Localized Minimum-Latency Broadcasting in Multi-Radio Multi-Channel Multi-Rate Wireless Mesh Networks

Junaid Qadir[1], Chun Tung Chou[1], Archan Misra[2], Joo Ghee Lim[1]
[1]School of Computer Science and Engineering, University of New South Wales, Australia
[2]IBM T J Watson Research Center, Hawthorne, New York, USA
Email: {junaidq, ctchou, jool}@cse.unsw.edu.au
archan@us.ibm.com

Technical Report UNSW-CSE-TR-0716

July 6, 2007

THE UNIVERSITY OF
NEW SOUTH WALES

SYDNEY • AUSTRALIA

Abstract

We address the problem of minimizing the worst-case broadcast delay in "multi-radio multi-channel multi-rate wireless mesh networks" *(MR^2-MC WMN)* in a distributed and localized fashion. Efficient broadcasting in such networks is especially challenging due to the desirability of exploiting the "wireless broadcast advantage" *(WBA)*, the interface-diversity, the channel-diversity and the rate-diversity offered by these networks. We propose a framework that calculates a set of forwarding nodes and transmission rate at these forwarding nodes irrespective of the broadcast source. Thereafter, a forwarding tree is constructed taking into consideration the source of broadcast. Our broadcasting algorithms are distributed and utilize locally available information. To the best of our knowledge, this works constitutes the first contribution in the area of distributed broadcast in multi-radio multi-rate wireless mesh networks. We present a detailed performance evaluation of our *distributed and localized algorithm* and demonstrate that our algorithm can greatly improve broadcast performance by exploiting the rate, interface and channel diversity of MR^2-MC WMNs and match the performance of centralized algorithms proposed in literature while utilizing only limited two-hop neighborhood information.

1 Introduction

Wireless mesh networks *(WMN)* [2] offer an exciting low-cost, decentralized, extended-area wireless networking paradigm, whereby a relatively static set of mesh nodes provide a multi-hop access infrastructure for urban and community environments. Increasing the capacity (or admissible traffic load without violating the requisite delay and loss constraints) of WMNs, however, remains a pressing research priority. While the use of a single common radio channel suffers from poor spatial reuse and has scalability problems [8], recent work has shown that the usage of multiple radios on each mesh node can significantly improve the WMN capacity [6] [10] [17]. Another feature that mesh nodes can employ to improve performance is the ability to transmit at multiple transmission rates according to the available channel and network conditions.

These features of channel and rate diversity have, however, been rarely exploited for the dissemination of data for *network-wide broadcast* or multicast applications. As many of our targeted broadcast-based applications, e.g. IP-TV, audio conferencing and multi-player games, are interactive or have strict latency requirements, we focus on how link-layer rate diversity, interface and channel diversity can be harnessed in multi-radio WMNs to improve the metric of *"broadcast latency"*, defined as the maximum delay between the transmission of a packet by a source node and its eventual reception by all WMN receiver nodes. Choosing latency as a performance measure implicitly rewards approaches that use the "wireless broadcast advantage" (*WBA* [23]) to reduce the number of distinct transmissions (since this reduction directly translates into lower contention induced delay) and also increases the throughput achieved over the WMN. We refer to the problem of constructing wireless broadcasting trees that minimize the broadcast latency as the "minimum latency broadcast" *(MLB)* problem, which is known to be NP-hard in general [7]. Prior heuristics for approximately solving the MLB problem (e.g., [5] [13]) are, however, centralized in nature–while they do significantly lower the broadcast latency, they incur large global communication overhead and require the modification of the entire tree even when the WMN topology changes very slightly.

In this paper, we focus on the design and performance evaluation of *localized* and *distributed* rate-diversity aware tree construction techniques that require only 2-hop topology information, and assume no knowledge of the global WMN topology. Our objective is to computing the broadcast forwarding trees in a purely distributed fashion, with low message overhead, and yet try to achieve the low broadcast latencies demonstrated by previously proposed centralized heuristics.

1.1 Contributions of This Paper

This paper makes the following contributions in the area of "decentralized multi-radio rate-diversity aware" WMN broadcasting:

- It presents a 4-staged heuristic framework called **MRDT**, which represents the first distributed solution to the MLB problem for "multi-radio multi-rate multi-channel" *(MR^2-MC)* WMNs.
- It analytically determines the complexity of the presented heuristics and demonstrates that our heuristics can scale to large networks.
- It demonstrates through detailed simulations (using an underlying 802.11 MAC) that MRDT performs comparably to the best-performing heuristics proposed previously, especially for realistic WMN settings.

The rest of the paper is organized as follows: The related work is presented in Section 2. This is followed by the introduction of our network model in Section 3. We describe MRDT, our distributed broadcasting framework comprising four distinct stages, in Section 4. We then present performance results of MRDT in Section 5. We finally conclude our work in Section 6.

2 Related Work

A variety of distributed data-broadcasting algorithms ([25] [16] [19] [22]) compute a set of 'backbone' nodes that are responsible for forwarding broadcast data packets. These algorithms essentially try to compute a reduced set of nodes constituting a "connected dominating set" (*CDS*). The CDS of a topology, represented by a graph $G = (V, E)$, is a connected subgraph of G spanned by the nodes of $V' \subseteq V$ such that every node in the network is at most one hop distant from a node in V'. All of these 'backbone-based routing algorithms' are directly applicable only to the more primitive case of "single-rate, single-channel" *(SR-SC)* WMNs. They fail to consider that data forwarding in MR^2-MC WMNs should utilize other available opportunities, such as the maximization of transmission rates at a forwarding node (to reduce the transmission delay) and the effective use of interface diversity (to reduce the contention delay).

The MRDT heuristic presented in this paper is based on significant modifications to two underlying (rate-diversity unaware) techniques that both calculate a 'small' CDS by first computing a large CDS and then pruning away redundant transmissions. ***Firstly,*** the ***Wu-Li*** algorithm [25] is a simple localized technique that uses only 2-hop neighborhood information to compute a CDS as follows. Initially, all vertices (nodes) are "unmarked". The marking process uses the following simple rule: any vertex having two unconnected neighbors (not connected directly) is "marked" as a dominator. The set of marked vertices form a rather large CDS V'. Two pruning techniques are then used to reduce the CDS size. A node u can be removed from V' if there exists a node v with higher ID such that the closed neighbor set[1] of u is a subset of the closed neighbor set of v. For the same reason, a node u will be deleted from V' when two of its connected neighbors in V' with higher IDs can cover all of u's neighbors. This pruning idea is generalized to the following rule [25]: a node u can be removed from S if there exist k connected neighbors with higher IDs in S that can cover all u's neighbors. ***Secondly,*** the "multi-point relaying" (***MPR***) technique [1] can be used to locally compute a CDS. The MPR technique requires each node u to first elect a 'multi-point relay set" *MRS*) [16] [4] from its one-hop neighbors that cover all its two-hop neighbors. Finding a MRS with minimum size is NP-Complete [16]. The CDS is calculated as follows [1]: each node first compute a MRS, a subset of one-hop neighbors that can cover all its two-hop neighbors. After each node has determined its MRS, a node decides that it is in the connected dominating set by matching either *Rule 1:* the node is smaller than all its neighbors *or Rule 2:* it is multipoint relay of its smallest neighbor. Although neither of these two relatively simple algorithms necessarily form the smallest CDS, we shall see that for MR^2-MC networks, the *subsequent steps of rate and channel diversity maximization (which are new facets of the MRDT algorithm that we present) turn out to be much more important than the precise computation of the initial CDS.*

Our distributed algorithms shall utilize the concept of link-rate diversity, whereby different forwarding nodes broadcast at different transmission rates (a concept first demonstrated in [5] for a *SR-SC* WMN). In general, a node attempting to reach multiple downstream neighbors through a single broadcast transmission is constrained to use the lowest of the individual link rates. For example, if a node n is to multicast to two neighboring nodes m_1 and m_2, and if the maximum unicast rates from n to m_1 and m_2 are, respectively, r_1 and r_2, then the maximum rate n can use is $\min\{r_1, r_2\}$. Based on this insight, the resulting conflict between the goals of faster transmission rates and greater node coverage per transmission is reconciled through the use of a "rate-area-product" *(RAP)* metric [5]. The RAP of a transmission rate is defined as the product of the transmission rate and the transmission area that it covers. Our previous work [5] showed that a transmission rate whose RAP is higher is more efficient in reducing broadcast latency; we also proposed a centralized heuristic "weighted CDS" *(WCDS)* for SR-SC WMNs that utilizes the RAP metric and attempts to balance the conflicting goals of achieving a small forwarding-set with lower latency weight (the conflict arising since lower latency implies the usage of links that have higher rates *but* smaller ranges) [5].

[1] closed neighbor set is the union of the node itself and its neighbors

Subsequently, centralized heuristic algorithms have been proposed for MR^2-MC WMNs to exploit its rate, interface and channel diversity [13]. The *MWT*, "multi-radio WCDS tree", is a direct extension of the rate-diversity aware WCDS algorithm, but does not exploit the additional transmission concurrency offered by the availability of multiple radio interfaces on each WMN node. The *LMT* ("locally parallelized, multi-radio WCDS tree") algorithm parallelizes transmissions *locally* on a node's different interfaces (tuned to orthogonal channels) to have multiple overlapped transmissions [13]. The ***PAMT*** ("parallelized approximately-shortest multi-radio WCDS tree) algorithm improves on LMT's performance by extending the parallelization scope to also include interfaces on *other* nodes. *In general, PAMT provides the best-known low-latency broadcast routing heuristic in a MR-MC WMN*; it performs well due to its adaptive nature by resembling a WCDS tree for a SR-SC WMN, and a"shortest path tree" (*SPT*) for an WMN with large number of orthogonal interfaces per node.

