

Low-Latency Broadcast in Multirate Wireless Mesh Networks

Chun Tung Chou, *Member, IEEE*, Archan Misra, *Member, IEEE*, and Junaid Qadir

Abstract—In a multirate wireless network, a node can dynamically adjust its link transmission rate by switching between different modulation schemes. In the current IEEE802.11a/b/g standards, this rate adjustment is defined for unicast traffic only. In this paper, we consider a wireless mesh network (WMN), where a node can dynamically adjust its link-layer multicast rates to its neighbors, and address the problem of realizing low-latency network-wide broadcast in such a mesh. We first show that the multirate broadcast problem is significantly different from the single-rate case. We will then present an algorithm for achieving low-latency broadcast in a multirate mesh which exploits both the wireless multicast advantage and the multirate nature of the network. Simulations based on current IEEE802.11 parameters show that multirate multicast can reduce broadcast latency by 3–5 times compared with using the lowest rate alone. In addition, we show the significance of the product of transmission rate and transmission coverage area in designing multirate WMNs for broadcast.

Index Terms—Broadcast, multicast, multirate transmissions, quality-of-service, wireless mesh networks (WMNs).

I. INTRODUCTION

WIRELESS MESH NETWORKS (WMNs) are increasingly being viewed as a cheap and easily deployable extended-area access technology for suburban and urban community-based networks. In these scenarios, the nodes in a WMN often act as both *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 personal digital assistants (PDAs). In fact, one popular use of WMNs is to extend the benefits of wide-area connectivity to a larger community, by using the multihop 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 wired access medium.

Two aspects of WMN research seem to be especially popular at present.

- 1) *Use of multichannel, multiradio mesh nodes*, especially as recent research (e.g., [15]) demonstrates that the use of multiple radios on a single node, each tuned to possibly distinct channels, can significantly increase the overall network capacity.

Manuscript received October 1, 2005; revised March 4, 2006 and May 1, 2006. This work was supported in part by the Australian Research Council Discovery under Grant DP0664791 and in part by a University of New South Wales Faculty Research Grant.

C. T. Chou and J. Qadir are with School of Computer Science and Engineering, University of New South Wales (UNSW), Sydney 2052, Australia (e-mail: ctchou@cse.unsw.edu.au; junaidq@cse.unsw.edu.au).

A. Misra is with IBM T. J. Watson Research Center, Hawthorne, NY 10532 USA (e-mail: archan@us.ibm.com).

Digital Object Identifier 10.1109/JSAC.2006.881621

- 2) *Multirate MAC protocols*, especially as researchers move beyond IEEE 802.11-based single-rate medium access control (MAC), and study the throughput and fairness issues that arise when adaptive modulation schemes dynamically modify the data rate on a particular link in response to variation in signal-to-noise ratio (SNR).

Such advances should solve one of the fundamental problems of existing MANET-oriented networks, namely, the sharp drop in multihop throughput to a few kilobit per second, even though individual wireless links are operated at higher speeds such as 54 Mb/s. However, most research on both these aspects have focused on the *unicast* traffic scenario, e.g., [6] demonstrates the use of new unicast routing metrics that account for both the intraflow interference and differential link transmission rates in multichannel, multirate mesh environments.

In this paper, we study the problem of efficient routing and packet distribution for *broadcast* (or, more generally, multicast) traffic flows in a multirate WMN, where each node is equipped with one radio tuned to a common channel. (The problem for the multiradio multichannel multirate case is treated separately in [14]). We assume that the MAC layer of future WMNs will provide some form of multicast support, where the transmitter may be able to specify the transmission rate of the MAC-layer multicast, and, either explicitly or implicitly, the recipients of the multicast. To the authors' best knowledge, such multirate multicast capability has not been studied in the literature before. While the current IEEE 802.11a/b/g standards mandate the transmission of the control frames [e.g., ready-to-send/clear-to-send/acknowledgment (RTS/CTS/ACK) at the lowest rate (e.g., 6 Mb/s for IEEE 802.11a), transmission rates for broadcast data are typically implementation-specific. Clearly, future MAC protocols may permit more flexibility, e.g., relatively simple techniques have been proposed (e.g., [3]) to support such selective-broadcast at the wireless link layer.

Assuming such MAC-layer multirate multicast capability, our goal is to study how low-latency (and possibly high-throughput) network-layer broadcast¹ of *data traffic* can be realized. Our focus on the latency of data broadcasts distinguishes us from much of earlier research, where broadcasting was used principally for relatively infrequent control traffic (e.g., route establishment in MANETs). We believe that our focus on developing algorithms for low-latency data broadcast is important for many practical WMN applications, e.g., 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.

¹We assume that multicast can be realized by pruning the broadcast tree.

Multicast routing algorithms for multihop wireless networks have primarily focused on *energy efficiency* [19]. While energy efficiency is important for battery-powered nodes, it is less relevant in many WMN scenarios, where the nodes are relatively static (e.g., mounted on rooftops) and directly connected to regular power outlets. In mesh environments designed to support potentially high-bit rate multicast multimedia streams or interactive multiuser multimedia applications, it is therefore necessary to develop new routing techniques that allow for low-latency, high-throughput multihop wireless packet broadcast. We formalize our design goals for high-performance multihop broadcast by using the metric: *the broadcast latency*, computed as the maximum delay between the transmission of a packet by a source node and its eventual reception (over multihop paths) by all the intended receivers. Our current research addresses the following questions.

- **Effect of multirate links on efficient broadcasting:** Is multirate multicast at the link layer necessary for realizing a low broadcast latency? By what factor can the introduction of multirate multicast reduce the broadcast latency compared with single-rate multicast?
- **Choice of transmission rates in multirate networks:** If multirate multicast is to be introduced, how many different transmission rates do we need? How should they be chosen? Are some rates more efficient than others?

Our primary goal will be to show that the presence of multirate schemes, opens up new possibilities for broadcast traffic distribution that do not seem to have been explored before. Our contributions in this paper are the following.

- 1) Demonstrating that the broadcast latency is not necessarily minimized by tree-based packet distribution topologies, where each intermediate node forwards a packet by a single broadcast to its set of child nodes. Rather, optimal or efficient packet broadcasting is often achieved by having an intermediate node perform *multiple multicasts* of the same packet, each of which is directed towards a different subset of child nodes.
- 2) Designing wireless broadcast algorithms that exploit the *wireless multicast advantage* (WMA),² [19] as well as the multirate nature of wireless meshes.
- 3) Proposing the use of the product of transmission rate and transmission coverage area to measure the efficiency of using a particular transmission rate in achieving low broadcast latency in multirate environments.

II. IMPACT OF MULTIRATE LINKS ON EFFICIENT BROADCASTING

Effective packet broadcasting in a multirate WMN depends strongly on the interaction between the routing and MAC layers. Intuitively, a pure flooding strategy, where each intermediate node rebroadcasts 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 (also known as (a.k.a), the broadcast storm problem [18]). Thus, efficient broadcast strategies typically aim to build a distribution tree, where

²Due to the broadcast nature of the wireless medium, a single transmitting node can reach multiple one-hop neighboring nodes with a single transmission.

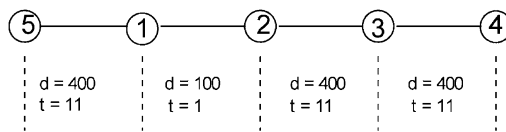


Fig. 1. Motivating example for broadcasting in a multirate WMN.

