad (4)$$

This equation shows that the throughput is limited not only

by each node's multicast latency, but by the maximum time it takes to cross an individual collision domain (i.e. the interference region around each node). (This is a generalisation of the throughput of a unicast path in a multirate wireless network [1]).

VI. SIMULATED PERFORMANCE STUDIES

In this section we study the performance of BIB to solve the low latency network-wide broadcast problem in a multi-rate wireless mesh network when each node can multicast a packet at most once. For the purpose of comparison, we will study altogether 3 heuristics. All these 3 heuristics have the same structure, computing a broadcast tree and then followed by the scheduling algorithm in Section V-B. In other words, these algorithms only differ in how the broadcast tree are computed. The algorithms to be considered are:

- 1) Algorithm BIB: Uses BIB in Section V-A to compute the broadcast tree.
- 2) Algorithm SPT: The broadcast tree is the shortest path tree (SPT) given by Dijkstra's algorithm. (This algorithm does not exploit the broadcast advantage while computing the tree; however, during transmission, each node transmits to its child nodes in a single transmission).
- 3) Algorithm CDS: This heuristic assumes that all broadcasts are at the lowest rate. It first computes a broadcast tree where only the lowest broadcast rate is used and then followed by scheduling algorithm in section V-B. The broadcast tree is computed using a greedy approximation of the minimum connected dominating set (CDS). The algorithm starts with the source broadcasting. In each round of the algorithm, a new node is chosen to broadcast. This is iterated until all nodes are covered, i.e. having received the packet. The greedy algorithm chooses, in each round, the node whose broadcast will maximize the number of currently "uncovered" nodes.

The simulations in this section are based on the range-rate relationship in Table I.

A. Small, Regular grid topology

We consider a regular 2-by-4 planar grid network whose physical topology is given in Figure 7. (This topology is deliberately chosen to be a simple and small thus allowing us to compute the optimal broadcasting solution via the integer programming). The horizontal and vertical separations between the nodes are denoted by, respectively, L_x and L_y . The value of L_y is 360 while L_x can be 120, 220 or 320. The values of L_x are designed to give different connectivity pattern and transmission link rates. For small values of L_x , each node can have 6 or 7 neighbours, but, for large values of L_x each node may have only 2 or 3 neighbours.

For each given physical topology (i.e. given values of L_x and L_y) and each possible choice of broadcast source node s , we compute the worst case delay given by the heuristics (denoted as $d_{\text{BIB}}(L_x, L_y, s)$, $d_{\text{SPT}}(L_x, L_y, s)$ and $d_{\text{CDS}}(L_x, L_y, s)$) and the optimal solution given the integer programming formulation (denoted as $d_{\text{OPT}}(L_x, L_y, s)$). As

a measure of performance of each heuristic, we compute, for each given physical topology, the following indices:

$$r_{\text{METHOD}}(L_x, L_y) = \left(\prod_{s=1, \dots, 8} \frac{d_{\text{METHOD}}(L_x, L_y, s)}{d_{\text{OPT}}(L_x, L_y, s)} \right)^{\frac{1}{8}}$$

where METHOD = BIB, SPT or CDS. The results are tabulated in Table II. It can be seen that BIB performs best out of the three heuristics for all the 3 different topologies used. This is due to the fact that BIB is able to exploit both the differential increment in link rates as well as the wireless broadcast advantage. As an illustration, let us consider the case where $L_x = 120$ and source node is 6. In Figure 8, we show the connectivity of Node 4 for this topology. The connectivity of the other nodes can be readily deduced from this. The number next to the arrow shows the relative cost in packet transmission delay. The minimum cost is 1 (when transmitted at 11Mbps) and the maximum is 11 (when transmitted at 1 Mbps). The latency given by the heuristics BIB, SPT and CDS are, respectively, 6.5, 17.5 and 11 time units. The optimal is 6.5 which is achievable by BIB. *The results demonstrate that the use of a multi-rate aware broadcast tree can reduce the broadcast latency by 50 – 60%.*

Figures 9 and 10 show, respectively, the broadcast tree given by SPT and BIB. Although Node 6 (the source) multicasts to the same set of neighbours in both cases, there are significant differences after that. For SPT, in addition to the source, 3 other nodes (Nodes 5, 7 and 8) will multicast, making the total number of multicasts four. In addition, these three multicasts interfere with each other so their transmission cannot take place in parallel. This result in poor performance of SPT.