To our knowledge, there have been three distributed protocols previously proposed for multi-rate broadcasting in wireless networks (*all for the SR-SC environment*). The "rate-adaptive multicast" (*RAM*) algorithm was proposed for SR-SC multi-rate MANETs in [12]. This protocol, which is based on the "on-demand multicast routing protocol" (***ODMRP***) [11] is designed for highly mobile networks and does not explicitly exploit the WBA or the interface and channel diversity available in a WMN. An alternative algorithm called 'Multi-radio delayed-pruning Wu-Li' (***MDW***) has been proposed in [14] for SR-SC multi-rate WMNs. Lastly, a 'Distributed Rate-First' algorithm has recently been proposed for use in SR-SC multi-rate WMNs [20]. *All these three algorithms do not utilize the interface or the channel-diversity available in a multi-radio WMN.*

3 Network Model

We assume that there are C orthogonal channels in the system with each node equipped with Q interfaces where $Q \leq C$. We assume that two radio interfaces at the same node are not tuned to the same channel. We represent the nodes in the network by V. The total number of nodes ($|V|$) in the network is represented by N. A channel assignment $\mathcal{C}$ assigns each vertex v in V, Q different channels denoted by the set: $\mathcal{C}(v) = \{\mathcal{C}(v_1), \mathcal{C}(v_2), \ldots, \mathcal{C}(v_Q) : \mathcal{C}(v_i) \neq \mathcal{C}(v_j), \forall i \neq j\}$ where $\mathcal{C}(v_i)$ denotes the channel assigned to interface i of v. The topology defined by $\mathcal{C}$ is represented by $G = (V, E, \Pi, \Lambda)$ where V, E, Π, Λ are the set of nodes, links, rates of links and channel of links, respectively. The quickest-rate transmission supported between u and v is denoted by $\pi(u, v)$. The channels used for communication between two nodes u and v is denoted by $\lambda(u, v)$ where $\lambda(u, v) \in \mathcal{C}(u) \bigcap \mathcal{C}(v)$. The network topology is represented by G in the following natural way: an edge

V	Set of vertices (or, nodes)
E	Set of edges (or, links)
Π	Set of transmission-rate of links in E
Λ	Set of channel of links in E
Q	Number of radio interfaces at each node
C	Total number of orthogonal radio channels
$\mathcal{C}(u)$	Set of channels node u's Q interfaces use
N	Total number of nodes in network (=$\|V\|$)
u_i	i^{th} interface of node u
ρ_i	i^{th} highest transmission rate supported by MAC
$\rho(u_i)$	Current transmission-rate of u_i
ρ_0	Rate of a non-transmitting interface
L	Number of distinct rates supported by MAC
$\mathcal{C}(u_i)$	Channel interface u_i is tuned to
$N(u_i)$	1-hop neighbors that u_i is currently covering
$N_{\rho_k}(u_i)$	1-hop neighbors of u_i (on rate ρ_k)
$r(u_i)$	Set of rates u_i having a "rate-limiting-node"
$r_0(u_i)$	Rate of u_i when it is not transmitting
bin_i	Set of neighbors of a node's i^{th} interface
$\rho(bin_i)$	transmission-rate of a node's i^{th} interface
$\pi(u,v)$	Highest transmission-rate link (u,v) can use
$\lambda(u,v)$	Channels link (u,v) can use
m	Number of marked-nodes
d	maximum number of neighbors of a marked-node

Table 1: Index of mathematical symbols used

$e = (u, v)$ between two nodes u and v exists on channel k ($\lambda(e) = k$) is in G, *if and only if*, $d(u, v) \leq r$ and $\lambda(e) \in \mathcal{C}(u) \bigcap \mathcal{C}(v)$. The rate of the edge e is the fastest transmission rate supported on e. The set Π contains the rate of each edge in E; similarly the set Λ contains the channel used on each edge in E. Note that G may be a *multi-graph*, with multiple edges between the same pair of nodes, when the node pair shares two or more channels. We assume that the MAC layer supports L different transmission rates, represented by $\rho_1, \ldots, \rho_L$ where $\rho_1 > \rho_2 > \ldots > \rho_L$. For mathematical compactness, we denote the transmission rate of a non-transmitting interface as ρ_0. The interface i at node u is represented by u_i, and its transmitting rate by $\rho(u_i)$. $N_{\rho_k}(u_i)$ refers to the neighbors of u_i that share a channel with u_i and can use a maximum rate of ρ_k to connect to u_i.

We assume that channel assignment is performed independently of our broadcasting framework. Further, we assume that each node knows its neighbors as well as the interfaces and rates it can use to reach them. The rate-adaptation, for example, can be performed using any of the frame-error based adaptation [24], throughput-based adaptation [3], or the SNR-based adaptation [9] techniques. In order to maintain bidirectional connectivity, the rate that nodes u and v can use to reach each other is the minimum rate that can be used in the two directions. We assume that both broadcast and unicast traffic will coexist on the network; accordingly, the 'current rate' of a particular link between any two nodes can actually be inferred from the rates to which the unicast flows converge. The two-hop topology information can be built by having each node broadcast a packet containing information about sending node's node-ID, neighbors on different rates and neighbors on different channels.

4 The MRDT Framework

We will now propose a distributed and localized framework, called "multi-radio distributed tree" (***MRDT***), that calculates low latency trees for broadcast in MR^2-MC WMNs in four logically independent stages. ***Firstly***, the "initial marking" stage, which is unaware of rate, interface and channel diversity in the network, initially approximates the forwarding-node set and forms a CDS. ***Secondly***, the "neighbor grouping" *(NG)* stage, which is also rate, interface and channel diversity unaware, decides the neighboring nodes a marked node has to cover. ***Thirdly***, the "rate maximization" *(RM)* stage, comprising of two sub-stages, maximizes the transmission rates at all the marked nodes; the first sub-stage, called "local rate maximization" (*LRM*), attempts rate maximization locally at a marked node by parallelizing its transmissions over its interfaces, while the second sub-stage called "external rate maximization" (*ERM*) attempts rate maximization at a node by 'exporting' its neighbors, that are limiting its rate, to other marked nodes. ***Lastly***, the "tree construction" stage constructs a source-specific broadcast tree that takes into account WBA and prunes redundant transmissions retained in earlier stages. The stages of our broadcasting framework, and specific algorithms used during them, are covered in more detail in the following subsections.

4.1 Stage 0—Initial Marking:

This stage can initially approximate the forwarding set (also called the CDS) by using three methods. ***Firstly***, Wu-Li marking process (see Section 2) can be used in which a node is marked if it has two neighbors (at the lowest rate) that are not directly connected. ***Secondly***, MPR based marking process (see Section 2) can be used in which a rate-diversity aware algorithm like WCDS [5] is used to generate the MRS of each node i.e., each node executes WCDS algorithm with itself as the source on its 2-hop neighborhood subgraph to determine the set of its one-hop neighbors to cover its entire 2-hop neighborhood. ***Lastly***, all the nodes can be marked as eligible forwarders;

whereas such an approach results in a large CDS, we shall see in Section 5 that this approach returns good results as the rate-maximization steps (in following stages) can exploit the larger CDS.

4.2 Stage 1—Neighbor Grouping (NG)

We decide the neighboring nodes a marked node has to cover in the NG stage. The intuition is straight-forward, a marked node's transmission rate should not be constrained to a lower rate to cover a node that can be, alternatively, be 'better' covered by another marked node. This stage exploits the redundancy available in wireless networks where a node is potentially covered by many transmitters.

Algorithm 1 Neighborhood Grouping at a marked node u

$N_{\rho_L}(u) = \cup_{i=1,\dots,Q}\ N_{\rho_L}(u_i)$
for each one-hop neighbors $v \in N_{\rho_L}(u)$ **do**
  for each *marked* node $w \in N_{\rho_L}(u)\backslash\{v\}$ **do**
    if $1/\pi(u,v) > 1/\pi(u,w) + 1/\pi(w,v)$ **then**
      remove v from neighbor-list of u at rate $\pi(u,v)$
    end if
  end for
end for

In the NG stage, explained in Algorithm 1, at each node u, it is searched if there exists a 1-hop neighboring node v that can be 'better' covered by w (a 1-hop *marked* neighbor of u). The node v is said to be better covered by w if the aggregate throughput/rate of the path $u \to w \to v$ is better than the throughput of the path $u \to v$. The 1-hop neighborhood of each marked node is decided at the completion of this stage; each marked node is now responsible for ensuring coverage, by itself or through another *connected marked node*, of its 1-hop neighborhood. We illustrate the NG stage using a simple example. Let us assume node u can reach v and w using 1 Mbps and 11 Mbps link, respectively. Let us also assume that node v can be covered by w using a 11 Mbps link. Since the condition $1/\pi(u,v) > 1/\pi(u,w) + 1/\pi(w,v)$ of Algorithm 1 is satisfied (as $\frac{1}{1} > \frac{1}{11} + \frac{1}{11}$), v is removed from the neighbor-list of u at the rate of 1 Mbps.

Theorem 4.1 *The computational complexity of the NG stage at a marked node u is $O(d^2)$ where d represents the maximum degree of a marked node.*

Message Complexity: Assuming that 2-hop neighborhood information has been established prior to the NG stage, no message needs to be exchanged during the NG stage. After the NG stage finishes, each marked node will broadcast a *NEIGHBOR* packet for a total of m *NEIGHBOR* packets—the maximum size of which is $(1 + (Q + L)d)$ times the bytes required to represent a node-id—to convey the sending node's node-ID, neighbors on different channels and rates, to all its neighbors. We note that Q and L are small (constant) values since typically limited interfaces and rates are supported; the total message-complexity of the NG stage, therefore, is $O(md)$.

4.3 Stage 2—Rate Maximization (RM)

Before discussing the RM stage, we introduce the concept of "*rate-limiting-nodes*". We note that a lower-rate transmission can cover all nodes reachable at a higher-rate but not vice-versa; this implies that the maximum rate, a node u can use to reach all its 1-hop neighbors $N(u)$ collectively, is the minimum of the (maximum) rate u can use to reach each individual node in $N(u)$. To illustrate this concept, assuming a single radio interface, refer to Figure 1 for an example topology. Although, u can reach nodes w, x and y with a rate of 5.5, 11 and 54 Mbps, respectively, u is constrained to transmit at a lowest-rate of 2 Mbps to reach node v. Node v, for this topology, is referred to as a rate-limiting-node since its presence limits the rate of u to 2 Mbps and its absence can increase the rate of u to 5.5 Mbps.