TABLE I
MAXIMUM TRANSMISSION RANGE IN METERS FOR DIFFERENT IEEE802.11b TRANSMISSION RATES

Transmission rate (Mbps)	Transmission range (m)	RAP (Mbps-km ²) (For Sec. VII)
1.0	483	0.73
2.0	370	0.86
5.5	351	2.13
11.0	283	2.76

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 unicast a packet separately to each of its downstream neighbors, is also wasteful. By failing to exploit the WMA, the all-unicast approach not only maximizes the forwarding latency at each intermediate node, but can induce additional contention-induced backoff delay at the MAC layer. Accordingly, the implicit assumption 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 that if each node in the distribution tree is limited to broadcasting a packet once, it can lead to suboptimal behavior in multirate WMN environments.

To understand the closely coupled nature of the broadcast tree formation and the MAC-layer scheduling, consider the topology in Fig. 1 with five nodes, labeled as Nodes 1–5, arranged in a straight line. For simplicity, we will refer Node 1 as N_1 , and so on, in the text. In Fig. 1, the d value between two nodes indicates the physical distance in meters between them. We assume each node is equipped with an 802.11b radio tuned to the same channel. By using the Qualnet simulator [17] as a reference, we obtain the transmission rate versus transmission range relationship in the first two columns in Table I, assuming a two-ray propagation model [16]. Note also that the interference range in Qualnet is 520 m. Thus, there are four links in the network configuration in Fig. 1. Link (1, 2) has a capacity of 11 Mb/s, while the other three links have a capacity of 1 Mb/s. 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 Fig. 1.

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 1 Mb/s 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

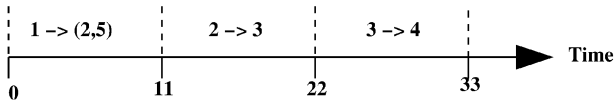


Fig. 2. Alt₁: Transmission schedule if each node can only broadcast a packet at most once.

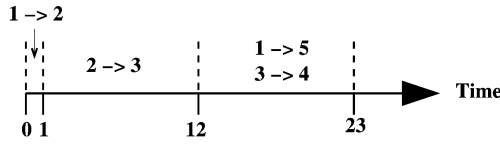


Fig. 3. Alt₂: Transmission schedule if each node can broadcast a packet more than once.

schedule depicted in Fig. 2, and leads to a broadcast latency of 33 time units.

In the second approach, which we call Alt₂, we allow each node to broadcast the same packet more than once. Fig. 3 depicts the transmission schedule. It shows the source N_1 transmitting the same packet twice. It first transmits to N_2 at 11 Mb/s (at time 0), taking 1 time unit. It then transmits the same packet again at time 12 to N_5 at a lower rate of 1 Mb/s. 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 Alt₁, the broadcast latency for Alt₂ is only 23 time units. This example illustrates the following important feature of broadcasting in multirate WMNs: If a node is to multicast to a number of its neighboring nodes simultaneously, the maximum permitted broadcast rate is constrained by the lowest of the individual link rates. Accordingly, if the objective is to improve the broadcast latency, a new degree-of-freedom (DoF) that can be used is *to allow a node to transmit the same packet more than once, to different subsets of its downstream neighbors*.

By exploiting this DoF (something that we believe has not been effectively used before), an intermediate node can transmit the packet at a higher rate to children that lie along the “more critical” subtrees (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” subtrees. Note that this new DoF can be combined with others that have already been proposed, namely, multiradio, multichannel [6], and network coding [20]. It is instructive to point out that if the objective is to minimize the total energy consumption, then transmitting the same packet more than once will always result in poorer performance.

III. RELATED WORK

Much work has been done in achieving *efficient* network-layer multicast and broadcast in multihop wireless networks and MANETs. The metrics typically used are energy consumption [19], the number of transmissions (which is equivalent to energy consumption if transmission power cannot be adjusted) [13] or the overhead in route discovery and management [9]. The limited work on broadcast latency has been based on single-rate wireless multihop networks.

The work that is most similar to ours is [7] which considers the problem of achieving minimum broadcast latency in a single-rate wireless ad hoc network. They show that their optimization problem is NP-hard and provided a polynomial-time approximation algorithm. If each node is allowed to multicast at most once, then our problem is a generalization of that in [7] to the multirate case. However, as we have argued in Section II, the multirate problem has a number of unique properties not present in the single-rate case.

The problem of routing in multirate multihop wireless networks has previously been studied in [2] and [6] but all of them focused on unicast routing. Assuming an infinite interference range, [2] shows that the unicast routing path that minimizes the total path delay will also maximizes the throughput between the source and destination. In order to deal with multirate links, [2] defines the medium-time metric (MTM) for each transmission rate. MTM essentially measures the time it takes to transmit a packet over a multirate link taking into account transmission delay, overheads of the RTS/CTS/ACK frames, and channel contention. Note that the inclusion of channel contention is needed to account for intraflow interference.

The rate adjustment problem at the MAC layer has been considered in, e.g., [10], but the work is focused on unicast, rather than multicast.

IV. MINIMUM BROADCAST LATENCY IN A MULTIRATE WMN

In this section, we formulate the problem of finding the optimal network-layer broadcast topology that results in the minimum broadcast latency in a multirate multihop WMN. Our modeling assumptions are as follows.

- 1) Each node is equipped with one radio, with all radios tuned to a common channel.
- 2) All nodes use the same transmission power for all transmission rates.
- 3) By adjusting the modulation scheme, a node can multicast at different data rates, with the transmission range a decreasing function of the data rate. Let s_{\max} denote the maximum transmission range. Also, while we use a disc model for the transmission range in our studies, our presented algorithms work with more generalized connectivity graphs.
- 4) A node’s neighbors are all the nodes that are reachable using the lowest possible transmission rate.
- 5) Let $\{i_1, \dots, i_k\}$ be a subset of the neighbors of a node x and the maximum (unicast) rates that node x can use to reach these nodes independently are r_1, \dots, r_k , respectively. The maximum rate that node x can use to multicast to i_1, \dots, i_k is $\min(r_1, \dots, r_k)$.
- 6) 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.³ We call κ the normalized interference range.
- 7) 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

³The interference range is independent of the transmission rate, since a higher rate implies not only a higher received signal power at a receiver, but also requires a correspondingly higher signal-to-interference-and-noise ratio (SINR) for packet reception.

multicast does not interfere with the intended recipients of node j 's multicast and vice versa.

- 8) We assume a centralized entity which schedules these multicasts so that, under the ideal MAC-layer assumption, no two multicasts will interfere with each other.
- 9) 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-layer broadcast is a sequence of link-layer multicasts instead of link-layer broadcasts. The use of link-layer multicasts is necessary especially when a node is to transmit the same message multiple times to different subset of neighbors, as illustrated in the motivating example in Section II.

Under the above modeling assumptions, the optimization objective is to minimize the broadcast latency, which is the time it takes all the nodes in the network to receive the packet. The key decisions in this optimization problem are: 1) whether a node should multicast and if so, how many times and to which of its neighbors and 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 in such a way that the minimum latency is achieved. Not surprisingly, this problem is NP-hard.

Theorem 1: The minimum latency network-layer broadcast problem with possibly multiple number of transmission per node in a multirate WMN is NP-hard.

Proof: A special case of our problem is the minimum broadcast latency problem in a single-rate WMN where each node can transmit at most once, which is NP-hard [7]. Therefore, our problem is NP-hard. \square

V. HEURISTIC ALGORITHMS FOR THE MINIMUM BROADCAST LATENCY PROBLEM IN MULTIRATE WMNS

In this section, we will present a heuristic algorithm to create "efficient" delivery trees for broadcast packets in a multirate WMN. Broadly speaking, any heuristic algorithm must make three important decisions. First, it has to decide whether a node should multicast. Second, the algorithm must decide the number of transmissions at each transmitting node and the neighboring nodes covered in each of these transmissions. Finally, the multicast transmissions of all nodes must be scheduled and their transmission time decided, while taking radio interference into account. It should be noted that these 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. With the hardness of the problem in mind, our algorithm is decomposed into three logically independent steps.