On the other hand, the BIB algorithm exploits the wireless multicast advantage and requires only two multicasts in total. The second multicast in BIB (see Figure 10) is performed by Node 2 and it reaches all the three remaining nodes in one go. This demonstrates that BIB is able to exploit wireless multicast advantage. Another feature of BIB is that it exploits incremental link rates. Consider the SPT tree in Figure 9, note that there is a simple modification of the tree which will result in a better latency. This can be seen by noticing that there are two shortest paths from Nodes 6 to 3: 6-7-3 and 6-2-3. Either one of these may be chosen by the Dijkstra's algorithm. However, if we replace the link (7,3) with cost 5.5 in Figure 9 by link (2,3) with cost 1, we still have a shortest path tree but the broadcast latency will be reduced. This is precisely what BIB does and chooses link (2,3) because it has a smaller incremental cost. It is also important to point out that the BIB tree in Figure 10 is in fact also a shortest path tree though one with better multicast property.

For the network in Figure 8, since all nodes are within the transmission range of Node 6. The CDS algorithm will use one broadcast which takes 11 time units.

B. Heuristic performance in random topology

In this section we compare the performance of the three heuristics using randomly generated topologies of different




Fig. 7. Regular grid topology used.

$L_y = 360$			
heuristic	$L_x = 120$	$L_x = 220$	$L_x = 320$
BIB	1.0231	1.0168	1.0595
SPT	2.6833	1.8516	1.2649
CDS	1.8974	1.8516	1.7889

TABLE II

THE PERFORMANCE OF THE HEURISTICS BIB, SPT AND CDS COMPARED TO THE OPTIMAL SOLUTION FOR A REGULAR 2-BY-4 GRID NETWORK.

sizes (as measured by the number of nodes in the network). For each network size, we generate 100 topologies whose nodes are uniformly randomly distributed in a square of 1 km². Since the network size that we use is at least 30, integer programming is not able to give us a reference in reasonable time. Instead, we choose to normalise the delay obtained from the heuristics by the delay given by the Dijkstra's algorithm (which is the shortest delay possible when there is no limit to the number of radios, channels and times a node can transmit a packet). Thus the minimum value of normalised delay is unity. The result that we will show is the geometric mean, over 100 network instances of a fixed size, of the normalised delay and the throughput (computed by Equation (4)).

The results are given in Figures 11 (for delay) and 12 (for throughput). It turns out that good performance for delay also means good performance for throughput and vice versa, since we have confined our study to the case where each node performs a link-layer multicast at most once. BIB performs best in these experiments and then followed by SPT and CDS, with BIB reducing the broadcast latency by 50% or more. It shows that BIB is able to exploit the multiple transmission rates available. While the SPT algorithm does not exploit the wireless broadcast advantage, CDS fails to exploit the multi-rate feature, thus leading to poorer performance.

The failure of SPT to exploit wireless broadcast can also




Fig. 8. Transmission cost of a regular grid topology.




Fig. 9. The SPT tree with $L_x = 120$ and $L_y = 360$.




Fig. 10. The BIB tree with $L_x = 120$ and $L_y = 360$.




Fig. 11. The geometric mean of the normalised latency of BIB, SPT and CDS.

be seen from Figure 13 which shows that SPT on average uses the most number of multicasts per tree out of the three heuristics. Although CDS uses the least number of multicasts per tree, it fails to exploit the higher transmission rates, thus resulting in the worst latency and the lowest throughput.