The two sub-stages of Stage 2, i.e. LRM and ERM, differ in how they deal with rate-limiting-nodes. LRM exploits interface-diversity at each marked node to 'export' the rate-limiting-node to a different interface *on the same node*, thereby increasing the original interface's rate; with this export, LRM also exploits the channel-diversity since different interfaces on a node are tuned to orthogonal channels. As the 'export' scope of rate-limiting-nodes in LRM is confined to local interfaces on the same node, LRM is similar to the LMT algorithm [13] which parallelizes transmissions locally over the node's interfaces. ERM, on the other hand, extends the scope of exporting rate-limiting-nodes to include interfaces at other nodes (subject to a few conditions) like the PAMT algorithm [13] in which an attempt is made to parallelize a node's transmissions in its neighborhood. The details of LRM and ERM follow next.

4.3.1 Local Rate Maximization (LRM)

The LRM stage distributes the neighbors of a node over its interfaces in a manner such that the rate, interface and channel diversity of the WMN is exploited. We define $\rho(u_i)$ as the current transmission rate of an interface u_i and $N(u_i)$ as the neighbors of u_i (that share a channel with $\mathcal{C}(u_i)$) that u_i is currently covering at its transmission rate $\rho(u_i)$. During the stage of LRM at any node u, we assign u's neighbors to u's interfaces such that u maximizes the sum, taken over all of u's interfaces, of the product of its interface's rate and neighbors on that interface (i.e., the metric $\sum_{i=1}^{Q} \rho(u_i) \times |N(u_i)|$ denoted

Algorithm 2 Local Rate Maximization at node u

$\rho(u_q) = 0$ and $bin_q = \emptyset$, $\forall q = 1$ to Q
for $q = 1$ to Q **do**
while $U \neq \emptyset$ **do**
$f(j,q)$ and $rate(j,q)$=$-\infty$, $\forall j = 1$ to $|U|$, $\forall q = 1$ to Q
——————————————————————————- {Finds the assignment of a node in U to an interface on u that maximizes $\delta(\texttt{sum})$}
for $j = 1$ to $|U|$ **do**
$\tilde{C}$= $\mathcal{C}(U(j)) \cap \mathcal{C}(u)$
for $k = 1$ to $|\tilde{C}|$ **do**
i = interface of u tuned to $\tilde{C}(k)$;
$\widehat{\mathcal{C}} = \mathcal{C} \backslash \tilde{C}(k)$
$I = \bigcup$ u's interfaces tuned to a channel in $\widehat{\mathcal{C}}$
if $\exists\ \rho_{I(x)}(u) = \pi(U(j), i)$, $\forall\ x = 1$ to $|I|$ **then**
$f(j,i) = -\infty$
$bin_{I(x)} = bin_{I(x)} \cup U(j)$
else
if $\rho(u_i) \neq 0$ **then**
$rate(j,i) = \min(\pi(U(j),u), \rho(u_i))$
else
$rate(j,i) = \pi(U(j),u)$
end if
$O = \bigcup_{\forall u_o \neq u_i} u_o$ **s.t.** $\rho(u_i) > \rho(u_o) > rate(j,i)$
$E = \cup n \in bin_i$ **s.t.** for any $u_o \in O$, $\pi(n,u_o) > rate(j,i)$
$EG = \sum_{\forall n \in E} \max_{\forall u_o \in O} \pi(n,u_o)$
$f(j,i) = (1 + |bin_i(u)| - |E|) \times rate(j,i) + EG$ - $(|bin_i(u)| \times \rho(u_i))$;
$O(j,i) = O$
end if
end for
end for
$[\widehat{j}, \widehat{i}] = \arg\max(f)$
if $\max(f) \neq -\infty$ **then**
{if assigned nodes should change bin to increase $\delta(\texttt{sum})$}
if $rate(\widehat{j}, \widehat{i}) < \rho(u_{\widehat{i}})$ **then**
$O = O(\widehat{j}, \widehat{i})$
while $O \neq \emptyset$ **do**
u_m= arg $\max_{\forall u_o \in O} \rho(u_o)$
$N = \bigcup n \in bin_{\widehat{i}}$ **s.t.** $C(n) \cap C(u_m) \neq \emptyset$
$bin_o = bin_o \cup N$;
$bin_{\widehat{i}} = bin_{\widehat{i}} \setminus N$
$O = O \backslash \{u_m\}$
end while
end if
$\rho(u_{\widehat{i}}) = rate(\widehat{j}, \widehat{i})$;
$bin_{\widehat{i}} = bin_{\widehat{i}} \cup U(\widehat{j})$;
end if
$U = U \backslash U(\widehat{j})$
end while
end for

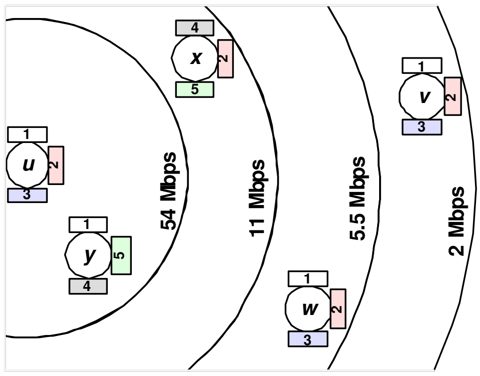

Figure 1: Before Local-Rate-Maximization at u

Step	bin$_1$	bin$_2$	bin$_3$	$\rho(u_1)$	$\rho(u_2)$	$\rho(u_3)$	$\delta(\texttt{sum})$
0	$\emptyset$	$\emptyset$	$\emptyset$	0	0	0	0
1	$\{y\}$	$\emptyset$	$\emptyset$	54	0	0	54
2	$\{y\}$	$\{x\}$	$\emptyset$	54	11	0	11
3	$\{y\}$	$\{x\}$	{w}	54	11	5.5	5.5
4	$\{y\}$	$\{x\}$	{w,v}	54	11	2	-1.5

Table 2: The steps of LRM at u of Figure 1

by "sum" is maximized); if u's j^{th} interface is unused, its transmission rate is zero (i.e., if there is no transmission on u_j, $\rho(u_j) = 0$). By maximizing sum, we ensure that nodes are covered in maximum rate transmissions that are parallelized. The LRM sub-stage, analogously to the function performed by the centralized LMT algorithm [5], performs *local* parallelization to utilize the interface and channel diversity available at the node.

We illustrate the working of LRM for a simple topology shown in Figure 1 in which each node has 3 interfaces. The channel to which u_i (the i^{th} interface of u) is tuned to is represented by $\mathcal{C}(u_i)$; also, $\mathcal{C}(u)$ represents the

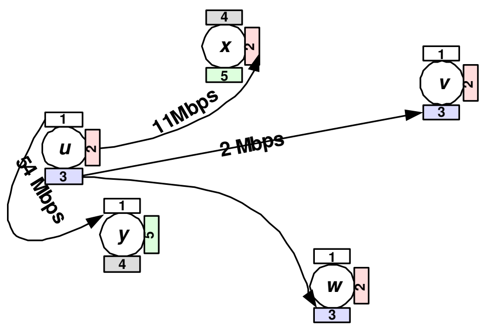

Figure 2: After Local-Rate-Maximization at u

channel set of u, i.e., $\mathcal{C}(u) = \cup_{i=1}^{Q}\mathcal{C}(u_i)$. In Figure 1, $\mathcal{C}(u) = \{1,2,3\}$, $\mathcal{C}(v) = \{1,2,3\}$, $\mathcal{C}(w) = \{1,2,3\}$, $\mathcal{C}(x) = \{2,4,5\}$ and $\mathcal{C}(y) = \{1,4,5\}$. Initially, the value of `sum` is equal to 0 as depicted in Table 2 since none of the interface is transmitting (i.e., $\forall i,\ \rho(u_i) = 0$). Node u connects to 4 neighbors; the interface neighbor set for first interface $N(u_1) = \{v,w,y\}$, similarly, $N(u_2) = \{v,w,x\}$ and $N(u_3) = \{v,w\}$. Furthermore, u can connect to $\{y\}, \{x\}, \{w\}$ and $\{v\}$ on 54 Mbps, 11 Mbps, 5.5 Mbps and 2 Mbps, respectively. The current rate $\rho(u_i)$ at bin_i is then determined as the minimum of the rates to individual nodes in bin_i where bin_i denotes the set of nodes assigned to the i^{th} interface of u. During each step, the node and *bin* combination that maximizes $\delta(\texttt{sum})$ is chosen. During step 1, y is added to bin_1 which gives $\texttt{sum} = 54$ $(54 \times 1 + 0 + 0)$ at the end of step 1 for maximum $\delta(\texttt{sum})$ of 54. Similarly, for step 2, x is added to bin_2 with $\texttt{sum} = 65$ $(54 \times 1 + 11 \times 1 + 0)$ for maximum $\delta(\texttt{sum})$ of 11. During the following step (step 3 in Table 2), w can be added to bin_1, bin_2 or bin_3 since w belongs to each of $N(u_1)$, $N(u_2)$ and $N(u_3)$. However, adding w to bin_1 and bin_2 would cause $\delta(\texttt{sum})$ to be -43 and 0, whereas adding w to $\delta(\texttt{sum})$ would make $\delta(\texttt{sum})$ equal to 5.5. To maximize the incremental $\delta(\texttt{sum})$, we assign w to bin_3. Similarly, in step 4, v can be assigned to bin_1, bin_2 or bin_3 with $\delta(\texttt{sum})$ for each assignment being -50, -7 and -1.5 respectively. The assignment of v to bin_3 is chosen since it maximizes the step $\delta(\texttt{sum})$. The final value of bin_i and $\rho(u_i)$, for all values of i from 1 to Q, is shown in Table 2. As seen in the last step, $\delta(\texttt{sum})$ can be negative; since when a *bin* is non-empty, addition of new nodes can never increase the *bin* rate as increase of rate will imply disconnection of the nodes in the *bin*. On the other hand, when a node is added to an empty *bin* (i.e., when the interface's rate is zero), the $\delta(\texttt{sum})$ will always be positive.