- 1) *Topology Construction:* The aim of this step is to compute a broadcast tree (or a spanning tree) T of the given multirate WMN, that exploits both the multirate nature of links and WMA. This step decides who the transmitting nodes

are (which are the non-leaf nodes of the tree) and the children/parent relation between different nodes. In this step, it is assumed that a node can transmit multiple times at different rates—the decision on the number of distinct-rate transmissions actually used at each node is deferred to the next step.

- 2) *Downstream Multicast Grouping:* The multicast grouping algorithm, which takes the broadcast tree T from the topology construction phase as the input, aims to determine the number of distinct-rate transmissions (where each corresponds to a separate link-layer multicast) each transmitting node should be making. For each distinct-rate transmission, this step also determines the transmission rate to use and the link-layer multicast recipients. Intuitively, the goal is to allow faster transmission to the more "critical" child nodes (i.e., those nodes that have leaf nodes with larger delivery latencies), at the expense of larger transmission latency to the other child nodes.
- 3) *Transmission Scheduling:* While we have decided the number of transmissions at each node on the tree, the exact timing of these transmissions (especially relative to different branches of the tree) still needs to be determined. The final step schedules all transmissions taking into account the fact that a node can only multicast after it has received the packet, and interfering multicast transmissions cannot occur concurrently. We are conceptually assuming a centralized scheduler in our current work, and shall investigate the more practical case of decentralized MAC scheduling in future work.

Clearly, this decomposition of the overall optimization problem is not optimal, e.g., it is only after the grouping phase that we obtain the multicast transmission sets, as well as the transmission rate associated with each link-layer multicast. However, as already noted, a joint optimization is computationally infeasible, except for trivially small mesh topologies.

We present in Section V-A a heuristic called weighted connected dominating set (WCDS) for the topology construction step. This is followed by the algorithmic description for the grouping and scheduling heuristics in, respectively, Sections V-B and V-C. Note that the grouping and scheduling heuristics can in fact work with any topology construction heuristic.

We first introduce some mathematical notation. The WMN 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} . 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 slowest rate.

A. Topology Construction Algorithm: WCDS

In this section, we present an algorithm based on the concept of WCDS to compute a broadcast tree rooted at source node s . Recall that for a graph $G = (V, E)$, a connected dominating set (CDS) Z of G is a subset of V such that: a) every element (node) of $V \setminus Z$ is in the neighborhood of at least one node in Z and b) the set Z is connected. Among all CDSs of G , the

one with minimum cardinality is the minimum connected dominating set (MCDS). Computing an MCDS in a unit graph is NP-hard [8]. The use of MCDS to achieve optimal flooding in a single-rate multihop wireless networks has been explored in [12] where the authors prove that the size of the optimal flooding tree (measured by the number of nodes performing broadcasts, not by broadcast latency) differs from the size of the MCDS by at most one. However, MCDS performs poorly in multirate environments because it does not account for multirate links in the tree construction. To extend MCDS to our multirate setting, we assume the network has k different rates r_1, r_2, \dots, r_k . Let $N(x, r_i)$ denote the nodes that are reachable from node $x \in V$ using rate r_i . We define the minimum WCDS problem whose aim is to find a subset $Y = \{y_1, y_2, \dots\}$ in V and the broadcast rate w_i (which are chosen from r_1, r_2, \dots, r_k) for node $y_i \in Y$ such that

- 1) Every element of $V \setminus Y$ is in $\cup_{y_i \in Y} N(y_i, w_i)$.
- 2) The set Y is connected.
- 3) The weighted sum $\sum_{y_i \in Y} (1)/(w_i)$ is minimal.

Note that when there is only one transmission rate, the minimum WCDS is equivalent to the MCDS. We expect the solution to the minimum WCDS problem to be similar to optimal broadcast tree for the multirate scenario. We use a greedy algorithm, depicted in Algorithm 1, to obtain an approximation of the minimum WCDS. The algorithm starts by making the source node s eligible to transmit. We say that a node is covered if it has already received a packet; the set C tracks the progressively larger set of covered nodes. Also, the set R is the set of all k transmission rates. For an eligible node c and rate $r \in R$, the quantity $|N(c, r) \setminus C|$ is the number of “not-yet-covered nodes” that are reachable by a broadcast by node c at rate r . Thus, in each round of the algorithm, we choose the (c, r) combination that maximizes the rate of increase of not-yet-covered nodes, as measured by $|N(c, r) \setminus C| \times r$. This metric reflects our desire to include as many nodes as possible in a single transmission, yet keep the transmission rate high (even though a higher transmission rate implies a smaller range, and thus, a smaller set of covered nodes). The algorithm returns T , the set of directed links in the broadcast tree.

Algorithm 1: WCDS Tree Construction

```

1: Input:  $G, s, R = \{r_1, \dots, r_k\}$ 
2:  $C = \{s\}, T = \emptyset$ 
3: while  $(V \setminus C \neq \emptyset)$ 
4:   For  $(c \in C)$  do
5:     For  $(r \in R)$  do
6:        $f(c, r) = |N(c, r) \setminus C| \times r$ 
7:     end for
8:   end for
9:    $(\hat{c}, \hat{r}) = \arg \max_{c \in C, r \in R} f(c, r)$ 
10:   $A \leftarrow N(\hat{c}, \hat{r}) \setminus C$ 
11:   $C \leftarrow C \cup A$ 
12:   $T \leftarrow T \cup (\cup_{a \in A} \{(\hat{c}, a)\})$ 
13: end while

```

B. Multicast Grouping Algorithm

After the broadcast tree is constructed, we must now decide on the number of times a transmitting node (i.e., non-leaf node of the broadcast tree) will multicast, as well as the recipients of each such link-layer multicast. If a node multicasts only once, all its child nodes must receive it. If a node is to multicast more than once, a different subset of child nodes will be reached in each multicast such that these subsets together form a partition of the set of child nodes.

Due to the complexity of the multicast grouping algorithm and space limitation, we can only give a brief description here, referring the reader to [5] for details. Since the transmission decision (i.e., how many times a node should transmit and to whom) to be made at a transmitting node depends on what happens downstream, the algorithm proceeds in a bottom-up manner. When the transmission decision at a transmitting node has been made, we are able to estimate the time it takes a packet to travel from that node to all its descendants. (Note that this estimation is done by ignoring the possible interference between different branches of the tree). For ease of reference, we call this time the *cardinal value* (CV) of a transmitting node. For a transmitting node whose downstream transmitting nodes have already made their transmission decision (i.e., their CVs have already been determined), its CV can be determined from those of its downstream transmitting nodes. The CVs of the transmitting nodes will also be used in determining the transmission schedule in the next algorithmic stage.

In order to determine the CV at a transmitting node n , we go through all the possible valid transmission sequences (VTSs) at node n to see which one will give the shortest possible time to send a packet from node n to all its descendants. We illustrate the concept of VTS by an example. Consider a transmitting node n which has two children c_1 and c_2 , which can be reached using a minimum latency of $d_1 (= 1)$ and $d_2 (= 2)$ time units, respectively. Node n can reach these nodes in a number of VTSs, e.g., it can first multicast to c_1 (latency 1) followed by another multicast to c_2 (latency 2). We denote this VTS as (d_1, d_2) . Another VTS is (d_2) which reaches both nodes in one multicast. These two are all the VTSs for this example. The sequence (d_1) is invalid because it does not reach all the child nodes. Also, (d_2, d_1) is invalid because the second transmission is unnecessary since both nodes are already reached by transmission d_2 whose coverage area is greater. We can readily show that if a node connects to its downstream nodes using k distinct transmission rates, then there are 2^{k-1} possible VTSs.