VII. CONCLUSIONS AND FUTURE WORK

We have introduced a novel type of wireless mesh network operation, where a node can multicast at the link layer to different subsets of neighbours at different transmission rates. In particular, we study the problem of realising low latency broadcast in such networks. We show that the multi-rate broadcast problem is significantly different from the single-rate case. Since the minimum latency multi-rate broadcast problem is NP-hard, we propose a heuristic which takes both wireless multicast advantage and multi-rate into consideration. Simulation studies using the ideal MAC layer assumption show that significant gain can be achieved by exploiting multiple rates available.




Fig. 12. The mean throughput of BIB, SPT and CDS.




Fig. 13. The mean number of multicasts per tree for BIB, SPT and CDS.

In ongoing work, we are evaluating the performance of BIB with a decentralized 802.11-type MAC to better understand the comparative benefit of cross-layer optimization between the routing and scheduling functions, and also evaluating alternative tree formation heuristics. Future work includes the development of heuristics for the case where a node is allowed to multicast a packet more than once.

REFERENCES

- [1] B. Awerbuch, D. Holmer, and H. Rubens. High throughput route selection in multi-rate ad hoc wireless networks. *Lecture Notes in Computer Science*, 2928:253–270, 2004.
- [2] S. Banerjee and A. Misra. Energy-efficient broadcast and multicast trees for reliable wireless communication. In *WCNC*, 2003.
- [3] S. Bansal, R. Shorey, and A. A. Kherani. Performance of tcp and udp protocols in multi-hop multi-rate wireless networks. In *WCNC*, volume 1, page 231, 2004.
- [4] M. Cagalj, J. Hubaux, and C. Enz. Minimum-energy broadcast in all-wireless networks: Np-completeness and distribution issues. In *MobiCom*, pages 172–182, 2002.
- [5] J. Cartigny, D. Simplot, and I. Stojmenovic. Localized minimum-energy broadcasting in ad-hoc networks. In *INFOCOM*, page 2210, 2003.
- [6] P. Chaporkar, A. Bhat, and S. Sarkar. An adaptive strategy for maximizing throughput in mac layer wireless multicast. In *Mobihoc*, pages 256–267, 2004.
- [7] C. Chou, A. Misra, and J. Qadir. Low latency broadcast in multi-rate wireless mesh networks. Technical Report TR 0514, School of Computer Science and Engineering, University of New South Wales, 2005.
- [8] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *Mobicom*, pages 114–128, 2004.
- [9] R. Gandhi, S. Parthasarathy, and A. Mishra. Minimizing broadcast latency and redundancy in ad hoc networks. In *Mobihoc*, pages 222–232, 2003.
- [10] C. Gui and P. Mohapatra. Scalable multicasting in mobile ad hoc networks. In *INFOCOM*, page 2119, 2004.
- [11] S. Lee, W. Su, and M. Gerla. On-demand multicast routing protocol in multihop wireless mobile networks. *Mobile Networks and Applications*, 7(6):441–453, 2002.
- [12] W. Lou and J. Wu. On reducing broadcast redundancy in ad hoc wireless networks. *IEEE Trans on Mobile Computing*, 1:111, 2002.
- [13] Scalable Networks Inc. <http://www.scalable-networks.com/>.
- [14] Y. Tseng, S. Ni, Y. Chen, and J. Sheu. The broadcast storm problem in a mobile ad hoc network. *Wireless Network*, 8(2/3):153–167, 2002.
- [15] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. On the construction of energy-efficient broadcast and multicast trees in wireless networks. In *INFOCOM*, volume 2, page 585, 2000.
- [16] Y. Wu, P. A. Chou, Q. Zhang, K. Jain, W. Zhu, and S. Kung. Network planning in wireless ad hoc networks: a cross-layer approach. *Selected Areas in Communications, IEEE Journal on*, 23(1):136, 2005.
- [17] K. Xu, M. Gerla, and S. Bae. Effectiveness of rts/cts handshake in ieee 802.11 based ad hoc networks. *Ad Hoc Networks*, 1(1):107, 2003.