The LRM algorithm, for any node u, is mathematically described in Algorithm 2. Initially, the transmitting rate at all the interfaces of all nodes is 0; for this reason, all the interface *bin*s at all nodes are also empty. U, denoting a set of unassigned (to any *bin*) neighbors, is initialized with all the neighbors of u. Our algorithm then, in each round, determines an unassigned node ($U(\widehat{j})$ in Algorithm 2) that can be placed in an interface *bin* to cause maximum increase in `sum`. We define as $f(j,i)$ (the $\delta(\texttt{sum})$ in Table 2) as the change in `sum` after the j^{th} element of U has been added to the i^{th} interface *bin*. Since in the stage of LRM, our aim is to exploit interface-diversity and WBA simultaneously; a node should suppress a transmission on rate ρ_i on a non-transmitting interface to cover a node that is in range of an already transmitting interface whose rate is ρ_i. Accordingly, the metric $f(j,i)$ is set to -∞, if bin_i is currently empty and $U(j)$ can be placed in a non-empty interface *bin* of u whose current rate is exactly $\pi(U(j),u)$, to eliminate a redundant transmission. After placing a node $U(j)$ in bin_i, the rate for that *bin* (represented by $rate(j,i)$) is calculated as the minimum of the (max) rate to an individual node in the interface *bin*. In case of an empty *bin*, this is equal to maximum rate $U(j)$ can support, i.e. $\pi(U(j),u)$, while in the case of a non-empty bin, the rate is $\min(\pi(U(j), u), \rho(u_i))$. Since $rate(j, i)$ can be less than $\rho(u_i)$, nodes already assigned to bin_i can transfer to alternate $bins$ (represented by set O in Algorithm 2) whose current rate is more than $rate(j, i)$ in case of a shared common channel to increase rate. The set of these nodes is called E and the sum of the new rates for each node in E is represented by EG. The $f(j, i)$ metric, defining the desirability of assigning U_j to bin_i, is now calculated as $((1 + |bin_i(u)| - |E|) \times rate(j, i) + EG$ - $(|bin_i(u)| \times \rho(u_i))$. An unassigned node $U(\widehat{j})$ and the interface $bin_{\widehat{i}}$ can be decided according to highest f value. If $rate(\widehat{j}, \widehat{i})$ is lower than $\rho(u_{\widehat{i}})$, all nodes assigned to $bin_{\widehat{i}}$ check if they can migrate to an interface in set $O(\widehat{j}, \widehat{i})$ which comprises of interfaces whose rate range from $rate(\widehat{j}, \widehat{i})$ to $\rho(u_{\widehat{i}})$. N, denoting nodes that can migrate to an interface u_o in O, is then added to bin_o and subtracted from $bin_{\widehat{i}}$. $U(\widehat{j})$ can now be added to $bin_{\widehat{i}}$ and taken out from the set of unassigned nodes U. The algorithm completes when U becomes an $\emptyset$ when all neighbors of u have been assigned to an interface bin.

Theorem 4.2 *The LRM algorithm, at any marked node, is a polynomial-time algorithm of* $O(d^2Q)$.

Message Complexity: No message needs to be exchanged during the LRM sub-stage. However, after the LRM sub-stage completes, each marked node broadcasts a *NEIGHBOR* packet—the maximum size of which is $(1 + dQ)$ times the bytes required to represent a node-ID—to convey the sender's node-ID and the neighbors it can cover on each of its interface. Each marked node also broadcasts a *RATE* packet which advertises the transmission rate of each of its Q interfaces; the size of *RATE* packet is sum of the bytes required to represent a node-ID and a rate-ID. The total number of message exchanged by LRM is $2m$ and thus the message-complexity of LRM is $O(m)$.

Theorem 4.3 *The LRM algorithm returns an optimal solution to the problem of allocating neighbors to node's interfaces to maximize* $\sum_{i=1}^{Q} \rho(u_i) \times |N(u_i)|$.

4.3.2 External Rate Maximization (ERM)

The objective of the ERM sub-stage is to find, for an interface u_i, neighboring forwarders to whom u_i's rate-limiting-nodes can be 'exported'. The utility of an export can be determined using, in particular, the "rate-area-product" (RAP) maximization principle described in [5]. The export of rate-limiting-nodes, in general, will increase an interface's transmitting rate, with a node unmarking itself if all its neighbors have been exported and the rate of *all* its interfaces has become ρ_0. Whereas, in the LRM sub-stage, rate-limiting-nodes were exported to another interface on the same node; the export, in ERM, is to an interface on a neighboring marked node. The challenge faced by ERM, due to the potential danger of rate-diversity making

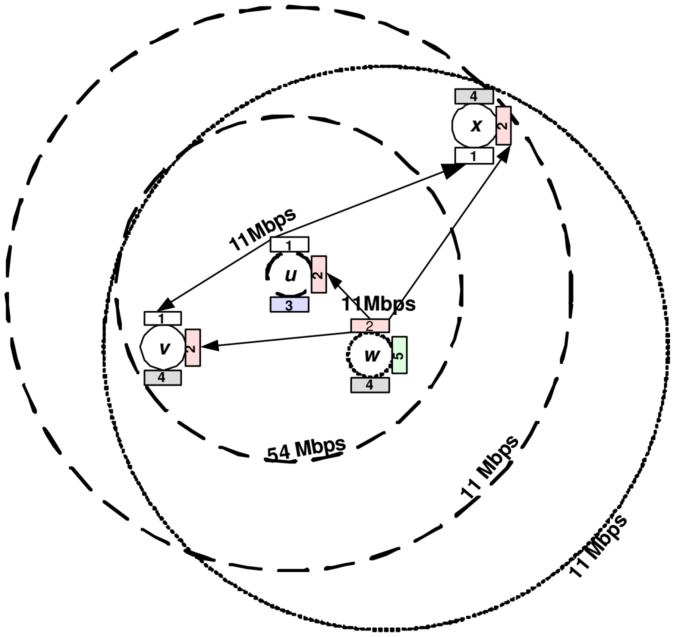

Figure 3: Before External-Rate-Maximization at u

links asymmetric[2], is maximizing rates at node's interfaces while preserving *strong* connectivity of the resulting dominating set. Since our framework determines forwarders and rates irrespective of the broadcast source (i.e., until Stage 4), it is important to ensure strong connectivity irrespective of the broadcast source.

To illustrate the concepts employed by ERM, we refer to Figure 3 for an example topology comprising of three nodes, each fitted with 3 interfaces. Node u can reach nodes $\{v, w\}$ and $\{x\}$ in a 54 Mbps and 11 Mbps transmission, respectively. However, since u shares channel 1 only with node v and x, only nodes $\{v, x\}$ are reached with the transmission of u_1 (u's interface 1 which is tuned to channel 1); for u to reach w, it must transmit on its interface 2 (tuned to the channel shared with node w). Node w, however, can reach nodes $\{v,u\}$ and $\{x\}$ on a single channel (channel 2) in a 54 and 11 Mbps transmission. We will study ERM sub-stage at node u. Interface u_1 is constrained to use a lower rate (of 11 Mbps) if both neighbors of u_1 (v and x) are to be covered in a single transmission. The rate-limiting-node of u_1 is x. Interface u_1 will look for an interface on a higher-id marked node[3] that can cover u_1's rate-limiting-node using its current rate and be reachable from u_1 after u_1 increases its rate; also, the sum of the uplink rates of

[2]e.g. it is possible for node u to reach v but not vice-versa (where $\rho(u) < \rho(v)$) due to different ranges for different rates

[3]the restrictive condition of only exporting to higher-ID neighbors is to avoid circular hand-offs

u_1's neighbors should improve after an export. We check now if u_1's rate-limiting-node x can be exported to w. Firstly, x is reachable through w_1's current transmission; secondly, w is reachable from u even after u_1's rate is increased to 54 Mbps through u_2 (which shares a channel with w_1); lastly, the sum of rates of u's neighbor increases with this transfer (54+11=65 instead of 11+11=22 before). Since all conditions are satisfied, the export of x can take place increasing the rate of u_1 to 54 Mbps as shown in Figure 4.

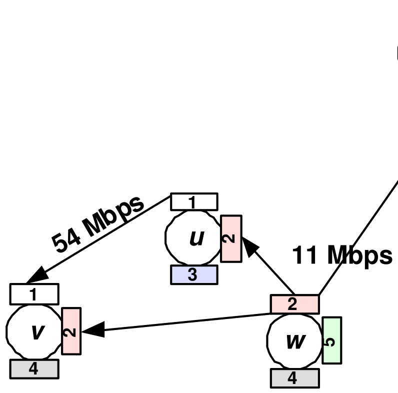

Figure 4: After External-Rate-Maximization at u

The ERM algorithm, for any node u, is mathematically described in Algorithm 3. Node u will attempt to increase the rate of its transmitting interfaces if it is currently a transmitting node (i.e. it has some rate-limiting-nodes). The token *continue* is initially equal to 1 which indicates that rate-increase can be attempted; a token *continue* valued 0, on the other hand, implies that the rate-limiting-nodes of the current rate are non-exportable and further rate-increase must not be attempted. Initially, E (denoting the rate gain for the exported nodes) is set to zero. We denote the rates on which an interface u_i has rate-limiting-nodes as $r(u_i)$. The total rates in $r(u_i)$ is not necessarily equal to the total number of rates L and is specific to each u_i. The rates in $r(u_i)$ are arranged in a descending order, i.e., $r_1(u_i) > r_2(u_i)$ and so forth. For mathematical compactness, $r_0(u_i)$ denotes the fact that u_i would not be transmitting since a non-transmitting interface has rate of zero. The index of u_i's current transmission rate, $\rho(u_i)$, in $r(u_i)$ is represented as k in Algorithm 3. The rate-limiting-nodes (RLN) is calculated as the difference between the neighbors of u_i at the current rate ($N_{r_k(u_i)}(u_i)$) and the next higher rate in $r(u_i)$ (i.e., $N_{r_{k-1}(u_i)}(u_i)$, if $r_{k-1}(u_i) \neq r_0(u_i)$). For each node rln in RLN, it is checked for every node $h \in H$ where H comprises of higher-ID marked neighbors of u excluding RLN if, *firstly*, rln is a neighbor of h (i.e., $\pi(h, rln) \geq \rho_q(h)$ for some interface h_q where h_q and rln share a common channel) and, *secondly*, if u is a neighbor of h (to ensure strong-connectivity). The maximum uplink rate rln can receive from a node $h \in H$