We now describe how to compute the time it takes to send a packet from a transmitting node to all its descendants for a given VTS by considering node n is using VTS (d_1, d_2, \dots) . We assume the transmissions will proceed as follows: Node n first transmits at latency d_1 reaching its downstream nodes, denoted by $N(n, d_1)$. If some of the nodes in $N(n, d_1)$ are transmitting nodes, they will then begin their transmissions to their respective downstream neighbors in parallel. (Note that we are ignoring interbranch interference here). Following from the definition of CV, the time it takes these transmissions from $N(n, d_1)$ to complete are their CV values. A particular assumption we make is that node n does not begin transmitting at latency d_2 immediately after finishing transmitting at d_1 . Rather, we assume that

node n waits until all the transmissions from $N(n, d_1)$ and their descendants have proceeded sufficiently so that the d_2 -transmission of n does not interfere with those of $N(n, d_1)$ and their descendants. This waiting time can be estimated from the CV values of the downstream nodes. This operation then repeats itself until all transmissions in the VTS have been made. Given this *modus operandi*, we can estimate the time to reach all the descendants of a node for a given VTS.

C. The Scheduling of Transmissions

After both topology construction and multicast grouping have been done, we know all the multicasts transmissions that have to be performed, *except* their timing. We approach the scheduling problem by formulating it 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 $V_b = \{b_1, b_2, \dots, b_k\}$ be the set of all the multicast transmissions decided by the multicast grouping algorithm in Section V-B. Each multicast transmission b_i have four attributes: 1) A sender (which is a non-leaf node of the broadcast tree). 2) A group of recipients (which is a subset of the child nodes of the sender in the broadcast tree). 3) The latency required by the transmission, denoted by $t(b_i)$, which is the minimum latency it takes the sender to reach all its designated recipients. 4) The CV value of a transmission which is the estimated time it takes a packet to travel from the sender to all its descendants. Since the CV value of transmission measures the time it takes a packet to reach the end of the tree, it is viewed as an urgency measure by the scheduling algorithm. In addition, we 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 1) the multicast of b_i interferes with the reception of the recipients in b_j or *vice versa* or 2) both multicasts b_i and b_j have the same sender. (This is a generalization of the conflict graph defined in [11] for the unicast case.)

Formally, a schedule can be defined as a mapping $\tau : V_b \rightarrow \mathbb{R}$ which gives the transmission starting time of $b_i \in V_b$. A valid schedule is one which meets the following constraints.

- 1) The source node multicasts at time zero.
- 2) A node can only multicast after it has received the packet: if the sender of b_j is a recipient of b_i , then $\tau(b_j) \geq \tau(b_i) + t(b_i)$
- 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 we use the same notation to denote both an open interval and an edge of a graph, the usage should be clear from the context.

The scheduling algorithm is depicted in Algorithm 2. The input to the algorithm is transmissions information (TX) which contains the attributes discussed earlier. The aim of the scheduling algorithm is to find out the starting time (τ) and ending time (δ) of all transmissions at each transmitting node. Initially, *time* depicting current running time is initialized to zero and E depicting *eligible transmissions* is initialized with all transmissions of the source node. A transmission is said to be eligible when the node performing this transmission receives the multicast from its parent, all transmissions of the source node are eligible at time zero. The scheduling process starts by scheduling the transmission with the largest CV value at the source node at

time zero. This transmission is added to the set T which contains all transmissions currently being performed. The starting time (τ) and ending time (δ) of transmissions are decided as they are added to T . The minimum of $\delta(\forall t \in T)$ is the earliest any transmission in T will finish and also the earliest a waiting eligible transmission can be scheduled and is called the next-stop time.

At the next-stop time, since the channel becomes available again due to completion of some transmission, a new transmission must be slotted for transmission. The transmission $t \in E$ having the maximum transmission CV is determined, and is assumed to be more “critical” as it connects to subtrees of higher broadcast delay. Thereafter, it is checked that t does not interfere with any of the transmissions in T . In case of no interference, t is added onto T and deleted from E . The starting time $\tau(t)$ and ending time $\delta(t)$ for the transmission t are also decided at this time. However, in case t interferes with any existing transmissions in T , it is held back until next-stop time. It is also ensured that a high-rate transmission does not follow a low-rate transmission at the same node.

After we have iterated through all eligible transmissions, i.e., all $t \in E$, the next-stop time is found by determining which transmission is going to finish the earliest. At the next-stop interval, the child nodes of the transmission finishing at next-stop interval receive the message, and thus are eligible for transmitting. Thus, at next-stop interval, the transmissions of these recently eligible nodes are added to the eligible transmissions E alongside those transmissions which were held back in the last round. We abide by the precedence constraint in this manner, i.e., by allowing a transmission to be added to E only after the transmission has been enabled where a transmission is said to be enabled when the node making the transmission has received from its parent. At the next-stop interval, all transmissions which are finishing are deleted from T . The algorithm runs in rounds and finishes when the starting time for all transmissions $\tau(\forall t \in V_b)$ and ending time for all transmissions $\delta(\forall t \in V_b)$ have been decided.

Algorithm 2: Scheduling

- 1: **Input:** TX
- 2: Set $\text{time} = 0$
- 3: Initialize $E \leftarrow \cup_{TX.\text{Node}=n_s} \{TX\}$
- 4: Initialize $T = \emptyset$
- 5: **while** ($E \neq \emptyset$ or $T \neq \emptyset$) **do**
- 6: **while** $E \neq \emptyset$ **do**
- 7: $t = \arg \max_{TX \in E} TX.CV$
- 8: $E = \{E \setminus t\}$
- 9: **if** $|T| \geq 1$ **then**
- 10: **if** $TX(t).node \ \& \ TX(\cup_{t' \in T \setminus \{t\}} \{t'\}).node$
do not interfere **then**
- 11: $T \leftarrow \{T \cup t\}$;
- 12: Set $\tau(t) = \text{time}$
- 13: Set $\delta(t) = \text{time} + TX(t).latency$
- 14: **else**

```

15:      $E_{\text{Next}} \leftarrow t$ 
16:     end if
17:     else if  $|T| < 1$  then
18:          $T \leftarrow \{T \cup t\}$ ;
19:         Set  $\tau(t) = \text{time}$ 
20:         Set  $\delta(t) = \text{time} + TX(t).\text{latency}$ 
21:     end if
22: end while
23:  $NextStop = \min(\delta(\cup_{t \in T}\{t\}))$ 
24:  $NextTrans = \{t : (\forall t)(\delta(t) = NextStop)$ 
25:  $E \leftarrow E \cup t_{\text{children}} \text{ of } NextTrans$ 
26:  $T = T - NextTrans$ 
27:  $E = E - NextTrans$ 
28:  $E = E \cup E_{\text{Next}}$ 
29:  $\text{time} \leftarrow NextStop$ 
30: end while
31: Output:  $\tau(t), \delta(t) \forall (1 \leq t \leq |TX|)$ 
    
```

D. Maximum End-to-End Throughput

The above discussion of the tree construction and scheduling algorithms focused on 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. For higher rate data flows, it is important to compute the maximum achievable throughput of a broadcast tree which utilizes the schedule computed in Section V-C. We will first define the meaning of maximum end-to-end throughput being used here.

Using the same notation as in Section V-C, the set of all multicast transmissions are $V_b = \{b_1, b_2, \dots, b_k\}$ and the schedule says that transmission b_i will take place during the time interval $[\tau(b_i), \tau(b_i) + t(b_i)]$. Assuming that packets are generated by the source node at regular time at $(m-1)\Delta$ (for $m = 1, 2, \dots$). Our goal is to maintain the same schedule computed earlier so that the transmitting node of multicast transmission b_i is expected to multicast the m th packet during $[(m-1)\Delta + \tau(b_i), (m-1)\Delta + \tau(b_i) + t(b_i)]$. The maximum throughput is achieved by the smallest possible Δ such that there is no conflict between the scheduling of all the packets. By defining

$$\mathcal{I}(m, b_i) = ((m-1)\Delta + \tau(b_i), (m-1)\Delta + \tau(b_i) + t(b_i))$$

we can formally express the above problem as

$$(\mathbf{P1}) : \min_{\Delta > 0} \Delta \text{ subject to}$$

$$\mathcal{I}(m_1, b_i) \cap \mathcal{I}(m_2, b_j) = \emptyset \forall m_1, m_2 = 1, 2, \dots \text{ if } (b_i, b_j) \in G_c$$

where G_c is the conflict graph defined in Section V-C. Since the schedules repeat themselves periodically, it is sufficient to examine possible conflicts in $[0, T_{\max}]$, where T_{\max} is the broad-

cast latency. Thus, Problem **(P1)** can alternatively be expressed as

$$(\mathbf{P2}) : \min_{\Delta \in [0, T_{\max}]} \Delta \text{ subject to}$$

$$\mathcal{I}(1, b_i) \cap \mathcal{I}(m, b_j) = \emptyset \forall m = 1, 2, \dots \text{ if } (b_i, b_j) \in G_c.$$

Assuming two transmissions b_i and b_j do interfere with each other, the constraint in Problem **(P2)** can alternatively be expressed as

$$(m-1)\Delta + \tau(b_j) \notin (\tau(b_i) - t(b_j), \tau(b_i) + t(b_i)).$$

The left-hand side of the above expression is the start transmission time of the m th packet by transmission b_j and it must not lie in the time interval given on the right-hand side in order to avoid conflict. This means that Δ cannot take certain values. Thus, by identifying all the values that Δ cannot take within $[0, T_{\max}]$, we can easily find the optimal value of Δ . This algorithm can find the optimal Δ in polynomial time.

VI. SIMULATED PERFORMANCE STUDIES

In this section, we study the performance of the algorithm proposed in Section V to solve the low-latency network-wide broadcast problem in a multirate WMN. For the purpose of comparison, we will study altogether four heuristics. All these four heuristics have the same structure: computing a broadcast tree, and then followed by the multicast grouping (Section V-B) and transmission scheduling (Section V-C). In other words, these algorithms only differ in how the broadcast tree are computed. The algorithms to be considered are the following.

- 1) Algorithm WCDS: Uses WCDS in Section V-A to compute the broadcast tree.
- 2) Algorithm BIB: Broadcast incremental bandwidth (BIB) (only the tree formation part) was proposed earlier by us in [4]; see [5] for details.
- 3) Algorithm SPT: The broadcast tree is the shortest path tree (SPT) computed by Dijkstra's algorithm.
- 4) Algorithm CDS: This heuristic assumes that all broadcasts are done at the lowest transmission rate. The broadcast tree can be computed by using WCDS in Section V-A with only the lowest rate allowed.

We compare the performance of these four heuristics using random topologies of different network sizes (measured by the number of nodes) and network area (which is the area over which the nodes are distributed, assuming to be a square of ℓ^2 km²). For each (network size, network area) combination, we generate 100 topologies whose nodes are uniformly randomly distributed in the network area. We then apply our algorithm in Section V to each topology to compute the broadcast latency and end-to-end throughput. We normalized the broadcast latency 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 normalized delay is unity. The result that we will show is the geometric mean, over 100 network instances, of the normalized delay and the throughput. For the simulation, the

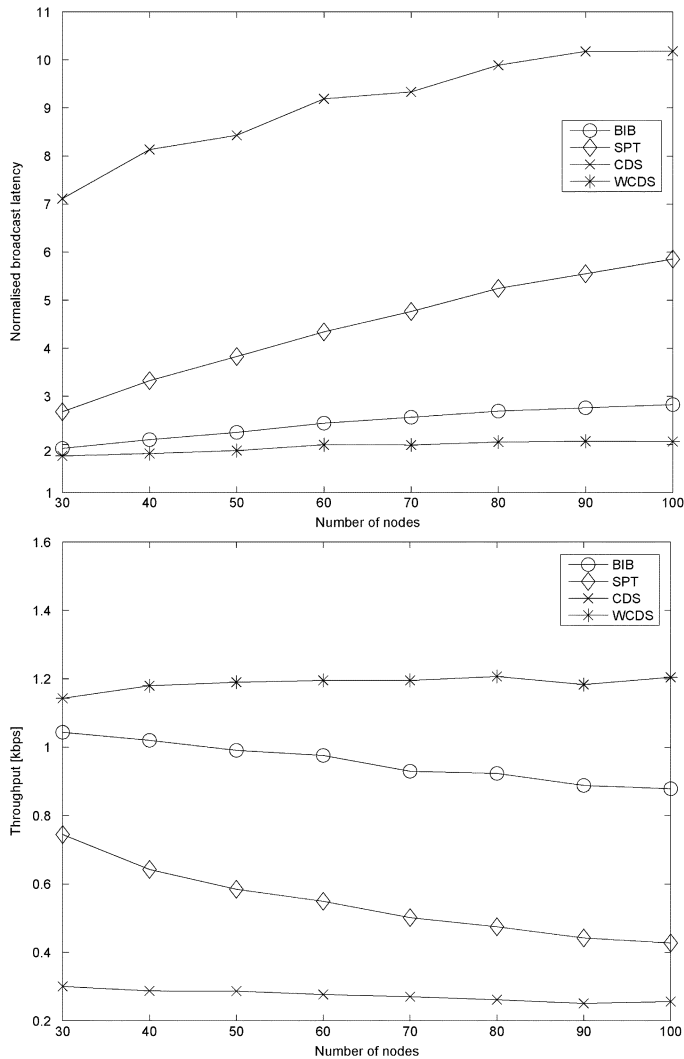


Fig. 4. This figure shows the geometric mean of the normalized latency (top graph) and throughput (bottom graph) of BIB, WCDS, SPT, and CDS.

rate-range curve in Table I is used. Our modeling assumption in Section IV states that the interference range is κ times of the transmission range of the lowest transmission rate. Unless otherwise specified, κ is 1.7 which is identical to that used in [21].