Algorithm 3 External-rate-maximization function at node u

for $i = 1$ to Q **do**
  $continue = 1$
  while $continue$ and $\rho(u_i) \neq 0$ **do**
    $E = 0$; $continue = 0$;
    $r(u_i)$ = rates at u_i sorted in descending order
    k = index of $\rho(u_i)$ in $r(u_i)$
    if $k = 1$ **then**
      $RLN = N_{r_k(u_i)}(u_i) \cup u$
    else
      $RLN = N_{r_k(u_i)}(u_i) \backslash N_{r_{k-1}(u_i)}(u_i)$
    end if
    H= all higher-ID *marked* neighbors of u (on any interface i= 1 to Q) $\backslash$ $\{RLN\}$
    {This part aims to find a neighbor's interface to export nodes in RLN while satisfying RAP condition}
    for $m = 1$ to $|RLN|$ **do**
      rln= $RLN(m)$; $rate\_new = -\infty$;
      for $n = 1$ to $|H|$ **do**
        $h = H(n)$
        for $q = 1$ to Q **do**
          if $rln \in N(h_q)$ **and** $u \in \cup_{i=1,\ldots,Q} N(h_i)$ **and** $\rho(h_q) > rate\_new$ **then**
            $rate\_new = \rho(h_q)$
          end if
        end for
      end for
      $rate\_diff = rate\_new - r_k(u_i)$
      $E = E + rate\_diff$
    end for
    if $E \geq 0$ **then**
      $continue = 1$; $\rho(u_i) = r_{k-1}(u_i)$
    end if
  end while
end for

fulfilling these conditions is stored in a variable called $rate\_new$ (that is initialized with -∞). The difference between the initial rate of rln and the $rate\_new$ is maintained in $rate\_diff$. The variable E contains the sum of $rate\_diff$ of all nodes in RLN. The nodes that cannot be exported have $rate\_diff$ of -∞. Thus, even for a single non exported rate-limiting-node at a particular rate, the value of E would be -∞. For each interface, if $E > 0$, its rate is increased and *continue* is set to 1; otherwise, if $E < 0$, *continue* is set to zero. The algorithm completes when increase in rate is not possible either due to export of all nodes, or due to *continue* token equal to zero.

Message Complexity: During the ERM sub-stage, each time an interface u_i of a marked node u is successful in increasing its rate, it would broadcast its new rate $\rho(u_i)$ to its neighbors in a $RATE$ message. The maximum number of $RATE$ messages exchanged is $((m-1) \times Q \times L)$ and the size of a $RATE$ packet is the sum of the bytes used to represent node-ID and rate-ID. Since Q and L are constants, total message-complexity of ERM is $O(m)$.

Theorem 4.4 *The ERM algorithm at a marked node is a polynomial-time algorithm of* $O(d^2Q)$

4.4 Stage 3—Tree construction:

The calculation of the forwarding interfaces (CDS) and their rates, till Stage 3 of our framework, is performed independent of the broadcast source. During this stage, tree relationship spanning all nodes is built and many redundant transmissions (retained during earlier stages) are eliminated. The decisions in this stage is restricted to the choice of the interfaces amongst the 'candidate' forwarding interfaces chosen earlier (along with their rates). The explicit aim of Stage 3 is to calculate a high performance tree that minimizes broadcast latency. Topology construction algorithms, for this objective, must incorporate WBA and rate-diversity, as noted in [5] [13], since the use of WBA reduces the number of transmissions to mitigate the adverse effects of interference in a wireless network, while the usage of rate-diversity can improve broadcast performance by employing higher rate links.

The tree construction (Stage 3) is mathematically described in Algorithm 4. Initially, the *label* of all nodes is equal to ∞. The source node s initially broadcasts a $RREQ$ message on each of its (transmitting) interface after setting $RREQ.source$ and $RREQ.sender$, $RREQ.interface$ and $RREQ.neighbors$ to s, the interface's id, and the neighbor set of s on that interface, respectively. The $RREQ.label$ for an interface s_i is set to its 'weighted latency', $\hat{l}(s_i)$, which is calculated as $1/\rho(s_i) \times 1/N(s_i)$. The weighted-latency metric $\hat{l}$ is based on the RAP concept presented in [5] which states that the efficiency of a rate (for broadcast latency) can be reasonably predicted from the product of the rate and its transmission range.

Algorithm 4 Distributed tree construction

```
Let label for all nodes in V = ∞ and u = id(node)
if u is the broadcast source s then
 --------------------------------
 for i = 1 to Q and ρ(s_i) ≠ 0 do
  RREQ.source = s;
  RREQ.sender = s; RREQ.interface = i
  RREQ.label = 1/ρ(s_i) × 1/N(s_i)
  RREQ.neighbors = N(s_i)
  send(RREQ) on interface i
 end for
end if
--------------------------------
if non-duplicate RREQ received on an interface î then
 Let R̃REQ = the received RREQ
 if R̃REQ.label < label(u) then
  P(u) = R̃REQ.sender
  PI(u) = R̃REQ.interface
  S(u) = R̃REQ.neighbors
  if RREPACK from RREQ.sender not received yet then
   RREP.nexthop = R̃REQ.sender
   send (RREP) on interface î to R̃REQ.sender
  end if
 end if
 --------------------------------
 if u is a marked-node that has not forwarded RREQ for R̃REQ.source before
 then
  for i = 1 to Q and ρ(u_i) ≠ 0 do
   RREQ.sender = u;
   RREQ.interface = i; RREQ.neighbors = N(u_i)
   l̂(u_i) = 1/ρ(u_i) × 1/(N(u_i)\S(u))
   RREQ.label = R̃REQ.label + l̂(u_i)
   send (RREQ) on interface i
  end for
 end if
 --------------------------------
 if received RREP and RREP.nexthop = u then
  Activate Forwarder flag for u
  RREP.nexthop = P(u)
  send (RREP) on interface PI(u)
 end if
end if
```

As a special case of the RAP principle, the efficiency of a rate is proportional to the product of that rate's transmission latency and the number of receivers (that have not received previously) in this rate's transmission area [5]. ***Firstly***, any node u that receives a non-duplicate $\tilde{RREQ}$ message on its interface $\widehat{i}$ will determine if $\tilde{RREQ}.label$ is less than $label(u)$; if so, u will choose the sender of the $\tilde{RREQ}$ as its parent ($P(u)$) and $\widehat{i}$ as the interface to connect to its parent ($PI(u)$). The neighbors of $P(u)$ (contained in $\tilde{RREQ}.neighbors$) are referred to as sibling nodes of u and are denoted by $S(u)$. Node u, if it knows that its parent $\tilde{RREQ}.sender$ is currently not a forwarder (we shall see later how), will send a $RREP$ to it on interface $\widehat{i}$ after setting $RREP.nexthop$ to $\tilde{RREQ}.sender$. ***Secondly***, if node u receives a $RREP$ message with the $RREP.nexthop$ set to u's ID, it will activate its $Forwarder$ flag and broadcast a $RREPACK$ message to announce to its neighbors that it is a forwarding node as well as send a $RREP$ back to its parent. After a node u has broadcasted a $RREPACK$, u's neighbors would not send a $RREP$ to it knowing that u is already a forwarder. ***Thirdly***, a marked-node u will forward a $RREQ$ message for a particular broadcast source $RREQ.source$ only once. Note that only marked-nodes can forward $RREQ$s. If after forwarding a $RREQ$ once, a $RREQ$ with a better $label$ (than the node's current $label$) is received, the node will modify its P, PI and S; the node will, however, not rebroadcast $RREQ$ another time after broadcasting $RREQ$ for a particular $RREQ.source$ before[4]. ***Lastly***, when marked-node u does broadcast a $RREQ$, it will generate a $RREQ$ message for each of its interface with u's ID in $RREQ.sender$ and $RREQ.label$ modified to the sum of $label(u)$ and the weighted latency $\hat{l}$ of the interface. Similarly, $RREQ.interface$ and $RREQ.neighbors$ is set to i and $N(u_i)$, respectively, for the case when interface u_i is used. The weighted-latency $\hat{l}$ for interface u_i is calculated as the product of u_i's rate and the new receivers on u_i with the new receivers on u_i approximated by subtracting $S(u)$ from $N(u_i)$. The tree construction is complete when each marked-node has relayed the RREQ once. The nodes that have the Forwarder flag activated will forward the broadcast data at pre-determined rates (decided prior to this stage).

Theorem 4.5 *The message complexity of our tree construction algorithm is O(mQ).*

5 Simulation Results:

In this section, we will present performance evaluation results for our algorithm. We have performed simulations on the Qualnet [18] simulator to see the performance of our broadcast algorithms with a decentralized MAC

[4]this enables our algorithm to minimize its message complexity

Protocol	Type	Interface and channel diversity	Rate diversity
MRDT	distributed	✓	✓
PAMT [13]	centralized	✓	✓
MDW [15]	distributed	✗	✓
ODMRP [11]	distributed	✗	✗

Table 3: Comparison of different protocols

scheduler (we have used 802.11 as our MAC scheduler). We implemented PHY 802.11a and 802.11b at the physical layer, which uses a pre-configured BER-based packet reception model. The MAC802.11 with Distributed Coordination Function (DCF) was chosen as the medium access control protocol. All default parameters are assumed unless stated otherwise. The simulations are performed using random topologies assuming a network of area 1×1 km^2.