A. Single Transmission Case

We first consider the case where we impose the limitation that each node can transmit a packet at most once. For the simulation, we set $\ell = 1.5$ and vary the number of nodes from 30 to 100. The normalized delay and throughput are given in Fig. 4. It turns out that good performance for delay also means good performance for throughput and *vice versa*. WCDS performs best, followed by BIB, SPT, and CDS. The results show that both BIB and WCDS are able to exploit the multirate capability. The SPT algorithm fails to exploit the WMA and results in higher latency and lower throughput, this can be seen from Fig. 5 which shows that SPT on average uses the most number of multicasts per tree out of the four 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. Note that the broadcast latency and

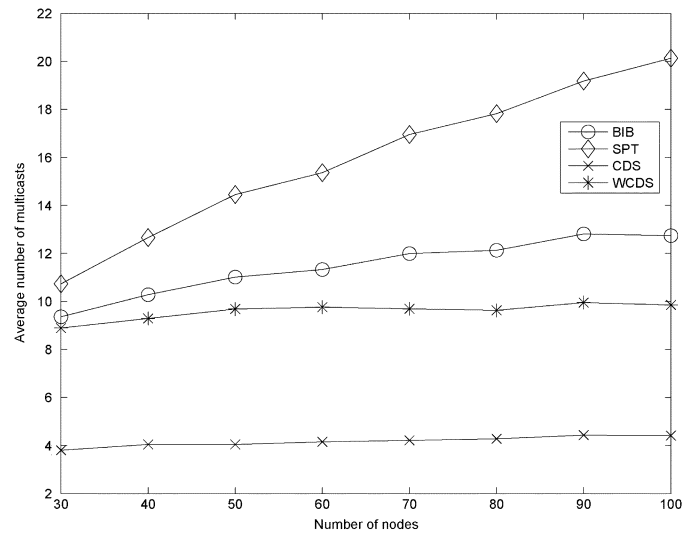


Fig. 5. Graph shows the mean number of multicasts per tree for BIB, WCDS, SPT, and CDS.

throughput of WCDS is 3–5 times lower than that of CDS, suggesting that the use of rate diversity in the broadcasting process can result in dramatic decrease in latency. Note also that the normalized broadcast latency for WCDS in Fig. 4 is relatively independent of the number of nodes. This is expected, since all network nodes will received their packets when the broadcasts from the transmitting nodes cover the entire network area. We also study the sensitivity of the results to the value of interference range and find that interference range has only a small effect on our results. See [5] for details.

B. Multiple Transmission Case

We now consider the case when the nodes are allowed to transmit the same packet multiple times but at different rates. We run our simulation in the same manner as in Section VI-A but we do not limit the number of times a node can transmit the same packet. Due to space limitation, we can only provide a key summary here, see [5] for details. We find that, over 100 random topologies of fixed number of nodes in a fixed area, multiple transmission do not significantly reduce the broadcast latency. Moreover, multiple transmissions were invoked by a fairly small number of topologies, e.g., for the WCDS algorithm, only 2 out of 100 topologies for a network area of 1 km² required multiple transmissions, and these resulted in a 10% reduction in broadcast latency. It appears that multiple transmission may not be required in the single-radio single-channel scenario. However, multiple transmissions is likely to be more useful in the multiradio multichannel environment, where the transmissions on different radio interfaces can proceed in parallel. From Fig. 4, we see that single transmission can result in a normalized broadcast latency of about 2. This means that the potential improvement offered by multiradio multichannel for latency reduction is still large, and should be investigated further.

VII. FUNDAMENTAL DESIGN PRINCIPLES OF BROADCAST IN MULTIRATE MESHES

In Section VI, we studied the performance of the heuristics using the transmission rate-transmission range characteristics

TABLE II
RATE-RANGE CURVE OF A HYPOTHETICAL MULTIRATE SYSTEM
PARAMETERIZED BY $\gamma < 1$ AND $\rho > 1$

Transmission rate (Mbps)	Maximum transmission range (m)
r_0	d_0
ρr_0	γd_0
$\rho^2 r_0$	$\gamma^2 d_0$
$\rho^3 r_0$	$\gamma^3 d_0$

(or rate-range curve for short) given in Table I and saw that multirate multicast using WCDS resulted in the lowest broadcast latency compared with CDS (which always multicasts at the lowest rate). In this section, we will study the sensitivity of this result to the choice of rate-range curves. The result of this investigation can help us to answer a number of fundamental design questions: 1) Given a multirate system with n different rates, is it necessary to use all the n different rates? 2) If not, which of the n different rates should we use and what is an efficient method to decide that?

A. The Transmission Rate-Transmission Coverage Area Product

In order to study the effect of rate-range curves on the broadcast latency, we use a family of hypothetical rate-range curves, as given in Table II. Our hypothetical system has a minimum transmission rate of r_0 Mb/s whose transmission range is d_0 m. Each subsequent transmission rate is a factor of $\rho (> 1)$ greater but whose transmission range is a factor of $\gamma (< 1)$ smaller.

Let us assume for the time being $\gamma = (1/2)$. Consider the transmission of a frame of size p bits. If the lowest rate is used, this packet will reach all nodes in an area of πd_0^2 in a time of $(p)/(r_0)$. However, if this is to be transmitted using the second lowest rate $r_1 = \rho r_0$, then each transmission will only cover an area of $(1/4)\pi d_0^2$ requiring a shorter time of $(p)/(\rho r_0)$ for each transmission. Therefore, four transmissions at rate r_1 can cover the same area as one transmission at rate r_0 . Furthermore, in the worst case where these four transmissions at rate r_1 are within the interference range of each other, then they can only take place one after the other and this will take a total time of $4(p)/(\rho r_0)$ to complete. Thus, if $\rho > 4$, it will be always be more efficient to transmit at rate r_1 .

Generalizing the argument used in the last paragraph, we propose to use the product of transmission rate and transmission coverage area (or rate-area product or RAP for short) as a measure of efficiency of a certain transmission rate. Thus, with the hypothetical system given in Table II, it will be more efficient to use the higher rate if $\gamma^2 \rho > 1$, otherwise, the lowest rate should be used instead. Alternatively, a transmission rate with a higher RAP is more efficient for broadcast. In order to verify this conjecture, we perform a number of simulations using the same method as in Section VI except that the rate-range curve in Table II is used.

In this set of simulations, $\rho = 2, r_0 = 1$, and $d = 500$. If the above conjecture holds, we expect that it will be more efficient to use the higher transmission rates if $\gamma > (1)/(\sqrt{\rho}) = 0.72$. Five different values of $\gamma = 0.5, 0.6, 0.7, 0.8$ and 0.9 are used. Only the WCDS and CDS heuristics are studied. We normalized the delay by using those of CDS. The normalized delay of WCDS is given in Fig. 6. It shows that WCDS gives a better latency than

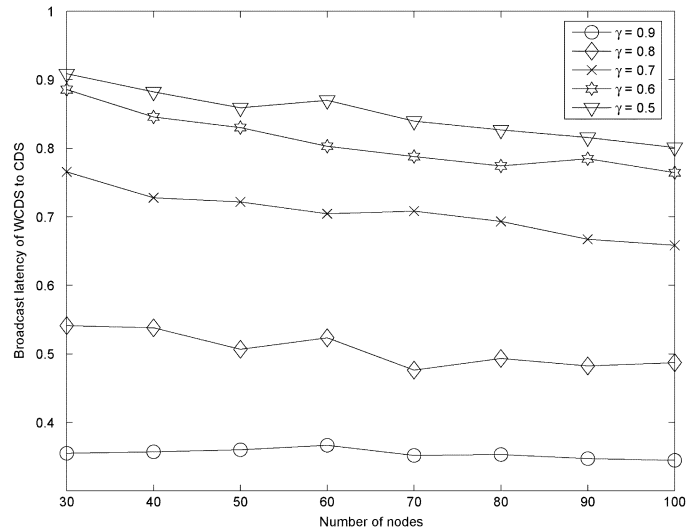


Fig. 6. Graph shows the geometric mean of the latency of WCDS to CDS over 100 randomly generated topologies of each network size.

TABLE III
AVERAGE PERCENTAGE OF TIMES EACH TRANSMISSION RATE IS USED
FOR DIFFERENT VALUE OF γ . THE PARAMETER $\rho = 2$,
THUS THE RATES DECREASES FROM LEFT TO RIGHT

γ	% of times each transmission rate is used			
	Rate $\rho^3 r_0$	Rate $\rho^2 r_0$	Rate ρr_0	Rate r_0
0.9	100	0	0	0
0.8	96	4	0	0
0.7	70	18	6	6
0.6	41	27	12	20
0.5	14	28	21	37

CDS for all values of γ . For $\gamma \geq 0.7$, WCDS exploits multirate and gives far better delay than CDS; for $\gamma < 0.7$, WCDS still performs better than CDS but the results are comparable. (We have also investigated the behavior of end-to-end throughput. For $\gamma < 0.7$, end-to-end throughput of WCDS is similar to that of CDS but for $\gamma > 0.7$, WCDS has better throughput; see [5] for details). We can understand this by looking at the average percentage of times that each transmission rate is used for each value of γ given in Table III. It shows that if the rate-range curve is favorable, then the higher transmission rates are used most of the time. However, even when the rate-range curve is less favorable, the higher rate transmissions are also used but less often. These observations therefore confirm our earlier conjecture. We have also used $\rho = 1.5$ and the results are similar, see [5] for details. We also found that the results are not sensitive to the value of the interference range, see [5].

Under the condition that network connectivity is not affected by using higher-rate multicast transmissions (which have shorter transmission range), the above results show that *the case for exploiting the higher transmission rates to reduce broadcast latency depends on the RAP of the transmission rates*. If the higher transmission rates have larger RAPs compare with the lowest rate, then using multirate link-layer multicasts can result in significant reduction in broadcast latency (provided, of course, that this does not affect the network connectivity). Applying this rule-of-thumb to the rate-range curve of 802.11b in Table I, it can be seen from the last column of the table that

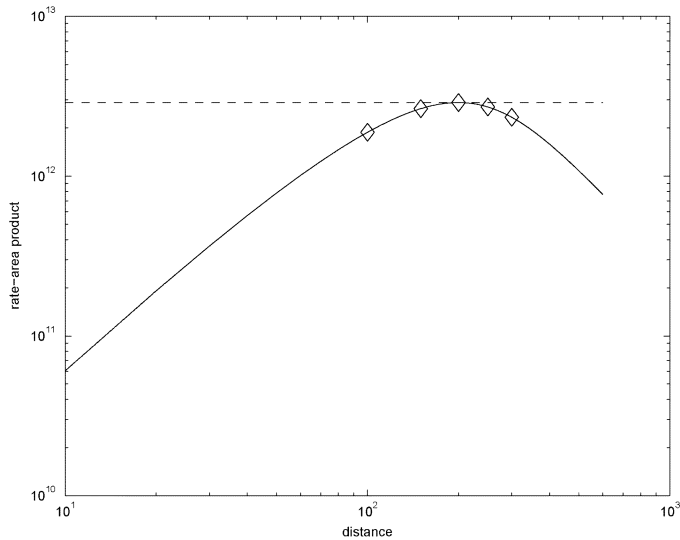


Fig. 7. Relation between distance and RAP with Shannon rate.

the RAP is larger for higher transmission rates and this agrees with the results in Section VI.

B. Channel Capacity and MultiRate Networks

In Section VII-A, we demonstrate that transmission rates with large RAP are good for achieving low broadcast latency. With improvement in coding, wireless signal processing, etc., the achievable wireless transmission rate is pushing closer to the Shannon capacity. An interesting question is to study the RAP if the transmission rate at a distance is given by the Shannon capacity. We consider a system where the bandwidth $B = 10$ MHz, the SNR at distance $d_0 = 50$ m is 30 dB. (We will see later that these parameter values will not affect the general discussion here). Assuming that the rate R at distance d is given by the Shannon capacity formula, as follows:

$$R = B \log_2 \left(1 + \text{SNR} \left(\frac{d_0}{d} \right)^\theta \right) \quad (1)$$

where θ is the path loss exponent. Assuming that $\theta = 4$, Fig. 7 shows RAP as a function of d where a large d corresponds to a small R and *vice versa*. It shows that the RAP increases for small values of d and decreases for large d . This is understandable since for small d , $R \sim \log_2(1/d)$ and for large d , $R \sim (1/d^\theta)$. It can be shown, via differentiating $R\pi d^2$, that the transmission rate (whose corresponding spectral efficiency is ψ) that maximizes the RAP is the solution to the equation $\psi - (\theta \log_2 e/2)(1 - 2^{-\psi})$. This shows that *the optimal ψ is a function of the path loss exponent θ only and not of other parameters*. For $\theta = 4$, the maximum RAP (indicated by the dashed lines in Fig. 7) occurs around a spectral efficiency of 2.3 b/s/Hz. The lowest transmission rates for both 802.11a/b has a spectral efficiency far lower than this, and therefore have poor RAP. By adding higher transmission rates with better RAP to 802.11b (see Table I), the broadcast latency of 802.11b is improved as seen in Section VI. However, the Shannon RAP predicts that RAP will eventually fall for higher transmission rates. From the technical specifications of a commercial 802.11b/g product in [1], we find that the outdoor transmission ranges for rates 1, 6,

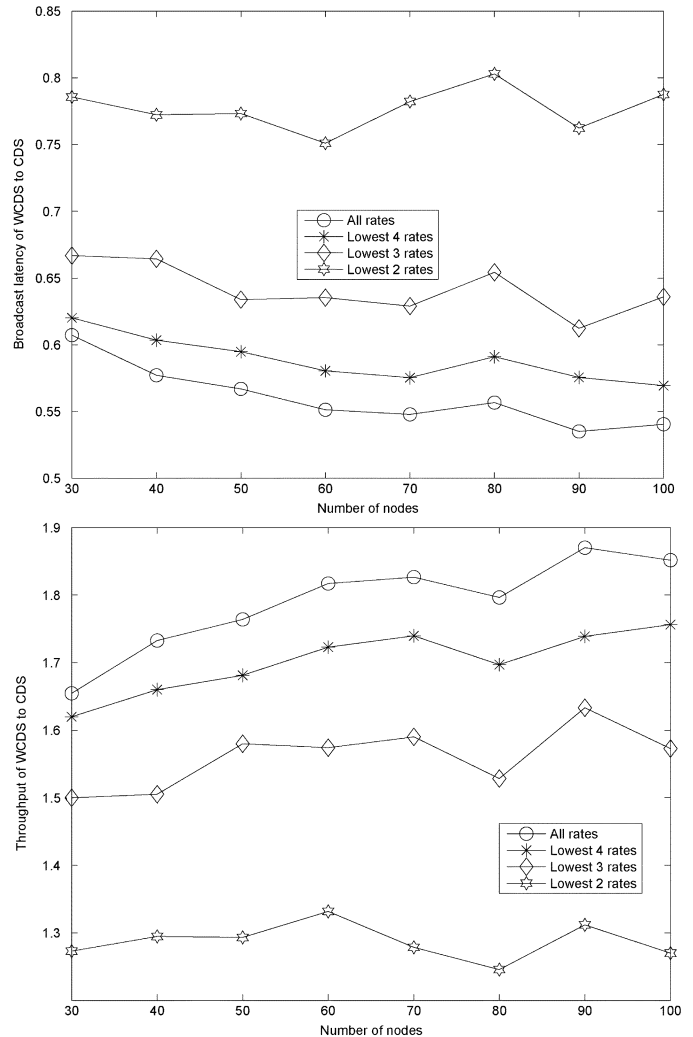


Fig. 8. Performance of WCDS algorithm (relative to CDS) using different number of transmission rates. The top graph shows latency, while the bottom one shows throughput.

11, 18, and 54 Mb/s are, respectively, 610, 396, 304, 183, and 76 m, giving RAP of 1.2, 3.0, 3.2, 1.9, and 1.0 Mb/s-km², which eventually falls for high transmission rates.