To evaluate the relative performance of MRDT, we compare it against the performance of PAMT [13], ODMRP [11], and MDW [14]. The characteristics of these protocols are tabulated in Table 3. Both ODMRP and MDW are distributed algorithms that do not exploit interface/channel diversity of a WMN; ODMRP, also, does not exploit the rate-diversity of WMN nodes unlike MDW which is rate-diversity aware. PAMT is the centralized algorithm that exploits both rate and interface diversity, and is the benchmark for the performance of our MRDT algorithm. PAMT, with Q and C as low as 2-3, has been shown to perform with in $\sim 30 - 40\%$ of the lower bound of Dijkstra's tree (*with infinite radio interfaces and channels*) when using an idealized MAC scheduler [13].

We use three *static* channel assignment strategies in our experiment. The "common channel approach" *(CCA)* [6] assigns all nodes a common set of channels. In contrast, in the "varying channel approach" *(VCA)*, each WMN node is assigned, at random, a distinct set of channels [17]. The third channel assignment approach is based on "interference survivable topology control' *(INSTC)* [21], in which channels are assigned so as to reduce the interference in the induced network topology. Unless stated otherwise, CCA scheme with Q and $C = 3$ should be assumed for our results. The ticks in the graphs represent confidence intervals based on the 5^{th} and 95^{th} percentiles, computed over 100 uniformly distributed random topologies. We will now proceed to discuss the results in the next few subsections.

5.1 The effect of node density

The effect of network's node density on MRDT's performance can be seen in Figures 5 and 6 for 802.11a and 802.11b networks, respectively. It is observed that distributed MRDT algorithm performs as well as the centralized PAMT

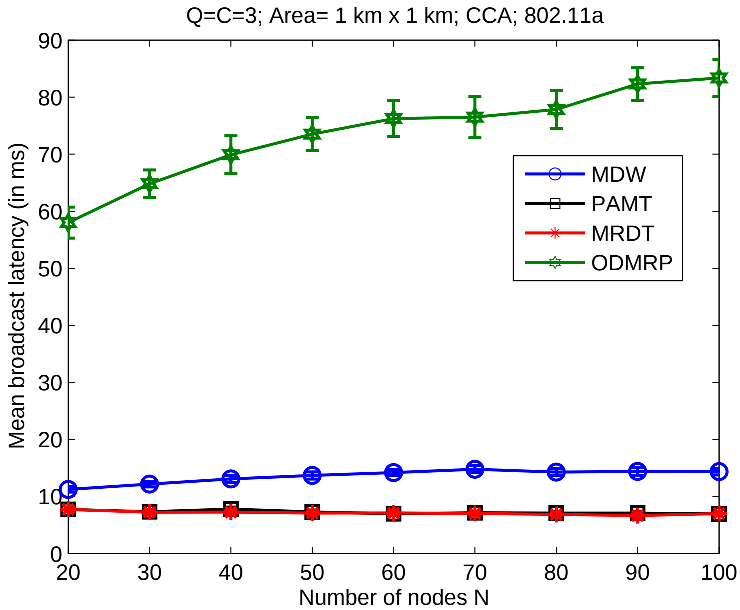

Figure 5: 802.11a with changing N

algorithm as seen in Figures 5 and 6 where the results of MRDT sit on top of PAMT's results for the range of N considered. It is also observed that MRDT improves the performance of MDW (another distributed multi-rate algorithm) by incorporating interface-diversity in its calculations. The ODMRP algorithm is included in our results as an example protocol that does not utilize rate or interface/channel diversity. The performance gain of MRDT over ODMRP and MDW algorithms is ~ 10 times and ~ 2 times, respectively (Figure 5).

5.2 The effect of number of radio interfaces

The effect of varying number of radio interfaces Q on MRDT's performance can be seen in Figures 7 and 8 for 802.11a and 802.11b networks, respectively. It is observed that, for values of Q as low as 2 or 3, MRDT matches the performance of PAMT algorithm. This improvement is possible since MRDT exploits interface-diversity by employing parallel transmissions on non-interfering interfaces. The difference in the performance of MRDT and MDW in Figure 8 demonstrates that MRDT, like PAMT and unlike MDW, is an *adaptive* algorithm that adapts to the number of radio resources available. The performance gap, between MRDT and MDW, widens as the number of radio interfaces is increased clearly demonstrating the benefit obtained by exploiting interface and channel diversity.

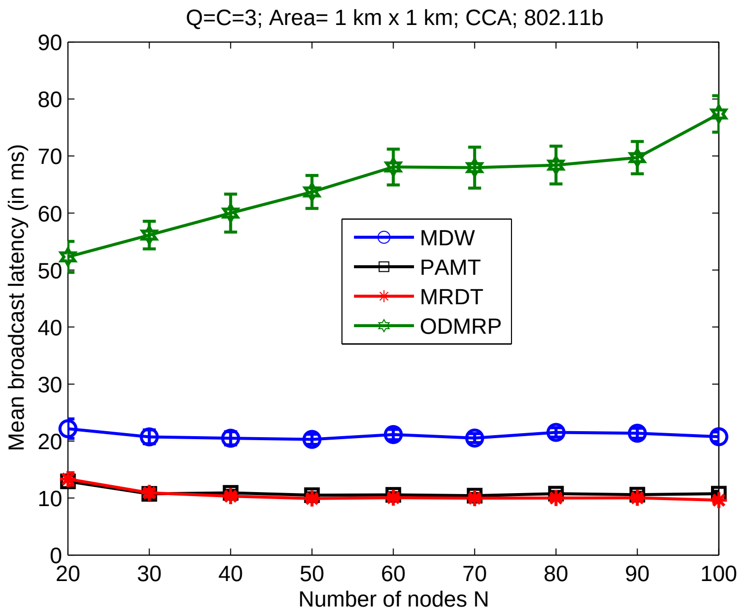

Figure 6: 802.11b with changing N

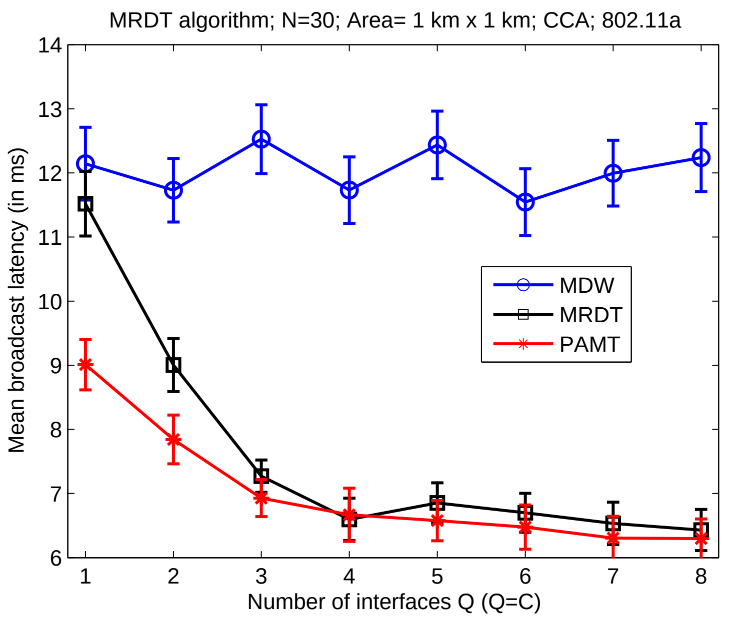

Figure 7: 802.11a with changing Q

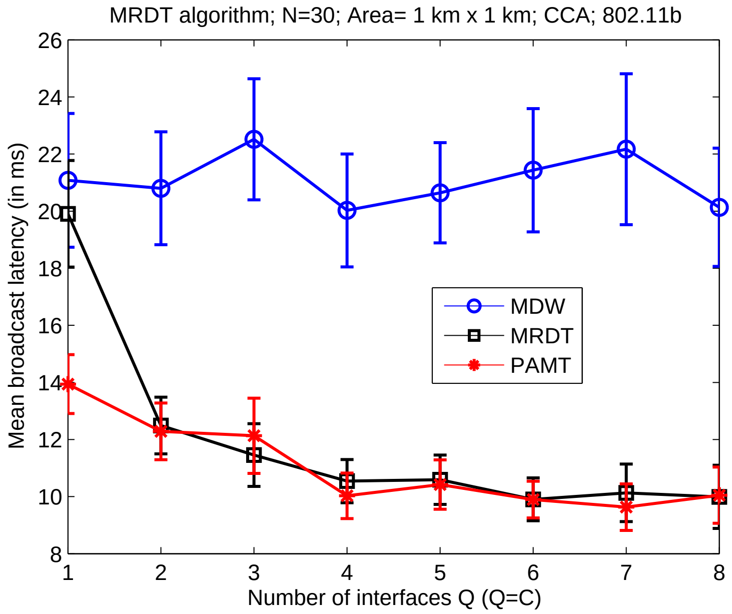

Figure 8: 802.11b with changing Q

5.3 Evaluation of different stages of MRDT

Amongst the different stages of MRDT, the LRM sub-stage offers the most improvement in broadcast-latency performance across the range of Q excepting the case when $Q = 1$ (since LRM can not perform its parallelizing function due to the lack of interface diversity then). Without the LRM sub-stage, MRDT does not remain adaptive to the available radio resources and cannot exploit the interface and channel-diversity. The improvement due to LRM is clearly evident in Figure 9. The Neighbor-Grouping (Stage 1: NG) also improves the performance markedly across the range of Q and especially for $Q = 1$, as shown in Figure 9. The use of ERM substage (of Stage 2), however, has limited benefits vis-a-vis broadcast latency reduction; its main impact, as we shall see later, is in reducing overhead during the stage of Tree-Construction (Stage 3). The different marking schemes (of Stage 0) also affect the broadcast latency performance as shown in Figure 10 which shows that marking by MPR, or marking all nodes in Stage 1 gives better results than marking by the Wu-Li method.