We assume a hypothetical multirate system by selecting five points from the Shannon rate-range curve indicated by the diamonds in Fig. 7. Since it is likely that future wireless systems will have rates with efficiency above and below 2.3 b/s/Hz, the rate that gives the maximum RAP is selected, as well as two points on each side of it. (Note also that the Shannon transmission rate can only be used if no other nodes are transmitting, or in other words, the interference range is infinity. Since we find that in the last section that the interference range has little impact on the result, we keep the normalized interference range as 1.7 as before). We use the same simulation set up as in Section VII-A except that we use the following five algorithms: WCDS with all the five rates, and the lowest 4, 3, 2, and 1 rate. Note that the last algorithm is in fact CDS with the lowest rate. We normalize the results for the various WCDS algorithms using those from CDS. The results are in Fig. 8. It can be seen that the best results are given by WCDS using all the five rates, thus again confirming that multirate is useful for reducing broadcast latency. Since the

third rate has the highest RAP, note that there is sizeable performance gap between using the lowest two rates and the lowest three rates.

VIII. CONCLUSION

We propose the novel concept of multirate link-layer multicast as a way to introduce low-latency network-layer multimedia broadcast (or multicast) in a WMN. We show that by exploiting both multiple transmission rates and WMA, we can get significant reduction in broadcast latency compared with using the lowest rate alone. For example, based on simulations using typical IEEE 802.11-based values, the use of our rate-aware WCDS heuristic results in a threefold to fivefold reduction in the broadcast latency compared with the CDS algorithm that always performs link broadcasts at the lowest rate. Moreover, at least for a single-channel, single-radio environment, it is more important to exploit the rate diversity than allow each individual node to engage in multiple transmissions—we conjecture that this will change when multiple radios are present on a mesh node.

In addition, we find that the efficiency of a particular transmission rate for reducing broadcast latency can be predicted by the product of the transmission rate and its transmission coverage area. This provides a rule-of-thumb that the designer of a multirate system can use to determine which transmission rates should be included in a multirate system. Investigation of theoretical Shannon limits suggest that the case for using at least a small subset of the available choice of rates for link-layer multicasts will become even more compelling, as better modulation and coding techniques are introduced.

REFERENCES

- [1] *Cisco Aironet 802.11a/b/g Wireless LAN Client Adapters (CB21AG and PI21AG) Installation and Configuration Guide*, 802.11a/b/g Standard, Cisco Systems, 2004.
- [2] B. Awerbuch, D. Holmer, and H. Rubens, "High throughput route selection in multi-rate ad hoc wireless networks," *Lecture Notes in Computer Science*, vol. 2928, pp. 253–270, 2004.
- [3] P. Chaporkar, A. Bhat, and S. Sarkar, "An adaptive strategy for maximizing throughput in MAC layer wireless multicast," in *Proc. Mobihoc*, 2004, pp. 256–267.
- [4] C. Chou and A. Misra, "Low latency multimedia broadcast in multi-rate wireless meshes," in *Proc. IEEE Workshop on Wireless Mesh Networks*, 2005, pp. 54–63.
- [5] C. Chou, A. Misra, and J. Qadir, "Low latency broadcast in multi-rate wireless mesh networks," UNSW, Tech. Rep. TR 0514, 2005. [Online]. Available: <ftp://ftp.cse.unsw.edu.au/pub/doc/papers/UNSW/0514.pdf>
- [6] R. Braves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *Proc. MobiCom*, 2004, pp. 114–128.
- [7] R. Gandhi, S. Parthasarathy, and A. Mishra, "Minimizing broadcast latency and redundancy in ad hoc networks," in *Proc. Mobihoc*, 2003, pp. 222–232.
- [8] S. Guha and S. Khuller, "Approximate algorithms for connected dominating sets," *Algorithmica*, vol. 20, pp. 347–387, 1998.
- [9] C. Gui and P. Mohapatra, "Scalable multicasting in mobile ad hoc networks," in *Proc. INFOCOM*, 2004, pp. 2119–2129.
- [10] G. Holland, N. Vaidya, and B. P., "A rate-adaptive MAC protocol for multi-hop wireless networks," in *Proc. MobiCom*, 2001, pp. 236–251.
- [11] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," in *Proc. MobiCom*, 2003, pp. 66–80.
- [12] H. Lim and C. Kim, "Flooding in wireless ad hoc networks," *Comput. Commun.*, vol. 24, no. 3–4, p. 353, 2001.
- [13] W. Lou and J. Wu, "On reducing broadcast redundancy in ad hoc wireless networks," *IEEE Trans. Mobile Comput.*, vol. 1, p. 111, 2002.

- [14] J. Qadir, A. Misra, and C. Chou, "Low latency broadcast in multi-radio multi-channel multi-rate wireless mesh networks," UNSW, Tech. Rep. TR 0608, 2006. [Online]. Available: <ftp://ftp.cse.unsw.edu.au/pub/doc/papers/UNSW/0608.pdf>
- [15] A. Raniwala and T. Chueh, "Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network," in *Proc. INFOCOM*, 2005, pp. 2223–2234.
- [16] T. Rappaport, *Wireless Communications*. Englewood Cliffs, NJ: Prentice-Hall, 1999.
- [17] Scalable Networks Inc. [Online]. Available: <http://www.scalable-networks.com/>
- [18] Y. Tseng, S. Ni, Y. Chen, and J. Sheu, "The broadcast storm problem in a mobile ad hoc network," *Wireless Network*, vol. 8, no. 2/3, pp. 153–167, 2002.
- [19] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides, "On the construction of energy-efficient broadcast and multicast trees in wireless networks," in *Proc. INFOCOM*, 2000, pp. 585–594.
- [20] Y. Wu, "Network planning in wireless ad hoc networks: A cross-layer approach," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 136–150, Jan. 2005.
- [21] K. Xu, M. Gerla, and S. Bae, "Effectiveness of RTS/CTS handshake in IEEE 802.11 based ad hoc networks," *Ad Hoc Networks*, vol. 1, no. 1, p. 107, 2003.



Chun Tung Chou (M'88) received the B.A. degree in engineering science from the University of Oxford, Oxford, U.K., and the Ph.D. degree in control engineering from the University of Cambridge, Cambridge, U.K.

He is a Senior Lecturer and the head of the Network Research Group at the School of Computer Science and Engineering, University of New South Wales, Sydney, Australia. He has published extensively in computer networking and system identification. His current research interests are in

wireless mesh networks, wireless sensor networks, and network optimization.



Archan Misra (M'02) received the B.Tech. degree in electronics and communication engineering from the Indian Institute of Technology (IIT), Kharagpur, in 1993, and the Ph.D. degree in electrical and computer engineering from the University of Maryland, College Park, in 2000.

He is a Research Staff Member at the IBM T. J. Watson Research Center, Hawthorne, NY, where he works on infrastructural components for pervasive networking and middleware for presence-based, context-aware computing. He has published extensively in the areas of wireless networking and pervasive computing. His current research projects include middleware for SIP-based "rich presence" in converged networks, high-performance wireless mesh networks and sensor data management techniques for healthcare and remote monitoring scenarios.

Dr. Misra is a coauthor on papers that received the Best Paper Awards at ACM WOWMOM 2002 and IEEE MILCOM 2001. He is presently on the Editorial Board of the *IEEE Wireless Communications Magazine*, and currently chairs the IEEE Computer Society's Technical Committee on Computer Communications (TCCC).



Junaid Qadir received the B.E. degree in electrical engineering from the University of Engineering and Technology, Taxila, Pakistan, in 2000 and the M.S. degree in computer engineering from the National University of Science and Technology (NUST), Rawalpindi, Pakistan, in 2003. He is currently working towards the Ph.D. degree at the School of Computer Science and Engineering, University of New South Wales (UNSW), Sydney, Australia.

He is broadly interested in the research areas of routing in multirate multihop wireless networks, network optimization, and QoS provisioning.