5.4 The overhead of Tree Construction (Stage 3)

The message overhead during Tree Construction (Stage 3) depends upon the marking scheme used in Stage 0. Recall that the message-complexity of Stage 3 is $O(m)$; therefore, the message overhead for marking schemes that return a smaller set of marked nodes m is lower. We show the total number

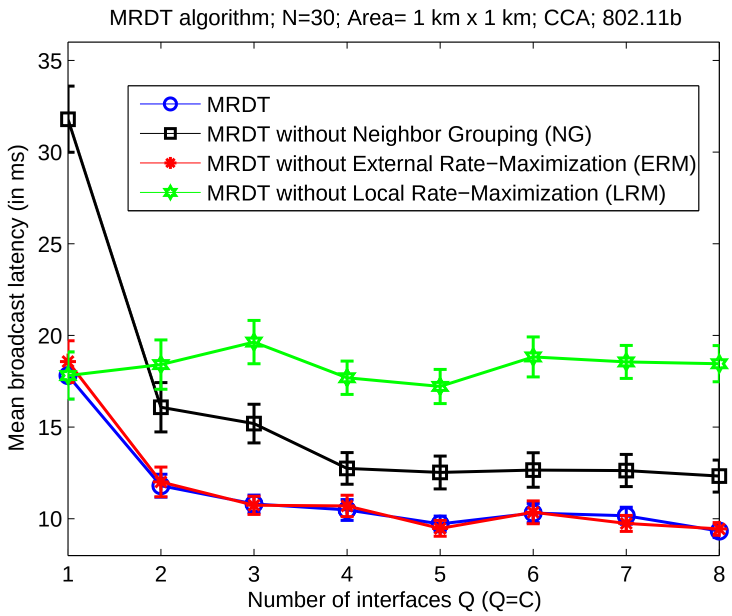

Figure 9: MRDT's different stages

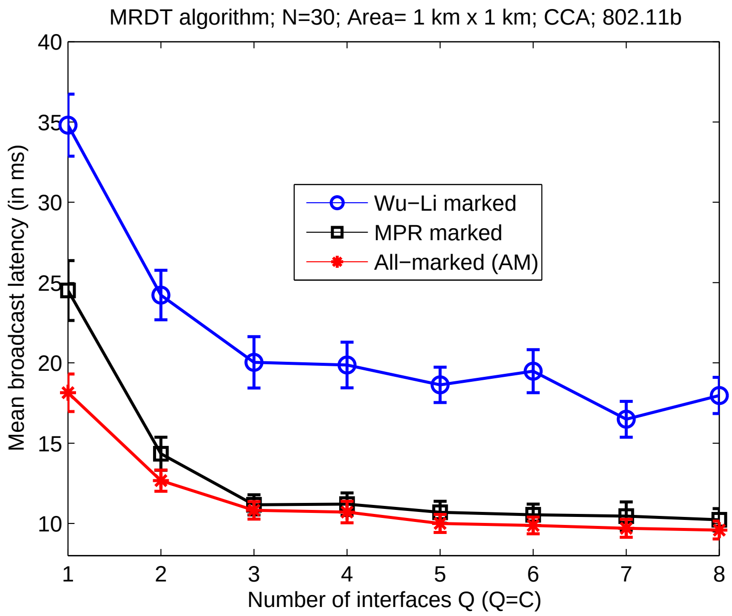

Figure 10: Stage 0 marking techniques

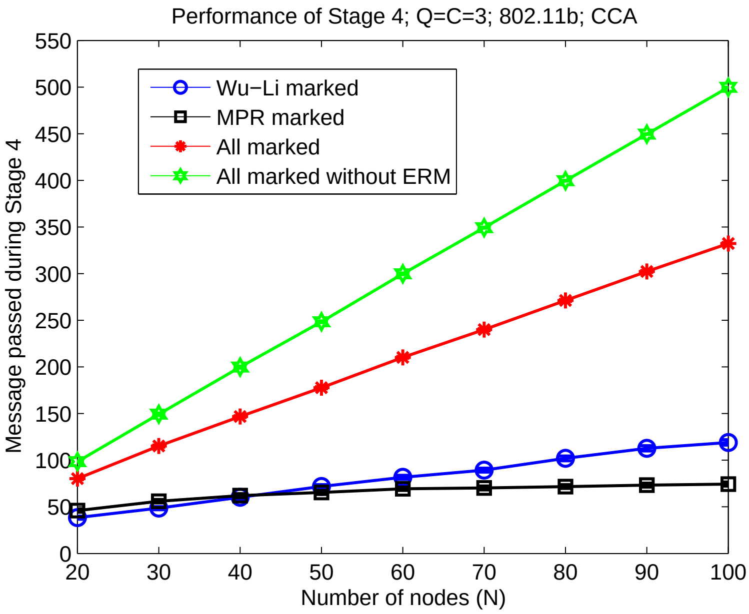

Figure 11: Messages exchanged (in Stage 3) for Stage 0 techniques with varying N

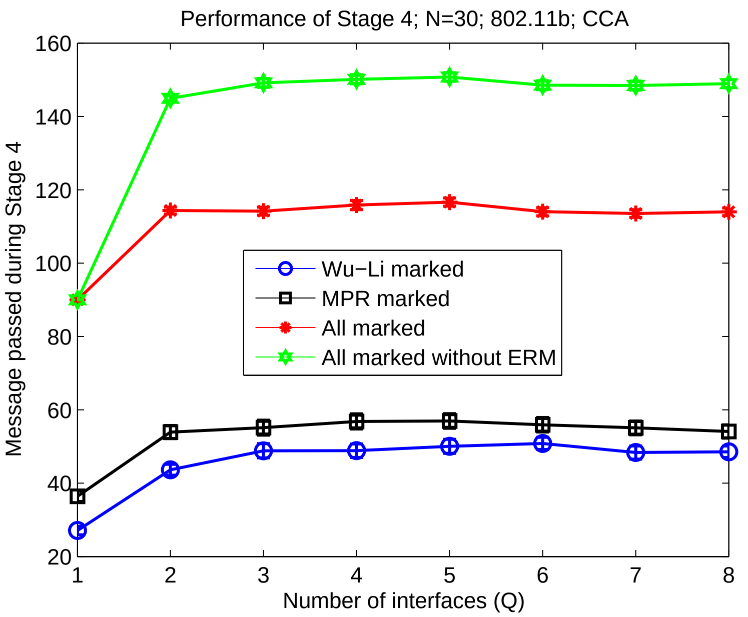

Figure 12: Messages exchanged (in Stage 3) for Stage 0 techniques with varying Q

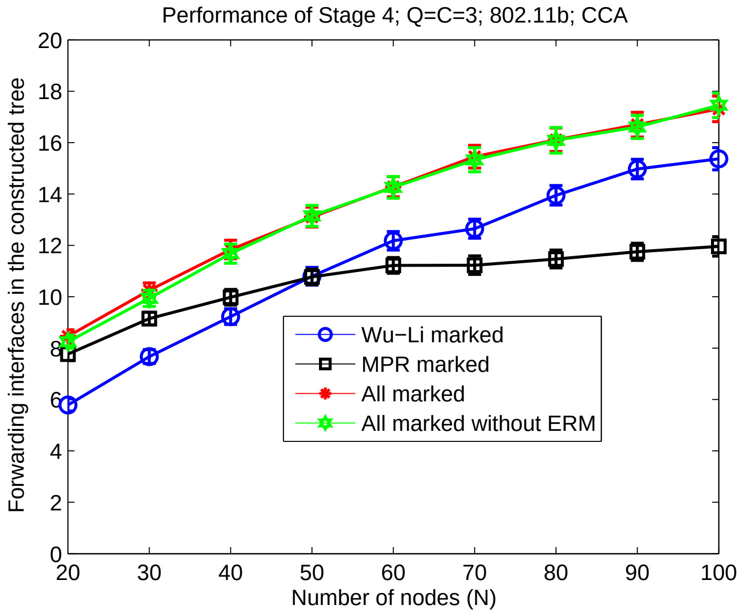

Figure 13: Forwarding interfaces in tree constructed by Stage 3

of messages exchanged during Tree Construction (Stage 3) in Figures 11 and 12 for varying number of N and Q. We see that fewer messages need to be exchanged when Wu-Li or MPR marking technique is used in Stage 0 as compared to when all nodes are marked. The benefit in overhead reduction due to ERM is evident in both Figures 11 and 12. We note that the overhead using all-node-marking scheme (especially without ERM) is particularly large. However, after the tree has been constructed by Stage 4, the forwarding interfaces are nearly the same regardless of the marking used in Stage 0 as shown in Figure 13.

5.5 Average latency vs. the 'broadcast latency'

We have compared the average-case and worst-case latency performance for MRDT and PAMT algorithms. This comparison is depicted in Figure 14 which displays the probability of nodes receiving the packet in given broadcast latency (in milliseconds). We see that the average-latency performance can also benefit from the increased parallelization of MRDT, as MRDT's performance is marginally better than the performance of PAMT.

5.6 Successful delivery probability results

The probability of successful delivery for our evaluated algorithms is shown in Figure 15. We see that MRDT's performance is marginally better than

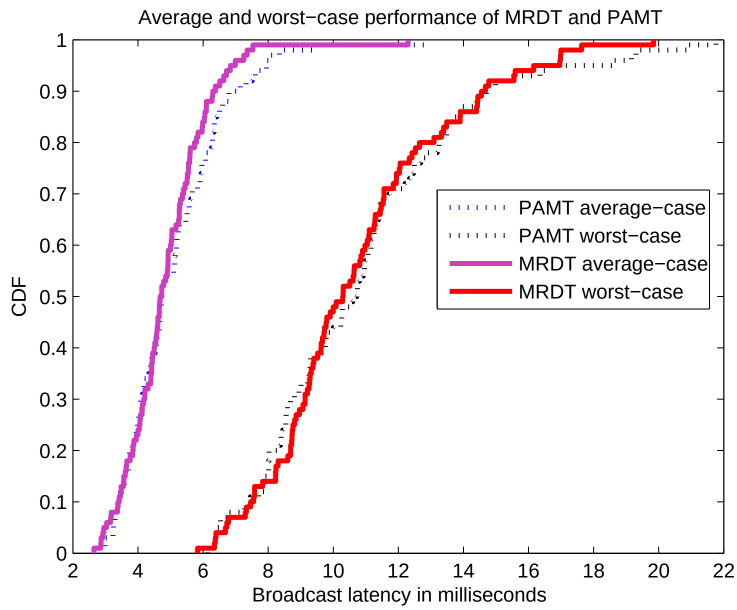

Figure 14: CDF of average and worst-case performance

the other algorithms. We note here that the delivery probability shown is for a single packet, i.e. for a single broadcasted packet, it shows the probability that all nodes in the network receive this packet.

5.7 The effect of channel-assignment

The performance of MRDT with different channel-assignment schemes is shown in Figure 16. We note that MRDT can benefit from increased 'connectivity' present in schemes like CCA which presents more opportunities of exploiting WBA. These results, however, are not specific to MRDT as PAMT presented similar features when different channels-assignment schemes were used [13].

5.8 The effect of transmission rate-range curve

For sensitivity analysis, we have performed our experiments in Qualnet for both 802.11*a* and 802.11*b* networks, both of which have different rate-range characteristics. Our results, Figures 5, 6, 7 and 8, seem to indicate that MRDT is relatively insensitive to the rate-range curve as it performs well in both 802.11*a* and 802.1*b* networks.

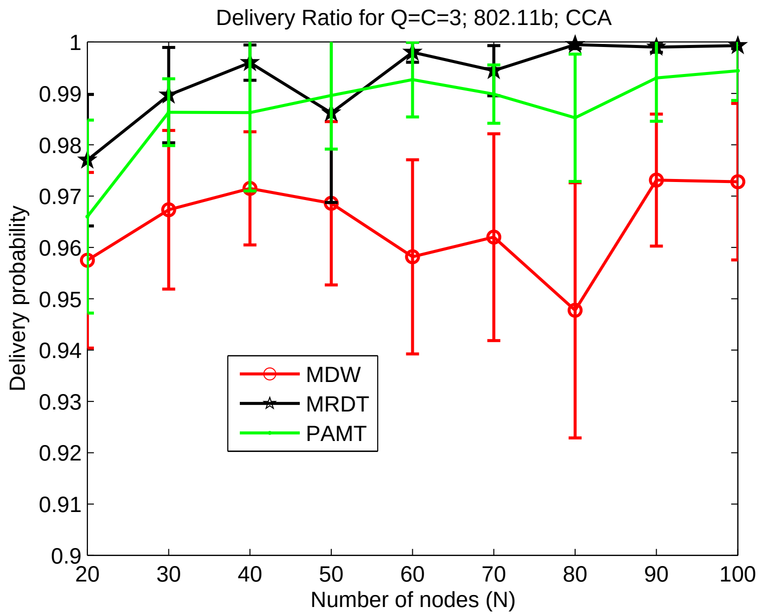

Figure 15: Probability of successful delivery to all nodes

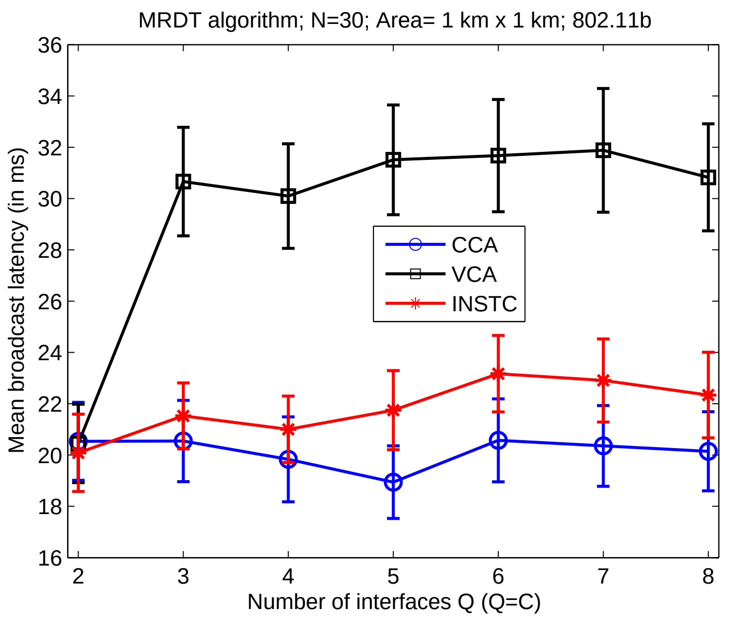

Figure 16: Effect of channel assignment

6 Conclusions and Future Work

In this paper, we have demonstrated how the link rate diversity and multi-radio architecture of individual mesh nodes deeply influence the efficiency of network-wide data broadcast in WMN environments. We have proposed *distributed and localized* broadcast heuristics that exploit "wireless broadcast advantage" and the rate, interface, channel diversity of a multi-radio multi-rate multi-channel (MR^2-MC) WMN. Our distributed algorithms match the performance of the centralized algorithms proposed in literature that require global information. We believe that work on efficient network-layer broadcasting in a MR^2-MC WMN is still at a very early stage and significant advances are needed to develop practical, distributed broadcast/multicast routing protocols that not only improve the overall network capacity but also prove robust in the face of dynamic link quality fluctuations that may be typical for many outdoor WMN deployments.

References

[1] C. Adjih, P. Jacquet, and L. Viennot. Computing connected dominated sets with multipoint relays. *Journal of Ad Hoc and Sensor Wireless Networks*, 1, 2005.

[2] I.F. Akyildiz, X. Wang, and W. Wang. Wireless Mesh Networks: a survey. *Computer Networks*, 47(4):445–487, 2005.

[3] J.C. Bicket. Bit-rate selection in wireless networks. *Master's thesis, Massachusetts Institute of Technology, 2005.*

[4] G. Călinescu, I.I. Măndoiu, P.J. Wan, and A.Z. Zelikovsky. Selecting forwarding neighbors in wireless ad hoc networks. *Mobile Networks and Applications, Kluwer Academic Publishers*, 9(2):101–111, 2004.

[5] C. T. Chou, A. Misra, and J. Qadir. Low latency broadcast in multi-rate wireless mesh networks. *IEEE JSAC Special Issue on Wireless Mesh Networks*, 2006.

[6] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *ACM MobiCom*, pages 114–128, 2004.

[7] R. Gandhi, S. Parthasarathy, and A. Mishra. Minimizing broadcast latency and redundancy in ad-hoc networks. In *ACM MobiHoc '03*.

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

[9] G. Holland, N. Vaidya, and P. Bahl. A rate-adaptive MAC protocol for multi-hop wireless networks. In *ACM MobiCom '01*, pages 236–251.

[10] P. Kyasanur and N.H. Vaidya. Routing and interface assignment in multi-channel multi-interface wireless networks. *IEEE WCNC*, Volume:4, 2005.

[11] S. J. Lee, W. Su, and M. Gerla. On-demand multicast routing protocol in multihop wireless mobile networks. *Mob. Netw. Appl.*, 7(6):441–453, 2002.

[12] U.T. Nguyen and X. Xiong. Rate-adaptive multicast in mobile ad-hoc networks. *IEEE WiMob*, pp. 352- 360 Vol. 3, 2005.

[13] J. Qadir, C.T. Chou, and A. Misra. Low latency broadcast in multi-radio multi-channel multi-rate wireless mesh networks. *Proceedings of IEEE SECON*, 2006.

[14] J. Qadir, C.T. Chou, A. Misra, and J.G. Lim. Localized minimum-latency broadcasting in multi-rate wireless mesh networks. In *IEEE WoWMoM, 2007.*

[15] J. Qadir, C.T. Chou, A. Misra, and J.G. Lim. Localized minimum-latency broadcasting in multi-rate wireless mesh networks. Technical Report TR, School of CSE, University of New South Wales., 2006.

[16] A. Qayyum, L. Viennot, and A. Laouiti. Multipoint relaying for flooding broadcast messages in mobile wireless networks. *System Sciences, 2002. HICSS.*, pages 3866–3875, 2002.

[17] A. Raniwala, K. Gopalan, and T. Chiueh. Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks. *Sigmobile MC2R*, 8(2):50–65, 2004.

[18] Scalable Networks Inc. `http://www.scalable-networks.com/`.

[19] I. Stojmenovic, M. Seddigh, and J. Zunic. Dominating sets and neighbor elimination-based broadcasting algorithms in wireless networks. *IEEE Transactions on Parallel and Distributed Systems*, 13(1):14–25, 2002.

[20] Z. K. Yang W. Q. Cheng T. Wang, B. Li. A minimized latency broadcast in multi-rate wireless mesh networks: distributed formulation and rate first algorithm. *IEEE ICME, Beijing, 2007.*

[21] J. Tang, G. Xue, and W. Zhang. Interference-aware topology control and QOS routing in multi-channel wireless mesh networks. In *ACM MobiHoc '05*, pages 68–77.

[22] P.J. Wan, K.M. Alzoubi, and O. Frieder. Distributed construction of connected dominating set in wireless ad hoc networks. *Mobile Networks and Applications*, 9(2):141–149, 2004.

[23] J.E. Wieselthier, G.D. Nguyen, and A. Ephremides. Energy-efficient broadcast and multicast trees in wireless networks. *Mobile Networks and Applications*, 7(6):481–492, 2002.

[24] S.H.Y. Wong, S. Lu, H. Yang, and V. Bharghavan. Robust rate adaptation for 802.11 wireless networks. *ACM MOBICOM 2006*.

[25] J. Wu and H. Li. On calculating connected dominating set for efficient routing in ad-hoc wireless networks. *Proceedings of the ACM DIALM '99*, pages 7–14, 1999.